

A MIDAS Approach to Modeling First and Second Moment Dynamics*

Daive Pettenuzzo
Brandeis University[†]

Allan Timmermann
UCSD, CEPR, and CREATES[‡]

Rossen Valkanov
UCSD[§]

April 24, 2015

Abstract

We propose a new approach to predictive density modeling that allows for MIDAS effects in both the first and second moments of the outcome. Specifically, our modeling approach allows for MIDAS stochastic volatility dynamics, generalizing a large literature focusing on MIDAS effects in the conditional mean, and allows the models to be estimated by means of standard Gibbs sampling methods. When applied to monthly time series on growth in industrial production and inflation, we find strong evidence that the introduction of MIDAS effects in the volatility equation leads to improved in-sample and out-of-sample density forecasts. Our results also suggest that model combination schemes assign high weight to MIDAS-in-volatility models and produce consistent gains in out-of-sample predictive performance.

Key words: MIDAS regressions; Bayesian estimation; stochastic volatility; out-of-sample forecasts; inflation forecasts, industrial production.

JEL classification: C53; C11; C32; E37

*We thank three anonymous referees, Francesco Ravazzolo, Elena Andreou, and Eric Ghysels (the Editor) for valuable comments and suggestions on an earlier version of the paper.

[†]Brandeis University, Sachar International Center, 415 South St, Waltham MA 02453, Tel: (781) 736-2834. Email: dpettenu@brandeis.edu

[‡]University of California, San Diego, 9500 Gilman Drive, MC 0553, La Jolla CA 92093. Tel: (858) 534-0894. Email: atimmerm@ucsd.edu.

[§]University of California, San Diego, 9500 Gilman Drive, MC 0553, La Jolla CA 92093. Tel: (858) 534-0898. Email: rvalkanov@ucsd.edu.

1 Introduction

The notion that financial variables measured at a high frequency (e.g., daily interest rates and stock returns) can be used to improve forecasts of less frequently observed monthly or quarterly macroeconomic variables is appealing and has generated considerable academic interest in a rapidly expanding literature on mixed-data sampling (MIDAS) models. MIDAS models aggregate data sampled at different frequencies in a manner that has the potential to improve the predictive accuracy of regression models. By using tightly parameterized lag polynomials that allow the weighting on current and past values of the predictors to be flexibly tailored to the data, the MIDAS approach makes it feasible to include information on long histories of variables observed at a higher frequency than the outcome variable of interest.¹

Empirical studies in the MIDAS literature have analyzed the dynamics in variables as diverse as GDP growth (Andreou et al. (2013), Carriero et al. (2015), Clements and Galvao (2008), Clements and Galvao (2009), Kuzin et al. (2011), Kuzin et al. (2013), Marcellino et al. (2013)), stock market volatility (Ghysels et al. (2007), Ghysels and Valkanov (2012)) and the relation between stock market volatility and macroeconomic activity (Engle et al. (2013) and Schorfheide et al. (2014)). Such studies typically use simple and compelling designs to introduce high frequency variables in the conditional mean equation and frequently find that the resulting point forecasts produce lower out-of-sample root mean square forecast errors (RMSFEs) than benchmarks ignoring high frequency information.

This paper introduces MIDAS-in-volatility effects using a Bayesian modeling approach that allows for stochastic volatility (SV) dynamics and thus treats the underlying volatility as an unobserved process. Previous studies either estimate MIDAS-in-mean models and allow for SV effects (but not MIDAS-in-volatility) or estimate a MIDAS model on an observable proxy for the volatility such as the realized variance. The premise of our approach is that there are good reasons to expect information in variables observed at a high frequency to be helpful in predicting the volatility of monthly or quarterly macroeconomic variables. Studies such as Sims and Zha (2006) and Stock and Watson (2002) document that the volatility of macroeconomic

¹A common alternative is to use an average of recent values, e.g., daily values within a quarter. However, this overlooks that recent observations carry information deemed more relevant than older observations. Alternatively, one can use only the most recent daily observation. However, this may be suboptimal, particularly in the presence of measurement errors.

variables varies over time. Moreover, accounting for dynamics in the volatility equation can lead to more efficient estimators and improved point forecasts as shown by [Clark \(2011\)](#), [Carriero et al. \(2012\)](#), [Carriero et al. \(2015\)](#), and [Clark and Ravazzolo \(2014\)](#).²

The Bayesian modeling approach offers several advantages in our setting. Notably, the objective of the Bayesian analysis is to obtain the predictive density given the data, as opposed to simply computing a point forecast. Such density forecasts can be used to evaluate a range of measures of predictive accuracy, including RMSFEs, log scores, and the continuously ranked probability score of [Gneiting and Raftery \(2007a\)](#). Measures of the accuracy of density forecasts are more likely to have power to detect improvements in volatility forecasts than measures based on point forecasts such as the RMSFE. Moreover, our forecasts account for parameter estimation error. This can be important in empirical applications with macroeconomic variables for which data samples are short and parameters tend to be imprecisely estimated. Finally, because we construct the predictive density for a range of different models, we can compute forecast combinations that optimally weigh the individual models. Such forecast combinations provide a way to deal with model uncertainty since they do not depend on identifying a single best model. As new data arrive, the combination weights get updated and models that start to perform better receive greater weight in the combinations.

Our paper exploits that the MIDAS lag polynomial can be cast as a linear regression model with transformed daily predictors. For models with constant volatility and normal innovations, Bayesian estimation can therefore be undertaken using a two-block Gibbs sampler. For models with stochastic volatility we propose a specification that adds a MIDAS term to the log conditional volatility equation. Conditional on the sequence of log-volatilities and the parameters determining the stochastic volatility dynamics, our MIDAS specification reduces to a standard linear regression model. In turn, to obtain the sequence of log-volatilities and the stochastic volatility parameters we rely on the algorithm of [Kim et al. \(1998\)](#), extended by [Chib et al.](#)

²[Carriero et al. \(2015\)](#) develop a Bayesian method for producing current-quarter forecasts of GDP growth that is closely related to the U-MIDAS approach proposed by [Forni et al. \(2013\)](#), and allow for both time-varying coefficients and stochastic volatility in the estimation. [Ghysels \(2012\)](#) extends the standard Bayesian VAR approach to allow for mixed frequency lags and MIDAS polynomials. [Marcellino et al. \(2013\)](#) develop a mixed frequency dynamic factor model featuring stochastic shifts in the volatility of both the latent common factor and the idiosyncratic components. [Rodriguez and Puggioni \(2010\)](#) cast a MIDAS regression model as a dynamic linear model, leaving unrestricted the coefficients on all the high frequency data lags.

(2002) to allow for exogenous covariates in the volatility equation.³ Hence a four-block Gibbs sampler can be used to produce posterior estimates for the model parameters. Our approach is straightforward to implement and well-suited for generating sequences of recursively updated density forecasts.

We illustrate our approach in empirical applications to U.S. inflation and growth in industrial production, both of which have been extensively studied in the literature. We use daily observations on eight predictor variables, including interest rates, stock returns and the business cycle measure proposed by [Aruoba et al. \(2009\)](#). Because we estimate both MIDAS-in-mean and MIDAS-in-volatility models, we can compare the contributions of high-frequency information towards predicting first and second moments.

We find few instances where the MIDAS-in-mean effects lead to systematic improvements in the point or density forecasts. In contrast, we find that MIDAS-in-volatility effects lead to significantly better density forecasts of inflation and industrial production growth both in-sample and out-of-sample. This holds even for benchmarks that are tough to beat such as factor models with stochastic volatility (non-MIDAS) dynamics. Our finding holds across multiple forecast horizons ranging from one through twelve months and across most of the daily predictor variables. Moreover, we find that model combination schemes produce better density forecasts than a range of benchmarks, suggesting that our results are robust to model uncertainty. Moreover, the [Geweke and Amisano \(2010\)](#) optimal prediction pool approach places at least 50 percent of the weight on MIDAS-in-volatility models, further underscoring their importance to out-of-sample performance in forecasting inflation and industrial production growth.

The outline of the paper is as follows. Section 2 introduces the MIDAS methodology and extends the model to include stochastic volatility effects and, as a new contribution, MIDAS terms in both the conditional mean and the log conditional volatility equation. Section 3 introduces our Bayesian estimation approach and discusses how to generate draws from the predictive density using Gibbs sampling methods. Section 4 describes our empirical applications, while Section 5 covers different forecast combination schemes. Section 6 concludes.

³Posterior simulation of the whole path of stochastic volatilities under an arbitrary second moment MIDAS lag polynomial would require the use of a particle filter.

2 MIDAS regression models

This section outlines how we generalize the conventional regression specification to account for MIDAS effects in the volatility equation, while also allowing for stochastic volatility.

2.1 MIDAS Setup

Suppose we are interested in forecasting some variable y_{t+1} which is observed only at discrete times $t-1, t, t+1$, etc., while data on a predictor variable, $x_t^{(m)}$, are observed m times between $t-1$ and t . For example, y_{t+1} could be a monthly variable and $x_t^{(m)}$ could be a daily variable. In this case $m = 22$, assuming that the number of daily observations available within a month is constant and equals 22.

Let $H \geq 1$ be an (arbitrary) forecast horizon and suppose we use the direct forecasting approach to generate multi-period forecasts by projecting the period $\tau + H$ outcome on information known at time t . It is natural to consider using lagged values of $x_t^{(m)}$ to forecast y_{t+1} . We denote such lags of $x_t^{(m)}$ by $x_{t-j/m}^{(m)}$, where the m superscript makes explicit the higher sampling frequency of $x_t^{(m)}$ relative to y_{t+1} . To include such lags we could use a simple MIDAS model⁴

$$y_{\tau+H} = \beta_0 + \mathcal{B}\left(L^{1/m}; \boldsymbol{\theta}\right) x_{\tau}^{(m)} + \varepsilon_{\tau+H}, \quad \tau = 1, \dots, t - H \quad (1)$$

where

$$\mathcal{B}\left(L^{1/m}; \boldsymbol{\theta}\right) = \sum_{k=0}^{K-1} B(k; \boldsymbol{\theta}) L^{k/m}.$$

$L^{k/m}$ is a lag operator such that $L^{1/m} x_{\tau}^{(m)} = x_{\tau-1/m}^{(m)}$, and $\varepsilon_{\tau+H}$ is i.i.d. with $E(\varepsilon_{\tau+H}) = 0$ and $Var(\varepsilon_{\tau+H}) = \sigma_{\varepsilon}^2$. K is the maximum lag length for the included predictors. The distinguishing feature of MIDAS models is that the lag coefficients in $B(k; \boldsymbol{\theta})$ are parameterized as a function of a low dimensional vector of parameters $\boldsymbol{\theta} = (\theta_0, \theta_1, \dots, \theta_p)$. To use a concrete example, suppose again that y_{t+1} is a monthly series which gets affected by twelve months' of daily data, $x_t^{(m)}$. In this case, we would need $K = 264$ (22×12) lags of daily variables. Without any restrictions on the parameters in $\mathcal{B}(L^{1/m}; \boldsymbol{\theta})$ there would be $264 + 2$ parameters to estimate in (1). By making $\mathcal{B}(L^{1/m}; \boldsymbol{\theta})$ a function of a small set of $p + 1 \ll K$ parameters we can greatly reduce the number of parameters to estimate.

⁴For simplicity, our notation suppresses the dependence of the parameters on the forecast horizon, H .

It is sometimes useful to cast the MIDAS model as

$$y_{\tau+H} = \beta_0 + \beta_1 \mathcal{B}_1 \left(L^{1/m}; \boldsymbol{\theta}_1 \right) x_{\tau}^{(m)} + \varepsilon_{\tau+H}, \quad \tau = 1, \dots, t - H \quad (2)$$

where $\beta_1 \mathcal{B}_1 \left(L^{1/m}; \boldsymbol{\theta}_1 \right) = \mathcal{B} \left(L^{1/m}; \boldsymbol{\theta} \right)$ and β_1 is a scalar that captures the overall impact of lagged values of $x_{\tau}^{(m)}$ on $y_{\tau+H}$. Since β_1 enters multiplicatively in (2), it cannot be identified without imposing further restrictions on the polynomial $\mathcal{B}_1 \left(L^{1/m}; \boldsymbol{\theta}_1 \right)$, e.g., by normalizing the function $\mathcal{B}_1 \left(L^{1/m}; \boldsymbol{\theta}_1 \right)$ to sum to unity.⁵

The model in (1) can be generalized to allow for p_y lags of y_t and another p_z lags of r predictor variables $\mathbf{z}_t = (z_{1t}, \dots, z_{rt})'$ measured at the same frequency as y_t :

$$y_{\tau+H} = \alpha + \sum_{j=0}^{p_y-1} \rho_{j+1} y_{\tau-j} + \sum_{j=0}^{p_z-1} \gamma'_{j+1} \mathbf{z}_{\tau-j} + \mathcal{B} \left(L^{1/m}; \boldsymbol{\theta} \right) x_{\tau}^{(m)} + \varepsilon_{\tau+H}. \quad (3)$$

This regression requires the estimation of $(2 + p + p_y + p_z \times r)$ coefficients. The distributed lag term $\sum_{j=0}^{p_y-1} \rho_{j+1} y_{\tau-j}$ captures same-frequency dynamics in y_{t+H} , while the addition of the \mathbf{z}_t factors allows for predictors other than own lags. We refer to the model in (3) as the Factor-augmented AutoRegressive MIDAS, or FAR-MIDAS, for short. If the lagged factors are excluded from equation (3), the model has only autoregressive and MIDAS elements and is called AR-MIDAS. These abbreviations reflect the nested structure of the models that we consider.⁶

2.2 MIDAS weighting functions

The functional form of the MIDAS weights $\mathcal{B} \left(L^{1/m}; \boldsymbol{\theta} \right)$ depends on the application at hand and has to be flexible enough to capture the dynamics in how the high frequency data $x_{\tau}^{(m)}$ affect the outcome. We adopt a simple unrestricted version of $\mathcal{B} \left(L^{1/m}; \boldsymbol{\theta} \right)$, known as the Almon lag polynomial, which takes the form

$$B(k; \boldsymbol{\theta}) = \sum_{i=0}^p \theta_i k^i, \quad (4)$$

where $\boldsymbol{\theta} = (\theta_0, \theta_1, \dots, \theta_p)$ is a vector featuring $p + 1$ parameters. Under this parameterization, (3) takes the form

$$y_{\tau+H} = \alpha + \sum_{j=0}^{p_y-1} \rho_{j+1} y_{\tau-j} + \sum_{j=0}^{p_z-1} \gamma'_{j+1} \mathbf{z}_{\tau-j} + \sum_{k=0}^{K-1} \sum_{i=0}^p \theta_i k^i L^{k/m} x_{\tau}^{(m)} + \varepsilon_{\tau+H}. \quad (5)$$

⁵Normalization and identification of β_1 are not strictly necessary in a MIDAS regression but can be useful in settings such as those of Ghysels et al. (2005) and Ghysels et al. (2007) where β_1 is important for economic interpretation of the results.

⁶The FAR-MIDAS model is called FADL-MIDAS (for factor augmented distributed lag MIDAS) in Andreou et al. (2013).

Define the $(p + 1 \times K)$ matrix \mathbf{Q}

$$\mathbf{Q} = \begin{bmatrix} 1 & 1 & 1 & \dots & 1 \\ 1 & 2 & 3 & \dots & K \\ 1 & 2^2 & 3^2 & \dots & K^2 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & 2^p & 3^p & \dots & K^p \end{bmatrix}, \quad (6)$$

and the $(K \times 1)$ vector of high frequency data lags $\mathbf{X}_\tau^{(m)}$

$$\mathbf{X}_\tau^{(m)} = \left[x_\tau^{(m)}, x_{\tau-1/m}^{(m)}, x_{\tau-2/m}^{(m)}, \dots, x_{\tau-1}^{(m)}, \dots, x_{\tau-(K-1)/m}^{(m)} \right]'. \quad (7)$$

Given the linearity of (4) and (5), we can rewrite (5) as

$$y_{\tau+H} = \alpha + \sum_{j=0}^{p_y-1} \rho_{j+1} y_{\tau-j} + \sum_{j=0}^{q_z-1} \gamma'_{j+1} z_{t-j} + \boldsymbol{\theta}' \widetilde{\mathbf{X}}_\tau^{(m)} + \varepsilon_{\tau+H}, \quad (8)$$

where $\widetilde{\mathbf{X}}_\tau^{(m)} = \mathbf{Q} \mathbf{X}_\tau^{(m)}$ is a $(p + 1 \times 1)$ vector of transformed daily regressors. Once estimates for the coefficients $\boldsymbol{\theta}$ are available, we can compute the MIDAS weights from (4) as $\hat{B}(k; \boldsymbol{\theta}) = \sum_{i=0}^p \hat{\theta}_i k^i$. We can also impose the restriction that the weights $\hat{B}(k; \boldsymbol{\theta})$ sum to one by normalizing them as

$$\tilde{B}(k; \boldsymbol{\theta}) = \frac{\hat{B}(k; \boldsymbol{\theta})}{\sum_{i=1}^K \hat{B}(i; \boldsymbol{\theta})}. \quad (9)$$

In forecasting applications, this normalization does not provide any advantages. Hence, we work with the unrestricted expression (8) for which the MIDAS parameters $\boldsymbol{\theta}$ can conveniently be estimated by OLS after transforming the daily regressors $\mathbf{X}_\tau^{(m)}$ into $\widetilde{\mathbf{X}}_\tau^{(m)}$.

It is useful to briefly contrast the Almon weights in (4) with other parameterizations in the MIDAS literature. These include the exponential Almon lag of (Ghysels et al. (2005), Andreou et al. (2013))

$$B(k; \boldsymbol{\theta}) = \frac{e^{\theta_1 k + \theta_2 k^2 + \dots + \theta_p k^p}}{\sum_{i=1}^K e^{\theta_1 i + \theta_2 i^2 + \dots + \theta_p i^p}},$$

and the normalized Beta function of Ghysels et al. (2007)

$$B(k; \boldsymbol{\theta}) = \frac{\left(\frac{k-1}{K-1}\right)^{\theta_1-1} \left(1 - \frac{k-1}{K-1}\right)^{\theta_2-1}}{\sum_{i=1}^K \left(\frac{i-1}{K-1}\right)^{\theta_1-1} \left(1 - \frac{i-1}{K-1}\right)^{\theta_2-1}}.$$

Estimation of MIDAS models with these parameterizations requires non-linear optimization.

A third alternative is the stepwise weights proposed in Ghysels et al. (2007) and Corsi (2009)

$$B(k; \boldsymbol{\theta}) = \theta_1 I_{k \in [a_0, a_1]} + \sum_{p=2}^p \theta_p I_{k \in [a_{p-1}, a_p]},$$

where the $p + 1$ parameters a_0, \dots, a_p are thresholds for the step function with $a_0 = 1 < a_1 < \dots < a_p = K$ and $I_{k \in [a_{p-1}, a_p]}$ is an indicator function, with

$$I_{k \in [a_{p-1}, a_p]} = \begin{cases} 1 & \text{if } a_{p-1} \leq k < a_p \\ 0 & \text{otherwise} \end{cases}.$$

Provided that the threshold values are known, estimation of MIDAS models with these weights can also be undertaken using OLS. A final alternative is the Unrestricted polynomial (U-MIDAS) approach proposed by Foroni et al. (2013). In this case, all the high frequency lag coefficients are left unconstrained, and estimation can be undertaken using OLS.

2.3 MIDAS in volatility

We next generalize the constant-volatility MIDAS models in the previous subsection to allow for time-varying volatility and MIDAS-in-volatility effects. This generalization is potentially important because it is well established that the use of high frequency variables leads to better in-sample fit and out-of-sample forecasting performance for many financial and macroeconomic variables; see Andersen et al. (2006), Ghysels et al. (2007), Engle et al. (2013), and Schorfheide et al. (2014).

We generalize the constant volatility model in two steps. First, consider extending the FAR-MIDAS model (8) to allow for stochastic volatility as

$$y_{\tau+H} = \alpha + \sum_{j=0}^{p_y-1} \rho_{j+1} y_{\tau-j} + \sum_{j=0}^{p_z-1} \gamma'_{j+1} z_{\tau-j} + \boldsymbol{\theta}' \widetilde{\mathbf{X}}_{\tau}^{(m)} + \exp(h_{\tau+H}) u_{\tau+H}, \quad (10)$$

where $h_{\tau+H}$ denotes the log-volatility of $y_{\tau+H}$ and $u_{\tau+H} \sim \mathcal{N}(0, 1)$. It is commonly assumed that the log-volatility evolves as a driftless random walk

$$h_{\tau+H} = h_{\tau} + \xi_{\tau+H}, \quad (11)$$

where $\xi_{\tau+H} \sim \mathcal{N}(0, \sigma_{\xi}^2)$ and u_t and ξ_s are mutually independent for all t and s . We refer to (10) and (11) as the FAR-MIDAS SV model. This type of specification is considered by Carriero

et al. (2015), and Marcellino et al. (2013), but with a different parameterization of the MIDAS weights.⁷

The SV model in (11) permits time varying volatility but does not allow high frequency variables, $v_\tau^{(m)}$, to affect the conditional log-volatility. To accomplish this, we generalize (11) to include second moment MIDAS effects

$$h_{\tau+H} = \lambda_0 + \lambda_1 h_\tau + \sum_{k=0}^{K-1} B_h(k; \boldsymbol{\theta}_h) L^{k/m} v_\tau^{(m)} + \xi_{\tau+H}. \quad (12)$$

The daily variables $v_\tau^{(m)}$ need not be the same as those entering the first moment in (10), $x_\tau^{(m)}$. Similarly, the MIDAS weights $B_h(k; \boldsymbol{\theta}_h)$ need not be the same as those in the conditional mean equation. The specification in (10) and (12) is a FAR-MIDAS with MIDAS stochastic volatility or FAR-MIDAS SV-MIDAS model. In addition to allowing the high frequency lags to enter the log-volatility equation, (12) relaxes the random walk assumption and introduces autoregressive dynamics for the log-volatility.⁸ The stochastic volatility MIDAS specification is an analogue to the MIDAS-in-mean specification, (3).

To complete the model in (12), we need to specify the SV-MIDAS weights, $B_h(k; \boldsymbol{\theta}_h)$, and the $v_\tau^{(m)}$ variables. We focus on specifications for which the first moment and second moment MIDAS variables are the same, $x_\tau^{(m)} = v_\tau^{(m)}$, and apply Almon lag polynomials for both the first and second moments. Under these assumptions, we can rewrite (12) as

$$h_{\tau+H} = \lambda_0 + \lambda_1 h_\tau + \boldsymbol{\theta}'_h \widetilde{\boldsymbol{X}}_\tau^{(m)} + \xi_{\tau+H}, \quad \xi_{\tau+H} \sim \mathcal{N}(0, \sigma_\xi^2). \quad (13)$$

We use the FAR-MIDAS SV-MIDAS model comprised of (10) and (13) in the estimation and forecasting sections.

3 Bayesian estimation and forecasting

This section explains how we use Bayesian methods to estimate the MIDAS forecasting models and generate density forecasts.

⁷The link between MIDAS models and time varying volatility has also been explored by Engle et al. (2013) who use a MIDAS-GARCH approach to link macroeconomic variables to the long-run component of volatility. Their model uses a mean reverting daily GARCH process and a MIDAS polynomial applied to monthly, quarterly, and biannual macroeconomic and financial variables.

⁸The addition of exogenous covariates in the log-volatility equation has been studied by Chib et al. (2002).

3.1 MIDAS models with constant volatility

Let Φ denote the regression parameters in the constant-volatility FAR-MIDAS model (3), excluding the MIDAS coefficients θ , i.e., $\Phi = (\alpha, \rho_1, \dots, \rho_{p_y}, \gamma'_1, \dots, \gamma'_{q_z})'$. Conditional on θ , the MIDAS model reduces to a standard linear regression and one only needs to draw from the posterior distributions of Φ and the variance of $\varepsilon_{\tau+H}$, σ_ε^2 . Assuming standard independent Normal-inverted gamma priors on Φ and σ_ε^2 and normally distributed residuals, $\varepsilon_{\tau+H}$, drawing from the posterior of these parameters is straightforward and simply requires using a two-block Gibbs sampler.

The same logic extends to estimation of the MIDAS parameters θ in cases where the transformed high frequency variables $\widetilde{\mathbf{X}}_t^{(m)}$ have a linear effect on the mean. Such cases include the (non-normalized) Almon lag polynomial specification in (8), the step function polynomial specification of Ghysels et al. (2007), and the U-MIDAS approach of Foroni et al. (2013). Suppose that $\varepsilon_{\tau+H}$ is normally distributed along with conjugate priors for the regression parameters and error variance. For such cases a two-block Gibbs sampler can be used to obtain posterior estimates for the parameters Φ , θ , and σ_ε^2 .⁹

To see how this works, rewrite (8) as

$$\begin{aligned} y_{\tau+H} &= \mathbf{Z}_\tau \Psi + \varepsilon_{\tau+H} \\ \tau &= 1, \dots, t-1 \end{aligned} \tag{14}$$

where $\Psi = (\Phi', \theta')'$ and $\mathbf{Z}_\tau = \left(1, y_\tau, \dots, y_{\tau-p_y+1}, z'_\tau, \dots, z'_{\tau-q_z+1}, \widetilde{\mathbf{X}}_\tau^{(m)'}\right)'$. Following standard practice, suppose that the priors for the regression parameters Ψ in (14) are normally distributed and independent of σ_ε^2

$$\Psi \sim N(\underline{\mathbf{b}}, \mathbf{V}). \tag{15}$$

All elements of $\underline{\mathbf{b}}$ are set to zero except for the value corresponding to ρ_1 which is set to one. Hence, our choice of the prior mean vector $\underline{\mathbf{b}}$ reflects the view that the best prediction model is

⁹Under the U-MIDAS approach of Foroni et al. (2013), the matrix of transformed regressors, $\widetilde{\mathbf{X}}_t^{(m)}$, is the same as the original matrix, $\mathbf{X}_t^{(m)}$.

a random walk. We choose a data-based prior for $\underline{\mathbf{V}}$:¹⁰

$$\underline{\mathbf{V}} = \underline{\psi}^2 \left[s_{y,t}^2 \left(\sum_{\tau=1}^{t-1} \mathbf{Z}'_{\tau} \mathbf{Z}_{\tau} \right)^{-1} \right], \quad (16)$$

with

$$s_{y,t}^2 = \frac{1}{t-2} \sum_{\tau=1}^{t-1} (y_{\tau+1} - y_{\tau})^2.$$

In (16), the scalar $\underline{\psi}$ controls the tightness of the prior. Letting $\underline{\psi} \rightarrow \infty$ produces a diffuse prior on $\underline{\Psi}$. Our analysis sets $\underline{\psi} = 25$, corresponding to relatively diffuse priors.

For the constant volatility model we assume a standard gamma prior for the error precision of the return innovation, $\sigma_{\varepsilon}^{-2}$:

$$\sigma_{\varepsilon}^{-2} \sim \mathcal{G}(s_{y,t}^{-2}, \underline{v}_0(t-1)). \quad (17)$$

\underline{v}_0 is a prior hyperparameter that controls how informative the prior is. $\underline{v}_0 \rightarrow 0$ corresponds to a diffuse prior on $\sigma_{\varepsilon}^{-2}$. Our baseline analysis sets $\underline{v}_0 = 0.005$. This corresponds to a pre-sample of half of one percent of the actual data sample, again representing an uninformative prior.

Let \mathcal{D}^t denote information available at time t . Obtaining draws from the joint posterior distribution $p(\underline{\Psi}, \sigma_{\varepsilon}^{-2} | \mathcal{D}^t)$ of the constant-volatility MIDAS regression model is now straightforward. Combine the priors in (15)-(17) with the observed data to get the conditional posteriors:

$$\underline{\Psi} | \sigma_{\varepsilon}^{-2}, \mathcal{D}^t \sim \mathcal{N}(\bar{\mathbf{b}}, \bar{\mathbf{V}}), \quad (18)$$

and

$$\sigma_{\varepsilon}^{-2} | \underline{\Psi}, \mathcal{D}^t \sim \mathcal{G}(\bar{s}^{-2}, \bar{v}), \quad (19)$$

where

$$\begin{aligned} \bar{\mathbf{V}} &= \left[\underline{\mathbf{V}}^{-1} + \sigma_{\varepsilon}^{-2} \sum_{\tau=1}^{t-1} \mathbf{Z}'_{\tau} \mathbf{Z}_{\tau} \right]^{-1}, \\ \bar{\mathbf{b}} &= \bar{\mathbf{V}} \left[\underline{\mathbf{V}}^{-1} \underline{\mathbf{b}} + \sigma_{\varepsilon}^{-2} \sum_{\tau=1}^{t-1} \mathbf{Z}'_{\tau} y_{\tau+1} \right], \end{aligned} \quad (20)$$

and

$$\bar{s}^{-2} = \frac{\sum_{\tau=1}^{t-1} (y_{\tau+1} - \mathbf{Z}_{\tau} \underline{\Psi})^2 + (s_{y,t}^2 \times \underline{v}_0(t-1))}{\bar{v}}, \quad (21)$$

¹⁰Priors for the hyperparameters are often based on sample estimates, see [Stock and Watson \(2006\)](#) and [Efron \(2010\)](#). Our analysis can be viewed as an empirical Bayes approach.

where $\bar{v} = \underline{v}_0 + (t - 1)$.

For more general MIDAS lag polynomials, obtaining posterior estimates for $\boldsymbol{\theta}$ is only slightly more involved and requires a straightforward modification of the Gibbs sampler algorithm outlined above. As an example, Ghysels (2012) focuses on the case of normalized beta weights where $\boldsymbol{\theta} = (\theta_1, \theta_2)'$, and suggests using a Gamma prior for both θ_1 and θ_2

$$\theta_j \sim \mathcal{G}(f_0, F_0), j = 1, 2. \quad (22)$$

Here f_0 and F_0 are hyperparameters controlling the mean and degrees of freedom of the Gamma distribution. Next, to draw from the posteriors of θ_1 and θ_2 , Ghysels proposes utilizing a Metropolis-in-Gibbs step as in Chib and Greenberg (1995). The Metropolis step is an accept-reject step that requires a candidate $\boldsymbol{\theta}^*$ from a proposal density $q(\boldsymbol{\theta}^* | \boldsymbol{\theta}^{[i]})$, where $\boldsymbol{\theta}^{[i]}$ is the last accepted draw for the MIDAS parameters $\boldsymbol{\theta}$. For example, when the Gamma distribution is chosen for $q(\boldsymbol{\theta}^* | \boldsymbol{\theta}^{[i]})$, at iteration $i + 1$ of the Gibbs sampler

$$\theta_j^* \sim \mathcal{G}\left(\theta_j^{[i]}, c\left(\theta_j^{[i]}\right)^2\right), j = 1, 2, \quad (23)$$

where c is a tuning parameter chosen to achieve a reasonable acceptance rate. The candidate draw gets selected with probability $\min\{a, 1\}$,

$$\boldsymbol{\theta}^{[i+1]} = \begin{cases} \boldsymbol{\theta}^* & \text{with probability } \min\{a, 1\} \\ \boldsymbol{\theta}^{[i]} & \text{with probability } 1 - \min\{a, 1\} \end{cases} \quad (24)$$

where a is computed as

$$a = \frac{L(\mathcal{D}^t | \boldsymbol{\Phi}, \boldsymbol{\theta}^*)}{L(\mathcal{D}^t | \boldsymbol{\Phi}, \boldsymbol{\theta}^{[i]})} \frac{\mathcal{G}(\boldsymbol{\theta}^* | f_0, F_0)}{\mathcal{G}(\boldsymbol{\theta}^{[i]} | f_0, F_0)} \frac{\mathcal{G}(\boldsymbol{\theta}^{[i]} | \boldsymbol{\theta}^*, c(\boldsymbol{\theta}^*)^2)}{\mathcal{G}(\boldsymbol{\theta}^* | \boldsymbol{\theta}^{[i]}, c(\boldsymbol{\theta}^{[i]})^2)}. \quad (25)$$

$L(\mathcal{D}^t | \boldsymbol{\Phi}, \boldsymbol{\theta}^*)$ and $L(\mathcal{D}^t | \boldsymbol{\Phi}, \boldsymbol{\theta}^{[i]})$ are the conditional likelihood functions given the parameters $\boldsymbol{\Phi}, \boldsymbol{\theta}^*$ and $\boldsymbol{\Phi}, \boldsymbol{\theta}^{[i]}$, respectively.

3.2 MIDAS models with time-varying volatility

Next, consider estimation of the models that allow the volatility of $\varepsilon_{\tau+1}$ to change over time, as in either (11) or (13). We focus our discussion on the most general process for the log-volatility, (13), and note that when working with (11), λ_0, λ_1 , and $\boldsymbol{\theta}_h$ drop out of the model. For the FAR-MIDAS SV-MIDAS model in equations (10) and (13), we require posterior estimates for all mean

parameters in equation (10), $\Psi = (\Phi', \theta')'$, the sequence of log volatilities, $\mathbf{h}^t = \{h_1, h_2, \dots, h_t\}$, the parameters $(\lambda_0, \lambda_1, \boldsymbol{\theta}_h)$, and the log-volatility variance σ_ξ^2 .

We follow the earlier choice of priors for the parameters in the mean equation Ψ

$$\Psi \sim N(\underline{\mathbf{b}}, \underline{\mathbf{V}}). \quad (26)$$

Turning to the sequence of log-volatilities, $\mathbf{h}^t = (h_1, \dots, h_t)$, the error precision, σ_ξ^{-2} , and the volatility parameters λ_0 , λ_1 , and $\boldsymbol{\theta}_h$, we can write

$$p(\mathbf{h}^t, \lambda_0, \lambda_1, \boldsymbol{\theta}_h, \sigma_\xi^{-2}) = p(\mathbf{h}^t | \lambda_0, \lambda_1, \boldsymbol{\theta}_h, \sigma_\xi^{-2}) p(\lambda_0, \lambda_1, \boldsymbol{\theta}_h) p(\sigma_\xi^{-2}).$$

Using (13), we can express $p(\mathbf{h}^t | \lambda_0, \lambda_1, \boldsymbol{\theta}_h, \sigma_\xi^{-2})$ as

$$p(\mathbf{h}^t | \lambda_0, \lambda_1, \boldsymbol{\theta}_h, \sigma_\xi^{-2}) = \prod_{\tau=1}^{t-1} p(h_{\tau+1} | \lambda_0, \lambda_1, \boldsymbol{\theta}_h, h_\tau, \sigma_\xi^{-2}) p(h_1), \quad (27)$$

with $h_{\tau+1} | \lambda_0, \lambda_1, \boldsymbol{\theta}_h, h_\tau, \sigma_\xi^{-2} \sim \mathcal{N}(\lambda_0 + \lambda_1 h_\tau + \boldsymbol{\theta}'_h \widetilde{\mathbf{X}}_\tau^{(m)}, \sigma_\xi^2)$. To complete the prior elicitation for $p(\mathbf{h}^t, \lambda_0, \lambda_1, \boldsymbol{\theta}_h, \sigma_\xi^{-2})$, we only need to specify priors for λ_0 , λ_1 , $\boldsymbol{\theta}_h$, the initial log-volatility, h_1 , and σ_ξ^{-2} . We choose these from the normal-inverted gamma family

$$h_1 \sim \mathcal{N}(\ln(s_{y,t}), \underline{k}_h), \quad (28)$$

$$\begin{bmatrix} \lambda_0 \\ \lambda_1 \\ \boldsymbol{\theta}_h \end{bmatrix} \sim \mathcal{N}(\underline{\mathbf{m}}_h, \underline{\mathbf{V}}_h), \quad (29)$$

and

$$\sigma_\xi^{-2} \sim \mathcal{G}(1/\underline{k}_\xi, \underline{v}_\xi(t-1)). \quad (30)$$

We set $\underline{k}_\xi = 0.01$, $\underline{v}_\xi = 1$, and $\underline{k}_h = 0.1$. These are more informative priors than our earlier choices. Setting $\underline{k}_\xi = 0.01$ and $\underline{v}_\xi = 1$ restricts changes to the log-volatility to be only 0.01 on average. Conversely, $\underline{k}_h = 0.1$ places a relatively diffuse prior on the initial log volatility state. We conduct a sensitivity analysis for these priors in a subsequent section.

Following Clark and Ravazzolo (2014) we set all the elements of the prior mean hyperparameter $\underline{\mathbf{m}}_h$ in (29) to zero, except for the parameter corresponding to the AR(1) coefficient λ_1 , which we set to 0.9. As for the prior variance hyperparameter $\underline{\mathbf{V}}_h$ in (29), we set it to an identity matrix with diagonal elements equal to 0.5², except for the element corresponding to

the AR(1) coefficient λ_1 , which we set to 0.01^2 . This corresponds to a mildly uninformative prior on the intercept and MIDAS coefficients, and a more informative prior on λ_1 , matching the persistent dynamics in the log volatility process.

To obtain posterior estimates for the mean parameters Ψ , the sequence of log volatilities \mathbf{h}^t , the stochastic volatility parameters $(\lambda_0, \lambda_1, \boldsymbol{\theta}_h)$, and the log-volatility variance σ_ξ^2 , we use a four-block Gibbs sampler to draw recursively from the following four conditional posterior distributions:

1. $p\left(\Psi \mid \mathbf{h}^t, \lambda_0, \lambda_1, \boldsymbol{\theta}_h, \sigma_\xi^{-2}, \mathcal{D}^t\right)$.
2. $p\left(\mathbf{h}^t \mid \Psi, \lambda_0, \lambda_1, \boldsymbol{\theta}_h, \sigma_\xi^{-2}, \mathcal{D}^t\right)$.
3. $p\left(\sigma_\xi^{-2} \mid \Psi, \mathbf{h}^t, \lambda_0, \lambda_1, \boldsymbol{\theta}_h, \mathcal{D}^t\right)$
4. $p\left(\lambda_0, \lambda_1, \boldsymbol{\theta}_h \mid \Psi, \mathbf{h}^t, \sigma_\xi^{-2}, \mathcal{D}^t\right)$

Simulating from the first three of these blocks is straightforward using the algorithms of [Kim et al. \(1998\)](#), extended by [Chib et al. \(2002\)](#) to allow for exogenous covariates in the volatility equation. The conditional posterior distribution of the SV-MIDAS parameters $p\left(\lambda_0, \lambda_1, \boldsymbol{\theta}_h \mid \Psi, \mathbf{h}^t, \sigma_\xi^{-2}, \mathcal{D}^t\right)$ in the fourth step can be expressed as

$$\lambda_0, \lambda_1, \boldsymbol{\theta}_h \mid \Psi, \mathbf{h}^t, \sigma_\xi^{-2}, \mathcal{D}^t \sim \mathcal{N}\left(\bar{\mathbf{m}}_h, \bar{\mathbf{V}}_h\right),$$

where

$$\bar{\mathbf{V}}_h = \left\{ \mathbf{V}_h^{-1} + \sigma_\xi^{-2} \sum_{\tau=1}^{t-1} \begin{bmatrix} 1 \\ h_\tau \\ \widetilde{\mathbf{X}}_\tau^{(m)} \end{bmatrix} \begin{bmatrix} 1, h_\tau, \widetilde{\mathbf{X}}_\tau^{(m)'} \end{bmatrix} \right\}^{-1}, \quad (31)$$

and

$$\bar{\mathbf{m}}_h = \bar{\mathbf{V}}_h \left\{ \mathbf{V}_h^{-1} \mathbf{m}_h + \sigma_\xi^{-2} \sum_{\tau=1}^{t-1} \begin{bmatrix} 1 \\ h_\tau \\ \widetilde{\mathbf{X}}_\tau^{(m)} \end{bmatrix} h_{\tau+1} \right\}. \quad (32)$$

3.3 Forecasts from MIDAS models

The objective of Bayesian estimation of MIDAS forecasting models is to obtain the predictive density for y_{t+1} . This density conditions only on the data and so accounts for parameter uncertainty. For example, working with the constant volatility MIDAS model (8), the predictive

density for y_{t+H} is given by

$$p(y_{t+H} | \mathcal{D}^t) = \int_{\Psi, \sigma_\varepsilon^{-2}} p(y_{t+H} | \Psi, \sigma_\varepsilon^{-2}, \mathcal{D}^t) p(\Psi, \sigma_\varepsilon^{-2} | \mathcal{D}^t) d\Psi d\sigma_\varepsilon^{-2}, \quad (33)$$

where $p(\Psi, \sigma_\varepsilon^{-2} | \mathcal{D}^t)$ denotes the joint posterior distribution of the MIDAS parameters conditional on information available at time t , \mathcal{D}^t .

Alternatively, when working with the FAR-MIDAS SV-MIDAS model in (10) and (13) the density forecast for y_{t+H} is given by

$$\begin{aligned} p(y_{t+H} | \mathcal{D}^t) &= \int_{\Psi, \mathbf{h}^{t+H}, \lambda_0, \lambda_1, \boldsymbol{\theta}_h, \sigma_\xi^{-2}} p(y_{t+H} | \Psi, h_{t+H}, \mathbf{h}^t, \lambda_0, \lambda_1, \boldsymbol{\theta}_h, \sigma_\xi^{-2}, \mathcal{D}^t) \\ &\quad \times p(h_{t+H} | \Psi, \mathbf{h}^t, \lambda_0, \lambda_1, \boldsymbol{\theta}_h, \sigma_\xi^{-2}, \mathcal{D}^t) \\ &\quad \times p(\Psi, \mathbf{h}^t, \lambda_0, \lambda_1, \boldsymbol{\theta}_h, \sigma_\xi^{-2} | \mathcal{D}^t) d\Psi d\mathbf{h}^{t+H} d\lambda_0 d\lambda_1 d\boldsymbol{\theta}_h d\sigma_\xi^{-2}. \end{aligned} \quad (34)$$

We can use the Gibbs sampler to draw from the predictive densities in (33) and (34). These draws, $y_{t+H|t}^{(j)}$, $j = 1, \dots, J$ can be used to compute objects such as point forecasts, $\hat{y}_{t+H|t} = J^{-1} \sum_{j=1}^J y_{t+H|t}^{(j)}$ or the quantile of the realized value of the predicted variable, $J^{-1} \sum_{j=1}^J I(y_{t+H} \leq y_{t+H|t}^{(j)})$, where $I(y_{t+H} \leq y_{t+H|t}^{(j)})$ is an indicator function that equals one if the outcome, y_{t+H} , falls below the j th draw from the Gibbs sampler.

3.4 Implementation of the Gibbs Sampler

We run the Gibbs samplers for 15,000 iterations, after a burn-in period of 1,000 iterations, thinning the chains by keeping one out of every three draws.¹¹ To evaluate convergence we compute the following diagnostics: (1) autocorrelation functions of the draws (the smaller the autocorrelations, the more efficient the samplers are); (2) inefficiency factors (IFs) for the posterior estimates of the parameters. The IF is the inverse of the relative numerical efficiency measure of Geweke (1992), i.e., the IF is an estimate of $(1 + 2 \sum_{k=1}^{\infty} \rho_k)$, where ρ_k is the k -th autocorrelation of the chain. In our application the estimate is performed using a 4% tapered window for the estimation of the spectral density at frequency zero. Values of the IFs below or around 20 are regarded as satisfactory; (3) Raftery and Lewis (1992) diagnostics on the total number of runs required to achieve a certain precision. The parameters for the diagnostics are specified as follows: quantile = 0.025; desired accuracy = 0.025; required probability of attaining

¹¹The calculations were performed on the High Performance Computing Cluster at Brandeis University.

the required accuracy = 0.95. The required number of runs should fall below the total number of iterations used.

We compute these diagnostics for all estimated models. To preserve space we only report results here for the FAR-MIDAS SV-MIDAS models, i.e., the most general of the models under consideration.¹² In all cases the 20th order sample autocorrelation of the draws is less than 0.06. Similarly, in most cases the Geweke (1992) IF of the posterior parameter estimates is around 3, and is never larger than 7. Finally, the Raftery and Lewis (1992) diagnostic on the total number of draws required is significantly below the total number of iterations used to estimate our models. In summary, all convergence diagnostics appear satisfactory.

4 Empirical Results

This section introduces our monthly data on U.S. growth in industrial production and inflation (our target variables), a set of macroeconomic factors, and the daily predictors. We then analyze the in-sample and out-of-sample predictive accuracy of our forecasts for the model specifications described in sections 2 and 3. We apply a range of measures to evaluate the predictive accuracy of our forecasts. As discussed above, one of the advantages of adopting a Bayesian framework is the ability to compute predictive distributions, rather than simple point forecasts, and to account for parameter uncertainty. Accordingly, to shed light on the predictive ability of the different models, we evaluate both point and density forecasts.

4.1 Data

Our empirical analysis uses monthly data on U.S. industrial production and the inflation rate. Specifically, Let I_τ denote the monthly seasonally adjusted Industrial Production Index (IPI) at time τ , obtained from the Federal Reserve of St. Louis FRED database, and define

$$y_\tau = 100 \times \ln(I_\tau/I_{\tau-1}). \quad (35)$$

Similarly, the monthly inflation rate is obtained as

$$y_\tau = 1200 \times \ln(P_\tau/P_{\tau-1}), \quad (36)$$

¹²These results are based on the specification that predicts the growth in industrial production for $H = 1$, using as daily predictor the federal funds rate. Similar results are obtained for other forecast horizons and daily predictors, and for either growth in industrial production or inflation.

where P_τ denotes the U.S. monthly seasonally adjusted Consumer Price Index for All Urban Consumers (All Items) at time τ , again downloaded from the Federal Reserve of St. Louis FRED database.

The monthly predictor variables are an updated version of the 132 macroeconomic series used in Ludvigson and Ng (2009) and extended by Jurado et al. (2014) to December 2011. The series are selected to represent broad categories of macroeconomic quantities such as real output and income, employment and hours, real retail, manufacturing and trade sales, consumer spending, housing starts, inventories and inventory sales ratios, orders and unfilled orders, compensation and labor costs, capacity utilization measures, price indexes, bond and stock market indexes, and foreign exchange measures.¹³ We follow Stock and Watson (2012) and Andreou et al. (2013) and extract two common factors (z_τ) from the 132 macroeconomic series using principal components.

We consider eight daily series, $x_\tau^{(m)}$, in this study: (i) the effective Federal Funds rate (Ffr), first-differenced to eliminate any trends; (ii) the interest rate spread between the 10-year government bond rate and the federal fund rate (Spr); (iii) value-weighted returns on US stocks (Ret); (iv) returns on the portfolio of small minus big stocks considered by Fama and French (1993) (Smb); (v) returns on the portfolio of high minus low book-to-market ratio stocks studied by Fama and French (1993) (Hml); (vi) returns on a winner minus loser momentum spread portfolio (Mom), (vii) the ADS daily business cycle variable of Aruoba et al. (2009); and (viii) the default spread measured as the difference in the yield on portfolios of BAA and AAA-rated corporate bonds.¹⁴ The interest rate series are from the Federal Reserve Bank of St. Louis database FRED. Value-weighted stock return data are obtained from CRSP and include dividends. Returns on the Smb, Hml, and Mom spread portfolios are downloaded from Kenneth French's data library.¹⁵ The ADS series are published by the Federal Reserve Bank of Philadelphia. Our data sample spans the period from 1962:01 to 2011:12.

To test whether a better in-sample fit and out-of-sample forecasts can be obtained by including the daily series in the forecasting model through MIDAS polynomials, we estimate several

¹³The data are available on Sydney Ludvigson's website at <http://www.econ.nyu.edu/user/ludvigsons/jlndata.zip>

¹⁴Using an international sample of data on ten countries, Liew and Vassalou (2000) find some evidence that Hml and Smb are helpful in predicting future GDP growth.

¹⁵We thank Kenneth French for making this data available on his website, at http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html

versions of the MIDAS specifications discussed above. These fall into one of three categories: (i) MIDAS-in-mean with constant volatility (8); (ii) MIDAS-in-mean with stochastic volatility (10 and 11); and (iii) MIDAS in both the mean and the volatility (10 and 13).

To explore the importance of MIDAS-in-mean effects, we compare AR-MIDAS and AR-MIDAS SV models to the corresponding non-MIDAS models. We do the same for the FAR-MIDAS and FAR-MIDAS SV models. Similarly, to assess the importance of MIDAS-in-volatility effects, we compare AR-MIDAS SV-MIDAS and FAR-MIDAS SV-MIDAS models to AR-MIDAS SV and FAR-MIDAS SV models, respectively. In total we consider six different MIDAS specifications.

Every MIDAS specification is estimated with one of the eight daily variables, Ffr, Spr, Ret, Smb, Hml, Mom, ADS, and Def, as defined above. Therefore, we have a total of 48 MIDAS forecasting models. In addition, we estimate four benchmark non-MIDAS models: a purely autoregressive model (AR); the same model with stochastic volatility (AR SV); a model that includes AR terms and factors (FAR); and a model with factors and stochastic volatility (FAR SV). In all cases our analysis assumes four AR lags of y ($p_y = 4$), two lags of the macro factors ($p_z = 2$) and twelve months of past daily observations.¹⁶

4.2 In-sample estimates and model comparisons

We first compare the fit of the different model specifications over the full sample, 1962:01-2011:12. In a Bayesian setting, a natural approach to model selection is to compute the Bayes factor, $B_{1,0}$, of the null model M_0 versus an alternative model, M_1 . The higher is the Bayes factor, the higher are the posterior odds in favor of M_1 against M_0 . We report two times the natural log of the Bayes factors, $2\ln(B_{1,0})$. To interpret the strength of the evidence, we follow studies such as Kass and Raftery (1995) and note that if $2\ln(B_{1,0})$ is below zero, the evidence supports M_0 over M_1 . For values of $2\ln(B_{1,0})$ between 0 and 2, there is “weak evidence” that M_1 is a more likely characterization of the data than M_0 . We view values of $2\ln(B_{1,0})$ between 2 and 6, 6 and 10, and higher than 10, as “some evidence,” “strong evidence,” and “very strong

¹⁶Our baseline results impose no restrictions on the coefficients of the four AR lags. We separately investigated the importance of this assumption by adding a stationarity restriction on the coefficients of the four AR lags, specifying a truncated normal prior for Ψ with support over the stationary region. We implemented this restriction by augmenting our Gibbs sampler with an accept/reject step that removes all explosive draws from the posterior of Ψ , and found the stationary restriction to have virtually no impact on the results.

evidence”, respectively, in support of M_1 versus the null, M_0 .

Table 1 shows pairwise model comparisons based on the transformed Bayes factors, $2 \ln(B_{1,0})$. First consider the results for the growth in industrial production shown in the left columns. Panels A and B display results for the AR-MIDAS and FAR-MIDAS models (8) relative to AR or FAR models, respectively, while Panels C and D reports the same statistics for models with SV dynamics in the second moment. Each of the MIDAS models is estimated by including a single daily predictor at a time, as shown in separate rows. At the shortest horizons ($H = 1, 3$ months) only the MIDAS models based on the Ret or ADS variables lead to strong improvements in forecasting performance relative to the corresponding benchmarks. These results carry over to the models that allow for stochastic volatility dynamics; by and large, ADS appears to be the only predictor that leads to better forecasts at the one-month horizon ($H = 1$).

Panels E and F in **Table 1** report Bayes factors for specifications with MIDAS effects in the volatility relative to non-MIDAS SV models, assuming that the mean already includes a MIDAS term. In other words, we test different versions of (10) and (13) versus the model implied by (10) and (11). We use the simpler SV volatility specification in our benchmark to reflect the popularity of this approach in empirical work; see Cogley et al. (2005), D’Agostino et al. (2013), Stock and Watson (2007), as well as the empirical evidence reported in Clark and Ravazzolo (2014). Also, empirical tests show that the RMSE performance generated by the two SV models in (13) and (11) is statistically indistinguishable.¹⁷

In general, for both the AR-MIDAS, SV-MIDAS and FAR-MIDAS, SV-MIDAS models the evidence strongly supports adding MIDAS effects to the volatility specification for the vast majority of comparisons. The evidence in favor of the IPI models with MIDAS in volatility grows stronger, the longer the forecast horizon.

For the inflation rate models shown in the right columns of **Table 1** the results mirror the IPI findings, namely, there is spotty evidence of improvement in the forecasting models due to the inclusion of MIDAS effects in the conditional mean (Panels A-D). In contrast, there is very strong evidence that the models that include MIDAS in the volatility equation lead to superior forecasts (Panels E-F). For the inflation series there is less evidence of systematic patterns in

¹⁷While the more general SV model (13) produces significantly better log probability scores than (11), this model is outperformed by the best SV-MIDAS models that include daily variables such as the ADS index.

the Bayes factors related to the forecast horizon compared with the data on growth in industrial production.

In summary, the (in-sample) evidence from the Bayes factors suggests that there are no gains in the performance of the prediction models from including MIDAS effects in the conditional mean. In contrast, we find strong evidence of improvements as a result of including MIDAS terms in the volatility equation. This holds across predictor variables and at both short and long forecast horizons.

4.2.1 Sensitivity to Choice of Priors

Readers may be concerned that the Bayes factors are sensitive to our choice of priors, particularly in the case of the volatility parameters which are chosen to be relatively informative. To address this point, we conduct a sensitivity analysis that highlights the robustness of our results by varying the key prior hyper-parameters, $\underline{\psi}$, \underline{v}_0 , and \underline{k}_ξ over a wide range. In particular, $\underline{\psi}$ varies between 100 and 2.5 (we set $\underline{\psi} = 25$ in our baseline results); \underline{v}_0 varies between 0.001 and 0.1 ($\underline{v}_0 = 0.005$ in our baseline results), and \underline{k}_ξ varies between 0.0001 and 1 ($\underline{k}_\xi = 0.01$ in our baseline results).

Table 2 reports Bayes factor results obtained under these alternative priors for the daily predictors Ffr, Ret, and ADS. The baseline results only change in meaningful ways when we change the prior on \underline{k}_ξ which significantly alters the dynamics in the volatility equation. For the other prior parameters, conclusions obtained under the other choices of priors are in line with our baseline results.

4.2.2 MIDAS weighting of daily predictors

An alternative, and more traditional, way to investigate the importance of the MIDAS weighting scheme applied to the daily predictor variables is to compare our MIDAS models with Almon polynomials to a set of alternative models featuring a simple time average of the high frequency data in either the mean and/or the volatility equation. The latter set of models are obtained by restricting the Q matrix in Equation (6) to only its first row. Thus, we can test the importance of the MIDAS weighting schemes in the mean and volatility by running an F-test on the coefficients

$(\theta_1, \theta_2, \dots, \theta_p)$ for the mean and $(\theta_{h,1}, \theta_{h,2}, \dots, \theta_{h,p})$ for the volatility.¹⁸ The results of such F-tests, reported in [Table 3](#), suggest that the simpler weighting scheme in many cases is rejected against the nesting MIDAS structure. Specifically, for the industrial production growth rate data we find that the interest rate spread, stock returns, ADS and the default spread variables require the more flexible MIDAS weighting schemes than the simple time average. While the evidence is somewhat weaker for the inflation rate data, we still find strong evidence that the more flexible weighting scheme is required for the models that include the interest rate spread, ADS, and default spread data.

4.3 Out-of-sample forecasts

We next turn our attention to out-of-sample forecasting performance. To generate these, we use the first twenty years of data as an initial training sample, i.e., we estimate our regression models over the period 1962:01-1981:12 and use the resulting estimates to predict the outcome in 1982:01. Next, we include 1982:01 in the estimation sample, which thus becomes 1962:01-1982:01, and use the corresponding estimates to predict the outcome in 1982:02. We proceed recursively in this fashion until the last observation in the sample, producing a time series of one-step-ahead forecasts spanning the time period from 1982:01 to 2011:12.

4.3.1 Point forecasts

First consider the performance of the point forecasts. For each of the MIDAS models we obtain point forecasts by repeatedly drawing from the predictive densities, $p(y_{\tau+H} | M_i, \mathcal{D}^\tau)$, and averaging across draws. We have added M_i in the conditioning argument of the predictive density to denote the specific model i , while τ ranges from 1981:12 to 2011:12- H . Following [Stock and Watson \(2003\)](#) and [Andreou et al. \(2013\)](#), we measure the predictive performance of the MIDAS models relative to a random walk (RW) model. Specifically, we summarize the precision of the point forecasts of model i , relative to that from the RW model, by means of the ratio of RMSFE values

$$RMSFE_i = \frac{\sqrt{\frac{1}{\bar{t}-t+1} \sum_{\tau=t}^{\bar{t}} e_{i,\tau}^2}}{\sqrt{\frac{1}{\bar{t}-t+1} \sum_{\tau=t}^{\bar{t}} e_{RW,\tau}^2}}, \quad (37)$$

¹⁸We use the means and variance-covariance of the posterior distributions for $(\theta_1, \theta_2, \dots, \theta_p)$ and $(\theta_{h,1}, \theta_{h,2}, \dots, \theta_{h,p})$ as inputs to these F tests.

where $e_{i,\tau}^2$ and $e_{RW,\tau}^2$ are the squared forecast errors at time τ generated by model i and the RW model, respectively. \underline{t} and \bar{t} denote the beginning and end of the forecast evaluation sample. Values less than one for $RMSFE_i$ indicate that model i produces more accurate point forecasts than the RW model.

The top panels in [Figure 1](#) plot the sequence of recursively generated point forecasts of the IPI growth rate generated by random walk, AR-MIDAS, AR-MIDAS SV, and AR-MIDAS SV-MIDAS models fitted to short ($H = 1$) and long ($H = 12$) horizons. The RW forecasts are quite different and notably more volatile than those from the three other models, particularly at the longest horizon. A similar pattern holds for the inflation rate data plotted in the lower panels, although the differences are attenuated for this series.

[Figure 2](#) plots volatility forecasts for the same set of models. These are again generated recursively and so are updated as new data arrives. This explains why the volatility forecast from the AR-MIDAS model displays a slight downward trend, reflecting the lower average volatility in the data following the start of the Great Moderation. The two SV models show considerable short-term variation in the conditional volatility. However, the dynamics are markedly different. For instance, the pure SV model generates more extreme volatility forecasts than the model that includes MIDAS effects in the volatility equation, notably during 2007-2009.

[Table 4](#) presents results for the RMSFE ratio in [\(37\)](#) using the IPI growth rate data. The first column (marked “No MIDAS”) shows the RMSFE ratio for the benchmark models that do not include MIDAS effects (i.e., AR, FAR, AR-SV, and FAR-SV models), while the subsequent columns show RMSFE values for the different MIDAS models, each of which includes a single daily predictor variable. We present results for different model specifications that add MIDAS effects to the mean (rows 1, 2, 3, and 5) or to the volatility (rows 4 and 6) and across forecast horizons ranging from $H = 1$ to $H = 12$ months.

An immediate observation from [Table 4](#) is that all RMSFE ratios are well below one, typically ranging from 0.70 to 0.80, thus suggesting that the models reduce the RMSFE by 20-30% relative to the random walk benchmark.¹⁹ At the short horizon ($H = 1$) 25% of this improvement comes from adopting an AR specification rather than imposing a unit root, while another

¹⁹The models’ RMSFE values are all significantly lower than that of the random walk benchmark so we do not report evidence of statistical significance in this table.

5% improvement comes from adding factors. Second, the results show that using the MIDAS approach to add daily stock returns (Ret) to the benchmark AR or FAR models results in small (up to 3%), but systematic reductions in the RMSFE for forecast horizons up to $H = 6$ months. At the shortest one-month horizon, MIDAS models based on the ADS reduce the RMSFE ratio quite substantially. However, consistent with our findings in [Table 1](#), these improvements seem limited to the shortest forecast horizon.

[Table 5](#) shows that reductions of 10-25% in the RMSFE of the random walk model are obtained for the inflation rate series. However, for this variable the evidence is much weaker that adding MIDAS to the mean and/or volatility equation results in more accurate out-of-sample point forecasts.²⁰

To test more formally if the MIDAS forecasts are more accurate than various competitors, [Table 6](#) reports Diebold-Mariano p -values under the null that a given model has the same predictive ability as the RW benchmark. Each case provides a pairwise comparison of a model with MIDAS-in-mean effects to the corresponding model without such MIDAS effects. Consistent with the results in [Table 4](#), only the ADS variable generates significant reductions in the MSFE values at the shortest one-month horizon for the IPI growth rate series (Panel A).²¹

4.3.2 Density forecasts

One limitation of the RMSFE values reported above is that they fail to capture the richness of the MIDAS models as they do not convey the full information in the density forecast $p(y_{t+1} | M_t, \mathcal{D}^t)$. Indeed, comparing the plots of the point forecasts and the volatility forecasts in [Figures 1](#) and [2](#), it is clear that there are much greater differences between the volatility forecasts generated by the different models. [Figure 3](#) shows that such differences give rise to very different density forecasts at two points in time. The left panels show the densities for 1994:01. Compared to AR-MIDAS models, SV dynamics compress the predictive density for this period. SV-MIDAS has a similar, but weaker, effect preserving some of the greater uncertainty associated with the constant volatility forecasts. The second snapshot, shown in the right panels in [Figure 3](#), occurs

²⁰Our empirical results are obtained using a limited set of daily predictor variables representing stock returns, interest rates and the daily business cycle index proposed by [Aruoba et al. \(2009\)](#). It is possible that further improvements could be obtained from MIDAS-in-mean effects by using a richer set of daily variables such as that proposed by [Andreou et al. \(2014\)](#).

²¹A comparison of MSE-values for the models with MIDAS-in-volatility effects to the models that omit such effects yielded very similar results and so are not reported here to preserve space.

in the middle of the global financial crisis (2008:12). Here the SV model generates a much higher conditional volatility than the constant-volatility benchmark, with the SV-MIDAS model again falling in the middle.

To measure the importance of such differences, we follow [Amisano and Giacomini \(2007\)](#), [Geweke and Amisano \(2010\)](#), and [Hall and Mitchell \(2007\)](#) and compute the average log-score differential (LSD)

$$LSD_i = \sum_{\tau=\underline{t}}^{\bar{t}} (LS_{i,\tau} - LS_{RW,\tau}), \quad (38)$$

where $LS_{i,\tau}$ ($LS_{RW,\tau}$) denotes the log-score of model i (RW), computed at time τ . Positive values of LSD_i indicate that model i produces more accurate density forecasts than the RW model and the larger the value, the greater the improvements.

[Table 7](#) and [Table 8](#) report results for the log-score measure for the IPI growth rate series ([Table 7](#)) and the inflation rate ([Table 8](#)). Once again the first column reports the value of the LSD for a given benchmark model that does not include MIDAS effects (i.e., AR, FAR, AR-SV, and FAR-SV models). Compared with the RMSFE measure, the LSD results identify more cases with improvements relative to the various benchmarks, particularly for the models that include MIDAS-in-volatility effects (rows 4 and 6 in each panel). This makes sense because the LSD measure reflects improved volatility forecasts whereas the RMSFE measure ignores such improvements.

To assess the statistical significance of these results, [Table 9](#) reports Diebold-Mariano p -values for tests of equal predictive ability using the log-scores. We report results only for the models that add MIDAS effects to the volatility since adding MIDAS to the mean does not result in significant improvements as we have seen. Adding MIDAS effects to the volatility equation produces significantly higher values of the log-score measure than the corresponding benchmark model without such effects. This holds for most of the predictive variables, across most forecast horizons, and for both the IPI growth rate and inflation rate series.

To see how the log score differential evolves over time, we compute the cumulative log score differential for model i versus the RW model

$$CLSD_{i,t} = \sum_{\tau=\underline{t}}^t (LS_{i,\tau} - LS_{RW,\tau}). \quad (39)$$

Positive and increasing values of $CLSD_{i,t}$ suggest that model i produces more accurate density forecasts than the RW model.

Figure 4 plots cumulative log score differentials for a benchmark AR SV model, an AR-MIDAS SV and an AR-MIDAS SV-MIDAS model, in each case selecting the MIDAS variable that is most informative. Top panels show results for the IPI growth rate while bottom panels show results for the inflation rate. At the 1-month horizon (left panel), the model that incorporates MIDAS-in-mean effects is notably better than the pure AR model. At the twelve-month horizon (right panel) there is less to distinguish between the MIDAS-in-mean and AR models, both of which are dominated by the SV-MIDAS model. For the inflation rate data shown in the bottom plots, there are no benefits from adding MIDAS-in-mean effects whereas adding MIDAS-in-volatility effects leads to better forecasting performance at both the one- and twelve-month horizons.

Following Gneiting and Raftery (2007b), Gneiting and Ranjan (2011) and Groen et al. (2013), we also compute the average continuously ranked probability score differential (CRPSD) of model i relative to the RW model²²

$$CRPSD_i = \frac{\frac{1}{\bar{t}-\underline{t}+1} \sum_{\tau=\underline{t}}^{\bar{t}} CRPS_{i,\tau}}{\frac{1}{\bar{t}-\underline{t}+1} \sum_{\tau=\underline{t}}^{\bar{t}} CRPS_{RW,\tau}}. \quad (40)$$

$CRPS_{i,\tau}$ ($CRPS_{RW,\tau}$) measures the average distance between the empirical cumulative distribution function (CDF) of y_τ (which is simply a step function in y_τ), and the empirical CDF associated with the predictive density of model i (RW). Values less than one for $CRPSD_i$ suggest that model i performs better than the benchmark RW model.

To save space we do not report the CRPSD results in separate tables but instead summarize our findings. For the IPI growth rate data the CRPSD results fall between those obtained for the RMSFE and LSD measures: for all MIDAS models based on the ADS variable we observe improvements in the forecasts at the shortest horizon ($H = 1$). In addition, we observe improvements across all horizons for the MIDAS-in-volatility specifications that include the ADS variable. Similar, if somewhat weaker, results hold for the MIDAS models based on daily aggregate stock returns (Ret). For the inflation rate series the CRPSD results are similar to

²²Gneiting and Raftery (2007b) explain how the CRPSD measure circumvents some of the problems of the logarithmic score, most notably the fact that the latter does not reward values from the predictive density that are close to, but different from, the realization.

the RMSFE ratio results reported in Table 5 and there is very little evidence of systematic improvements in forecasting performance due to adding MIDAS effects in the mean or volatility specifications.

Figures 5 and 6 illustrate how the ranking of IPI growth rate and inflation models varies across the RMSFE, LSD, and CRPSD criteria. Focusing on Figure 5, as the two figures contain a similar message, the left panel plots performance according to the RMSFE ratio on the horizontal axis against performance according to the log-score criterion on the vertical axis. The figure shows very clearly that the SV-MIDAS models perform better than the SV models which in turn outperform the linear models on the log score criterion. In contrast, there is not much to differentiate between the SV and SV-MIDAS models according to the RMSFE criterion. Thus, the benefits from using MIDAS models to incorporate daily information show up more strongly in the density forecast measures than in point forecasts. The right windows in Figure 5 show a close relation between the ranking by the RMSFE and CRPSD criteria within each class of models. Importantly, however, the CRPSD measure ranks the SV-MIDAS models better than the SV models which in turn perform better than the linear models, consistent with the ranking by the LSD criterion. Figure 6 conveys the same information for inflation models across the RMSFE, LSD, and CRPSD criteria.

4.3.3 Real time data

One concern about the factor augmented model is that many of the underlying macro variables are not observed in real time and are subject to data revisions. To address this point, we estimate models that replace the macroeconomic factors with the monthly NAPM PMI series which is available in real time. We find that the density forecasts implied by the resulting models are virtually indistinguishable from those we obtain from the baseline FAR models, thus suggesting that the role of data revisions in our setting is somewhat limited.²³

²³In particular, we find that for both the IPI growth and inflation rate series the RMSFE ratio obtained from comparing a FAR model to an AR model augmented with the NAPM PMI series is always statistically indistinguishable from one. Similarly, the log score differential between the same two sets of models is in almost all cases statistically indistinguishable from zero, with the exception of the inflation series for $H = 12$, for which we obtain an LSD value of 0.012 with a p -value of 0.030.

5 Forecast combinations

Our analysis covers multiple specifications that differ in terms of the identity of the daily MIDAS variables, inclusion of macro factors, as well as assumptions about volatility dynamics. In practice a forecaster will not know which, if any, model produces the best forecasts and so is confronted with model uncertainty. An attractive strategy in this situation is to combine forecasts from multiple models, rather than attempting to select a single best model or use a large model that nests all other specifications.

5.1 Forecast combination schemes

To see how model combination works in our setting, let M_i denote a specific model and suppose we have N different models with predictive densities $\{p(y_{t+1}|M_i, \mathcal{D}^t)\}_{i=1}^N$. Our starting point is the equal-weighted pool (EWP) which assigns equal weights to each model M_i

$$p(y_{t+1}|\mathcal{D}^t) = \frac{1}{N} \sum_{i=1}^N p(y_{t+1}|M_i, \mathcal{D}^t). \quad (41)$$

Next, consider the optimal predictive pool proposed by [Geweke and Amisano \(2011\)](#),

$$p(y_{t+1}|\mathcal{D}^t) = \sum_{i=1}^N w_{t,i}^* \times p(y_{t+1}|M_i, \mathcal{D}^t). \quad (42)$$

We use the past predictive performance of the N models to recursively determine the $(N \times 1)$ vector of weights $\mathbf{w}_t^* = [w_{t,1}^*, \dots, w_{t,N}^*]$. This requires determining \mathbf{w}_t^* by solving a maximization problem using only information available at time t ,

$$\mathbf{w}_t^* = \arg \max_{\mathbf{w}_t} \sum_{\tau=1}^{t-1} \log \left[\sum_{i=1}^N w_{\tau,i} \times S_{\tau+1,i} \right], \quad (43)$$

subject to $\mathbf{w}_t^* \in [0, 1]^N$. $S_{\tau+1,i} = \exp(LS_{\tau+1,i})$ is the exponential of the recursively computed log-score for model i at time $\tau + 1$. We follow [Geweke and Amisano \(2010\)](#) and compute the marginal likelihoods by cumulating the predictive log scores of each model, after conditioning on an initial warm-up estimation sample, as

$$\Pr(\mathcal{D}^t | M_i) = \exp \left(\sum_{\tau=\underline{t}}^t LS_{\tau,i} \right). \quad (44)$$

As $t \rightarrow \infty$ the weights chosen according to (43) minimize the Kullback-Leibler distance between the data generating process and the combined density. Finally, we consider Bayesian model averaging (BMA) which weighs the individual models by their posterior probabilities

$$p(y_{t+1} | \mathcal{D}^t) = \sum_{i=1}^N \Pr(M_i | \mathcal{D}^t) p(y_{t+1} | M_i, \mathcal{D}^t). \quad (45)$$

Here $\Pr(M_i | \mathcal{D}^t)$ denotes the posterior probability of model i , computed using all information available at time t ,

$$\Pr(M_i | \mathcal{D}^t) = \frac{\Pr(\mathcal{D}^t | M_i) \Pr(M_i)}{\sum_{j=1}^N \Pr(\mathcal{D}^t | M_j) \Pr(M_j)}. \quad (46)$$

$\Pr(\mathcal{D}^t | M_i)$ and $\Pr(M_i)$ are the marginal likelihood and prior probability for model i , respectively. We assume that all models are equally likely a priori and so set $\Pr(M_i) = 1/N$.

The idea of using forecast combination methods in MIDAS regressions is also explored by [Andreou et al. \(2013\)](#) in the context of predicting quarterly GDP growth by means of a host of daily financial variables. We note, however, that [Andreou et al. \(2013\)](#) combine point forecasts from the individual MIDAS models whereas our approach combines density forecasts.

5.2 Empirical results for forecast combinations

Panels A and B of [Table 10](#) report out-of-sample results for the RMSFE ratio and log-score differential using the three combination schemes discussed above and so are directly comparable to the values shown in [Tables 4-9](#). Panels C, D, E and F conduct Diebold-Mariano tests against the AR model (Panels C and D) or the AR-SV model (Panels E and F). Results for the IPI growth rate are shown in the top panels while the bottom panels show results for the inflation rate data.

For the IPI growth rate series the RMSFE ratios in Panel A are comparable to, or sometimes even better than, the best of the univariate models reported in [Table 4](#). A similar finding holds for the log-scores in Panel B which are comparable to the best models from [Table 7](#). At the shortest horizon ($H = 1$) the results are particularly strong for the BMA and optimal prediction pools, whereas the results tend to favor the equal-weighting scheme at horizons $H \geq 3$. This suggests that estimation errors in obtaining the combination weights matter more at the longer horizons.

Using the Diebold-Mariano statistic to formally test if the combined forecasts produce better out-of-sample results than specific benchmarks, we find that all three combination schemes generate significantly better performance at the one-month horizon than both an AR benchmark (Panels C and D) and an AR-SV benchmark (Panels E and F). The results are a little weaker at the longer horizons, although the equal-weighted combination continues to significantly outperform the two benchmarks across most of the horizons and for both performance metrics.

Quite similar results hold for the inflation rate data. Again we find that the model combinations generate significantly better forecasts than the random walk model. Moreover, all three combinations produce better forecasts than the AR benchmark, particularly when measured along the log-score metric. For horizons up to $H = 6$ months the BMA and optimal prediction pool produce significantly better forecasts than the AR-SV model as measured by the average log-score.

To gain insights into how adding MIDAS effects to the volatility specification helps improve the forecasts, [Figure 7](#) provides a recursive plot of the combination weights assigned to different classes of models. Except for the earliest part of the post-1982 sample, the AR-MIDAS and AR-MIDAS SV models receive very little weight. In the remainder of the sample the two models with SV-MIDAS dynamics receive most of the probability weight. The FAR-MIDAS SV-MIDAS model is particularly important for $H = 1$ whereas the AR-MIDAS SV-MIDAS receives significant weight for $H = 12$. The FAR-MIDAS SV model receives considerable weight around 2000 in the inflation rate forecasts, particularly at the one-month horizon. Looking at the figure from another perspective, the optimal weights assigned to the SV-MIDAS models is 50 percent or higher for the IPI growth rate and inflation at both horizons after an initial learning period. This underscores the superior out-of-sample predictive ability of these models.

Turning to the probability mass assigned to MIDAS models containing the eight different daily predictor variables, [Figure 8](#) shows that the ADS business cycle variable receives most of the weight in the short-run IPI growth rate forecasts (top panels), whereas the Ffr, Spr, Ret and Hml variables are more important at the 12-month horizon. Less clear-cut results are obtained for the inflation rate data (bottom panels).

A key observation from these plots is that the weighting of the different models fluctuates

considerably over time. [Del Negro et al. \(2014\)](#) develop a dynamic version of the prediction pool approach which they show works well for combinations of two models. Extending their approach to more than two models is computationally demanding so we stop short of implementing their approach empirically. However, it is worth noting that the evolution in the combination weights shown in [Figures 7 and 8](#) resembles the type of patterns found by [Del Negro et al. \(2014\)](#).

In summary, these results show clear advantages from using forecast combinations as a way to deal with model uncertainty. The superior predictive ability of SV-MIDAS models is manifest in large weights from an optimal prediction pool. Moreover, our results clearly show time variations in the optimal weights assigned to different MIDAS predictor variables and different types of model specifications. By recursively updating the model weights, the forecast combinations adaptively adjust to changes and produce better out-of-sample forecasts.

6 Conclusions

We develop a Bayesian approach to estimation of forecasting models that allows for MIDAS effects in both the first and second moments of the predicted series along with stochastic volatility dynamics. Our SV-MIDAS approach is easy to implement using conventional Gibbs sampling and generates predictive densities that only condition on the data available when the forecasts are made.

Empirical applications to monthly growth in industrial production and the monthly inflation rate in the U.S. show that the inclusion of MIDAS-in-volatility effects leads to significantly better in-sample and out-of-sample density forecasts compared with similar specifications without such MIDAS effects. We do not find similar improvements using root mean squared error measures based on the accuracy of point forecasts, nor do we find that adding MIDAS effects to the mean equation leads to systematically better forecasts. Importantly, our results hold across multiple forecast horizons stretching from one month through twelve months, and continue to hold even when a state-of-the-art model with factors in the mean and stochastic volatility dynamics is used as the benchmark.

Our results are obtained using daily data on individual variables. Some variables, notably market-wide stock returns and the business cycle index of [Aruoba et al. \(2009\)](#), work better than others when used in the MIDAS specification. This implies that model uncertainty could be an

issue. To handle this, we investigate a variety of model combination approaches. Empirically we find strong evidence that such combinations produce significantly better out-of-sample density forecasts.

Our empirical results open up a promising new venue for improving on macroeconomic and financial volatility forecasts using the MIDAS-in-volatility approach. Previous studies have found that it is difficult to improve out-of-sample volatility forecasts of variables such as stock returns relative to a model that simply uses lagged volatility as a predictor (e.g., [Paye, 2012](#)). However, such studies have not explored the combination of stochastic volatility dynamics with MIDAS-in-volatility effects as proposed here.

References

- Amisano, G. and R. Giacomini (2007). Comparing density forecasts via weighted likelihood ratio tests. *Journal of Business & Economic Statistics* 25(2), 177–190.
- Andersen, T. G., T. Bollerslev, P. F. Christoffersen, and F. X. Diebold (2006). Volatility and correlation forecasting. Volume 1 of *Handbook of Economic Forecasting*, pp. 777 – 878. Elsevier.
- Andreou, E., E. Ghysels, and A. Kourtellis (2013). Should macroeconomic forecasters use daily financial data and how? *Journal of Business & Economic Statistics* 31(2), 240–251.
- Aruoba, S. B., F. X. Diebold, and C. Scotti (2009). Real-time measurement of business conditions. *Journal of Business & Economic Statistics* 27(4), 417–427.
- Carriero, A., T. E. Clark, and M. Marcellino (2012). Common drifting volatility in large bayesian vars. CEPR Discussion Papers 8894.
- 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)*, n/a–n/a.
- Chib, S. and E. Greenberg (1995). Understanding the metropolis-hastings algorithm. *The American Statistician* 49(4), 327–335.

- Chib, S., F. Nardari, and N. Shephard (2002). Markov chain monte carlo methods for stochastic volatility models. *Journal of Econometrics* 108(2), 281 – 316.
- Clark, T. E. (2011). Real-time density forecasts from bayesian vector autoregressions with stochastic volatility. *Journal of Business & Economic Statistics* 29(3), 327–341.
- Clark, T. E. and F. Ravazzolo (2014). Macroeconomic forecasting performance under alternative specifications of time-varying volatility. *Journal of Applied Econometrics*, in press.
- Clements, M. P. and A. B. Galvao (2008). Macroeconomic forecasting with mixed-frequency data. *Journal of Business & Economic Statistics* 26(4), 546–554.
- Clements, M. P. and A. B. A. Galvao (2009). Forecasting us output growth using leading indicators: an appraisal using midas models. *Journal of Applied Econometrics* 24(7), 1187–1206.
- Cogley, T., S. Morozov, and T. J. Sargent (2005). Bayesian fan charts for u.k. inflation: Forecasting and sources of uncertainty in an evolving monetary system. *Journal of Economic Dynamics and Control* 29(11), 1893 – 1925.
- Corsi, F. (2009). A simple approximate long-memory model of realized volatility. *Journal of Financial Econometrics* 7(2), 174–196.
- 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., R. Hasegawa, and F. Schorfheide (2014). Dynamic prediction pools: an investigation of financial frictions and forecasting performance. Penn Institute for Economic Research Working Paper 14-034.
- Efron, B. (2010). *Large-Scale Inference: Empirical Bayes Methods for Estimation, Testing, and Prediction*. Cambridge University Press.
- Engle, R. F., E. Ghysels, and B. Sohn (2013). Stock market volatility and macroeconomic fundamentals. *The Review of Economics and Statistics* 95(3), 776–797.

- Froni, C., M. Marcellino, and C. Schumacher (2013). Unrestricted mixed data sampling (midas): Midas regressions with unrestricted lag polynomials. *Journal of the Royal Statistical Society: Series A (Statistics in Society)* 178(1), 57 – 82.
- Geweke, J. (1992). Evaluating the accuracy of sampling-based approaches to calculating posterior moments. In J. Bernardo, J. Berger, A. Dawid, and A. Smith (Eds.), *Bayesian Statistics 4*, pp. 169 – 193. Clarendon Press, Oxford, UK.
- Geweke, J. and G. Amisano (2010). Comparing and evaluating bayesian predictive distributions of asset returns. *International Journal of Forecasting* 26(2), 216 – 230.
- Geweke, J. and G. Amisano (2011). Optimal prediction pools. *Journal of Econometrics* 164(1), 130 – 141.
- Ghysels, E. (2012). Macroeconomics and the reality of mixed frequency data. Department of Economics, University of North Carolina Working Paper.
- Ghysels, E., P. Santa-Clara, and R. Valkanov (2005). There is a risk-return trade-off after all. *Journal of Financial Economics* 76(3), 509 – 548.
- Ghysels, E., A. Sinko, and R. Valkanov (2007). Midas regressions: Further results and new directions. *Econometric Reviews* 26(1), 53–90.
- Ghysels, E. and R. Valkanov (2012). Forecasting volatility with midas. In *Handbook of Volatility Models and Their Applications*, pp. 383–401. John Wiley & Sons, Inc.
- Gneiting, T. and A. E. Raftery (2007a). Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association* 102, 359–378.
- Gneiting, T. and A. E. Raftery (2007b). Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association* 102(477), 359–378.
- Gneiting, T. and R. Ranjan (2011). Comparing density forecasts using threshold- and quantile-weighted scoring rules. *Journal of Business & Economic Statistics* 29(3), 411–422.
- Groen, J., R. Paap, and F. Ravazzolo (2013). Real-Time Inflation Forecasting in a Changing World. *Journal of Business and Economic Statistics* 31(1), 29–44.

- Hall, S. G. and J. Mitchell (2007). Combining density forecasts. *International Journal of Forecasting* 23(1), 1 – 13.
- Jurado, K., S. C. Ludvigson, and S. Ng (2014). Measuring uncertainty. *American Economic Review*, forthcoming.
- Kass, R. E. and A. E. Raftery (1995). Bayes factors. *Journal of the American Statistical Association* 90(430), pp. 773–795.
- Kim, S., N. Shephard, and S. Chib (1998). Stochastic volatility: Likelihood inference and comparison with arch models. *The Review of Economic Studies* 65(3), 361–393.
- 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 – 542.
- Kuzin, V., M. Marcellino, and C. Schumacher (2013). Pooling versus model selection for nowcasting gdp with many predictors: Empirical evidence for six industrialized countries. *Journal of Applied Econometrics* 28(3), 392–411.
- Liew, J. and M. Vassalou (2000). Can book-to-market, size and momentum be risk factors that predict economic growth? *Journal of Financial Economics* 57(2), 221 – 245.
- Ludvigson, S. C. and S. Ng (2009). Macro factors in bond risk premia. *Review of Financial Studies* 22(12), 5027–5067.
- Marcellino, M. G., M. Porqueddu, and F. Venditti (2013). Short-term gdp forecasting with a mixed frequency dynamic factor model with stochastic volatility. Bank of Italy Temi di Discussione (Working Paper) No. 896.
- Paye, B. S. (2012). ‘Déjà vol’: Predictive regressions for aggregate stock market volatility using macroeconomic variables. *Journal of Financial Economics* 106(3), 527 – 546.
- Raftery, A. and S. Lewis (1992). How many iterations in the Gibbs sampler? In J. Bernardo, J. Berger, A. Dawid, and A. Smith (Eds.), *Bayesian Statistics 4*, pp. 763–773. Clarendon Press, Oxford, UK.

- Rodriguez, A. and G. Puggioni (2010). Mixed frequency models: Bayesian approaches to estimation and prediction. *International Journal of Forecasting* 26(2), 293 – 311.
- Schorfheide, F., D. Song, and A. Yaron (2014). Identifying long-run risks: A Bayesian Mixed-Frequency approach. Unpublished working paper, University of Pennsylvania, Philadelphia, PA.
- Sims, C. A. and T. Zha (2006). Were there regime switches in u.s. monetary policy? *American Economic Review* 96(1), 54–81.
- Stock, J. H. and M. W. Watson (2002). Has the business cycle changed and why? In M. Gertler and K. Rogoff (Eds.), *NBER Macroeconomics Annual*, Volume 17. MIT Press.
- Stock, J. H. and M. W. Watson (2003). Forecasting output and inflation: The role of asset prices. *Journal of Economic Literature* 41(3), 788–829.
- Stock, J. H. and M. W. Watson (2006). Forecasting with many predictors. Volume 1 of *Handbook of Economic Forecasting*, pp. 515 – 554. Elsevier.
- Stock, J. H. and M. W. Watson (2007). Why has u.s. inflation become harder to forecast? *Journal of Money, Credit and Banking* 39, 3–33.
- Stock, J. H. and M. W. Watson (2012). Generalized shrinkage methods for forecasting using many predictors. *Journal of Business and Economic Statistics* 30(4), 481–493.

Table 1. Bayes factors

<i>Panel A: AR-MIDAS vs. AR</i>										
Daily series	<i>IPI growth rate</i>					<i>Inflation rate</i>				
	<i>H = 1</i>	<i>H = 3</i>	<i>H = 6</i>	<i>H = 9</i>	<i>H = 12</i>	<i>H = 1</i>	<i>H = 3</i>	<i>H = 6</i>	<i>H = 9</i>	<i>H = 12</i>
Ffr	5.035	0.702	-1.227	-1.527	-2.085	12.566	2.534	-1.453	-9.971	5.990
Spr	-6.053	-26.141	-17.451	-8.096	3.646	11.873	7.453	-6.287	-13.761	-2.861
Ret	20.478	28.045	7.986	2.773	-1.942	-3.607	-5.921	-4.744	-5.305	-4.179
Smb	4.966	-0.061	-4.637	-5.252	-4.767	-1.058	-1.250	-4.831	-5.519	-6.226
Hml	-3.231	-4.199	-4.267	-2.500	-0.393	-5.216	-1.999	-1.706	-3.561	-7.859
Mom	-0.705	7.688	-5.670	-11.231	-16.720	-1.653	-5.624	-6.606	-13.098	-6.170
ADS	85.414	2.114	4.607	3.411	-2.384	-1.945	0.243	-6.489	-5.122	-3.329
Def	6.435	-3.377	-6.891	-9.013	-15.840	-3.451	-8.375	-14.600	-17.409	-13.951
<i>Panel B: FAR-MIDAS vs. FAR</i>										
Ffr	-2.099	-6.896	-1.572	-1.103	-4.457	-1.999	-0.852	-1.987	-7.694	-4.891
Spr	-0.961	-7.869	-12.290	-12.635	-12.137	-1.907	1.029	-3.540	-8.521	-2.466
Ret	3.747	7.341	1.402	-1.401	-15.852	0.324	-2.677	-6.320	-7.112	-5.014
Smb	-2.219	-5.505	-3.567	-5.353	-7.285	2.169	-1.577	-7.517	-6.575	-7.477
Hml	-3.477	-3.753	-5.265	-4.319	-2.975	-4.755	-1.518	-2.118	-2.393	-3.471
Mom	-3.274	-3.776	-11.284	-9.386	-9.294	-3.477	-3.570	-6.756	-9.167	-5.447
ADS	34.634	-2.313	2.936	-0.631	-6.367	2.265	15.480	-10.444	-13.588	-11.300
Def	0.993	-5.984	-8.698	-5.164	-8.037	-1.722	5.834	-10.001	-14.275	-10.424
<i>Panel C: AR-MIDAS SV vs. AR SV</i>										
Ffr	-4.202	-9.050	-19.678	-9.468	-28.776	-3.492	-10.224	-2.380	-3.028	-1.154
Spr	-19.166	-40.150	-28.134	-25.714	-6.167	-2.287	-4.322	-14.816	-39.427	-26.722
Ret	9.271	-2.800	-2.844	-13.269	-9.567	-2.388	-5.919	-28.664	-1.237	-7.630
Smb	1.561	-8.388	-12.046	-16.369	-17.933	-0.570	1.588	-2.780	4.508	-5.152
Hml	-9.638	-15.631	-31.106	-27.338	-22.168	-5.448	-5.963	-10.014	-2.291	-10.410
Mom	-4.287	-9.398	-39.450	-36.212	-30.285	-3.941	-1.296	-22.837	4.467	-2.078
ADS	66.997	4.527	9.210	6.420	-7.445	-2.960	-7.430	-18.625	-9.220	-18.920
Def	1.671	-11.992	-33.778	-20.861	-18.997	4.234	-2.469	-0.468	-3.536	-1.093
<i>Panel D: FAR-MIDAS SV vs. FAR SV</i>										
Ffr	-3.672	-8.625	-12.921	-3.371	-24.363	-7.179	-8.728	-2.815	12.864	-3.761
Spr	1.160	-11.630	-10.596	-25.894	-12.209	-3.641	-7.215	-15.783	4.471	-10.397
Ret	-0.930	-1.142	-7.295	-7.828	-25.213	-5.080	-12.999	-29.673	6.462	-15.077
Smb	-5.187	-13.988	-9.639	-9.327	-17.292	0.170	-1.993	10.936	14.268	-8.200
Hml	-6.695	-15.436	-26.916	-25.179	-24.520	-7.328	-8.311	-3.893	5.241	-8.380
Mom	-8.668	-15.073	-28.388	-20.395	-28.185	-2.634	-3.499	-30.219	14.843	-1.001
ADS	38.654	-11.055	0.162	3.597	-8.821	4.567	3.353	-7.000	-4.409	-13.625
Def	-4.720	-16.080	-33.443	-17.634	-11.370	10.607	6.908	3.056	4.640	-0.812
<i>Panel E: AR-MIDAS SV-MIDAS vs. AR-MIDAS SV</i>										
Ffr	9.962	19.527	38.172	41.127	59.056	21.726	31.325	61.524	27.415	27.252
Spr	0.553	-3.654	25.815	8.516	66.973	29.973	50.073	55.670	2.331	-6.704
Ret	27.660	38.825	49.499	35.354	69.606	19.531	36.202	54.915	30.917	21.801
Smb	19.832	25.954	43.312	30.447	62.111	20.785	36.453	60.542	34.179	23.220
Hml	7.954	18.359	26.601	26.642	67.721	9.155	36.214	59.045	33.045	18.866
Mom	13.922	31.694	23.702	13.589	43.405	12.493	39.514	47.950	28.720	27.367
ADS	79.708	36.043	56.360	44.895	70.462	15.739	28.742	43.332	14.847	5.825
Def	15.107	23.652	16.992	16.378	45.410	20.792	28.700	46.399	16.192	0.757
<i>Panel F: FAR-MIDAS SV-MIDAS vs. FAR-MIDAS SV</i>										
Ffr	22.758	26.663	29.059	43.199	57.174	26.241	34.665	67.709	30.086	29.776
Spr	30.299	32.302	30.451	14.262	58.036	33.900	41.432	62.925	13.746	16.824
Ret	27.510	42.061	33.666	35.622	46.425	31.173	32.218	53.554	25.713	22.762
Smb	24.115	28.187	32.785	32.607	60.811	33.460	33.349	66.882	34.285	27.958
Hml	20.403	24.142	22.600	27.687	59.379	19.027	31.942	64.711	24.572	21.799
Mom	22.759	29.314	11.173	18.235	49.743	24.889	39.622	52.387	23.861	32.933
ADS	57.213	33.414	39.822	41.413	64.113	33.828	44.358	46.618	15.392	18.570
Def	23.176	18.889	-8.501	25.032	57.831	37.629	47.702	59.418	24.439	26.984

This table reports pairwise model comparisons using twice the natural logarithm of the Bayes factor, $2 \times (\ln B_{1,0})$, where $B_{1,0}$ denotes the Bayes factor obtained from comparing model M_1 vs. M_0 $B_{1,0} = Pr(\mathcal{D}^t | M_1) / Pr(\mathcal{D}^t | M_0)$. Pairwise model comparisons are listed in the first column, where the notation ‘AR’ refers to an autoregressive model, ‘AR-MIDAS’ refers to an augmented distributed lag MIDAS model, ‘FAR’ refers to a factor augmented autoregressive model, and ‘FAR-MIDAS’ refers to a factor augmented distributed lag MIDAS model. The suffixes ‘SV’ and ‘SV-MIDAS’ denote models with stochastic volatility and MIDAS stochastic volatility, respectively. Column headers denote the daily predictor used in the MIDAS models, namely: the effective Federal Funds rate (Ffr), the interest rate spread between the 10-year government bond rate and the federal funds rate (Spr), value-weighted stock returns (Ret), the SML portfolio return (Smb), the HML portfolio return (Hml), the MOM portfolio return (Mom), the business cycle variable of [Aruoba et al. \(2009\)](#) (ADS), and the default spread between BAA and AAA rated bonds. Kass and Raftery (1995) suggest interpreting the results as follows: For $2 \times (\ln B_{1,0})$ less than 0, the evidence favors M_0 over M_1 ; for $2 \times (\ln B_{1,0})$ 0-2, 2-6, 6-10, and higher than 10, there is “weak evidence”, “some evidence”, “strong evidence”, and “very strong evidence” in favor of M_1 relative to M_0 . Bold numbers denote weak or stronger evidence in favor of M_1 .

Table 2. Bayes factors, alternative prior choices

AR-MIDAS SV vs. AR SV									
Daily series	(a) $\underline{\psi} = 100, \underline{v}_0 = 0.001, \underline{k}_\xi = 0.1$				(b) $\underline{\psi} = 2.5, \underline{v}_0 = 0.1, \underline{k}_\xi = 0.1$				
	IPI growth rate		Inflation rate		IPI growth rate		Inflation rate		
	H = 1	H = 12	H = 1	H = 12	H = 1	H = 12	H = 1	H = 12	
Ffr	-0.819	-29.199	-0.335	-20.352	-4.692	-2.073	2.811	-0.043	
Ret	10.779	-11.589	14.137	-5.611	-2.725	-4.973	-0.842	0.319	
ADS	69.926	-7.525	69.888	-2.774	-5.395	-26.804	0.897	2.450	
Daily series	(c) $\underline{\psi} = 25, \underline{v}_0 = 0.005, \underline{k}_\xi = 0.01$				(d) $\underline{\psi} = 25, \underline{v}_0 = 0.005, \underline{k}_\xi = 1$				
	Ffr	2.120	-9.915	4.528	-14.575	10.306	1.964	8.730	-10.657
	Ret	24.437	-0.700	26.182	4.071	2.975	-1.891	-0.574	-10.838
ADS	85.006	3.966	68.382	3.443	-0.186	-20.695	-5.031	-32.238	
FAR-MIDAS SV vs. FAR SV									
Daily series	(a) $\underline{\psi} = 100, \underline{v}_0 = 0.001, \underline{k}_\xi = 0.1$				(b) $\underline{\psi} = 2.5, \underline{v}_0 = 0.1, \underline{k}_\xi = 0.1$				
	Ffr	1.726	-23.044	-4.766	-20.588	-6.127	-8.232	1.072	-4.551
	Ret	3.879	-27.529	5.099	-21.932	-2.543	-18.765	-2.188	-1.835
ADS	42.650	-10.782	43.305	-9.361	2.625	-19.495	1.057	-1.304	
Daily series	(c) $\underline{\psi} = 25, \underline{v}_0 = 0.005, \underline{k}_\xi = 0.01$				(d) $\underline{\psi} = 25, \underline{v}_0 = 0.005, \underline{k}_\xi = 1$				
	Ffr	1.820	-6.595	-2.997	-7.441	-2.025	-6.236	-1.054	-4.355
	Ret	5.916	-19.761	17.027	-41.753	-0.607	-8.643	0.934	-24.196
ADS	38.818	-8.247	38.300	-12.530	5.672	-21.775	7.647	-6.888	
AR-MIDAS SV-MIDAS vs. AR-MIDAS SV									
Daily series	(a) $\underline{\psi} = 100, \underline{v}_0 = 0.001, \underline{k}_\xi = 0.1$				(b) $\underline{\psi} = 2.5, \underline{v}_0 = 0.1, \underline{k}_\xi = 0.1$				
	Ffr	8.316	60.694	6.066	53.630	20.346	22.493	10.200	14.390
	Ret	24.646	67.923	25.115	59.494	22.267	16.764	9.643	13.185
ADS	78.306	68.327	80.780	59.034	17.677	-6.731	10.482	10.441	
Daily series	(c) $\underline{\psi} = 25, \underline{v}_0 = 0.005, \underline{k}_\xi = 0.01$				(d) $\underline{\psi} = 25, \underline{v}_0 = 0.005, \underline{k}_\xi = 1$				
	Ffr	-27.646	-12.119	-47.857	12.815	1.542	1.919	29.526	-439.053
	Ret	-7.881	-10.743	38.489	41.670	-10.166	-6.641	15.838	15.640
ADS	54.511	-6.395	81.267	55.592	-10.243	-19.511	23.591	-6.302	
FAR-MIDAS SV-MIDAS vs. FAR-MIDAS SV									
Daily series	(a) $\underline{\psi} = 100, \underline{v}_0 = 0.001, \underline{k}_\xi = 0.1$				(b) $\underline{\psi} = 2.5, \underline{v}_0 = 0.1, \underline{k}_\xi = 0.1$				
	Ffr	20.182	53.282	23.200	53.929	27.107	29.778	10.369	19.705
	Ret	25.108	44.802	31.008	47.488	31.300	18.359	11.329	17.826
ADS	59.082	61.700	63.051	60.297	32.583	10.270	12.650	14.961	
Daily series	(c) $\underline{\psi} = 25, \underline{v}_0 = 0.005, \underline{k}_\xi = 0.01$				(d) $\underline{\psi} = 25, \underline{v}_0 = 0.005, \underline{k}_\xi = 1$				
	Ffr	-24.640	-8.949	-15.869	-95.626	4.465	4.132	-57.439	-195.555
	Ret	-15.360	-23.657	38.760	-10.402	7.004	-4.765	31.490	1.388
ADS	14.059	-9.982	38.815	19.207	11.788	-9.928	31.260	16.555	

This table reports pairwise model comparisons under four different prior choices. The table entries are equal to twice the natural logarithm of the Bayes factor, $2 \times (\ln B_{1,0})$, where $B_{1,0}$ denotes the Bayes factor obtained from comparing model M_1 vs. M_0 . $B_{1,0} = Pr(\mathcal{D}^t | M_1) / Pr(\mathcal{D}^t | M_0)$. Pairwise model comparisons are listed in the first column, where the notation ‘AR’ refers to an autoregressive model, ‘AR-MIDAS’ refers to an augmented distributed lag MIDAS model, ‘FAR’ refers to a factor augmented autoregressive model, and ‘FAR-MIDAS’ refers to a factor augmented distributed lag MIDAS model. The suffixes ‘SV’ and ‘SV-MIDAS’ denote models with stochastic volatility and MIDAS stochastic volatility, respectively. Column headers denote the daily predictor used in the MIDAS models, namely: the effective Federal Funds rate (Ffr), the interest rate spread between the 10-year government bond rate and the federal funds rate (Spr), value-weighted stock returns (Ret), the SML portfolio return (Smb), the HML portfolio return (Hml), the MOM portfolio return (Mom), the business cycle variable of [Aruoba et al. \(2009\)](#) (ADS), and the default spread between BAA and AAA rated bonds. Kass and Raftery (1995) suggest interpreting the results as follows: For $2 \times (\ln B_{1,0})$ less than 0, the evidence favors M_0 over M_1 ; for $2 \times (\ln B_{1,0})$ 0-2, 2-6, 6-10, and higher than 10, there is “weak evidence”, “some evidence”, “strong evidence”, and “very strong evidence” in favor of M_1 relative to M_0 . Bold numbers denote weak or stronger evidence in favor of M_1 .

Table 3. F test on the statistical significance of the MIDAS weighting schemes

Daily series	<i>IPI growth rate</i>					<i>Inflation rate</i>				
	$H = 1$	$H = 3$	$H = 6$	$H = 9$	$H = 12$	$H = 1$	$H = 3$	$H = 6$	$H = 9$	$H = 12$
Ffr	0.010	0.773	0.420	0.297	0.784	0.170	0.490	0.770	0.995	0.932
Spr	0.171	0.028	0.005	0.018	0.010	0.088	0.235	0.067	0.007	0.008
Ret	0.402	0.003	0.008	0.003	0.434	0.228	0.157	0.310	0.453	0.872
Smb	0.234	0.337	0.762	0.276	0.547	0.140	0.307	0.707	0.427	0.932
Hml	0.803	0.594	0.128	0.101	0.062	0.830	0.364	0.673	0.393	0.496
Mom	0.944	0.105	0.132	0.493	0.010	0.342	0.106	0.104	0.354	0.144
ADS	0.000	0.000	0.000	0.187	0.722	0.021	0.000	0.029	0.842	0.061
Def	0.003	0.066	0.064	0.031	0.093	0.037	0.144	0.155	0.196	0.063

This table reports the p-values of an F test that the MIDAS coefficients $(\theta_1, \dots, \theta_p)$ and $(\theta_{h,1}, \dots, \theta_{h,p})$ are jointly equal to zero. We use the posterior means and variance-covariances of the posterior distributions for $(\theta_1, \dots, \theta_p)$ and $(\theta_{h,1}, \dots, \theta_{h,p})$ as inputs in the F tests. Bold numbers indicate significance at the 10 percent level.

Table 4. Out-of-sample forecast performance for IPI growth rate - RMSFE

Model	No Midas	Ffr	Spr	Ret	Smb	Hml	Mom	ADS	Def
$H = 1$									
AR-MIDAS	0.750	0.750	0.769	0.728	0.747	0.754	0.747	0.669	0.741
FAR-MIDAS	0.683	0.687	0.682	0.679	0.684	0.686	0.684	0.647	0.679
AR-MIDAS SV	0.737	0.740	0.747	0.715	0.735	0.737	0.735	0.652	0.725
AR-MIDAS SV-MIDAS		0.745	0.750	0.718	0.738	0.742	0.738	0.657	0.728
FAR-MIDAS SV	0.675	0.681	0.673	0.668	0.678	0.675	0.678	0.635	0.672
FAR-MIDAS SV-MIDAS		0.683	0.675	0.672	0.680	0.678	0.680	0.639	0.673
$H = 3$									
AR-MIDAS	0.804	0.814	0.871	0.775	0.807	0.809	0.786	0.811	0.810
FAR-MIDAS	0.804	0.810	0.809	0.791	0.809	0.807	0.799	0.802	0.810
AR-MIDAS SV	0.804	0.820	0.837	0.774	0.807	0.805	0.800	0.794	0.811
AR-MIDAS SV-MIDAS		0.821	0.838	0.776	0.810	0.809	0.795	0.797	0.813
FAR-MIDAS SV	0.791	0.799	0.799	0.779	0.799	0.792	0.788	0.792	0.796
FAR-MIDAS SV-MIDAS		0.805	0.801	0.785	0.806	0.801	0.794	0.796	0.803
$H = 6$									
AR-MIDAS	0.789	0.813	0.839	0.784	0.794	0.792	0.794	0.779	0.801
FAR-MIDAS	0.778	0.783	0.796	0.776	0.781	0.784	0.788	0.775	0.792
AR-MIDAS SV	0.791	0.810	0.815	0.778	0.795	0.793	0.813	0.775	0.802
AR-MIDAS SV-MIDAS		0.813	0.818	0.779	0.795	0.793	0.808	0.776	0.804
FAR-MIDAS SV	0.763	0.774	0.776	0.760	0.765	0.765	0.776	0.761	0.774
FAR-MIDAS SV-MIDAS		0.779	0.779	0.764	0.769	0.768	0.779	0.762	0.781
$H = 9$									
AR-MIDAS	0.750	0.765	0.772	0.755	0.759	0.752	0.764	0.743	0.763
FAR-MIDAS	0.749	0.757	0.767	0.752	0.757	0.753	0.759	0.750	0.753
AR-MIDAS SV	0.755	0.769	0.768	0.752	0.759	0.758	0.784	0.747	0.765
AR-MIDAS SV-MIDAS		0.773	0.776	0.753	0.761	0.756	0.775	0.749	0.764
FAR-MIDAS SV	0.749	0.763	0.763	0.750	0.753	0.754	0.761	0.750	0.751
FAR-MIDAS SV-MIDAS		0.765	0.769	0.749	0.754	0.751	0.760	0.751	0.752
$H = 12$									
AR-MIDAS	0.703	0.730	0.719	0.705	0.708	0.702	0.735	0.705	0.732
FAR-MIDAS	0.694	0.723	0.726	0.713	0.703	0.695	0.712	0.700	0.703
AR-MIDAS SV	0.714	0.729	0.707	0.719	0.720	0.715	0.736	0.719	0.734
AR-MIDAS SV-MIDAS		0.732	0.711	0.715	0.717	0.712	0.733	0.716	0.736
FAR-MIDAS SV	0.699	0.720	0.709	0.712	0.704	0.703	0.713	0.701	0.704
FAR-MIDAS SV-MIDAS		0.723	0.715	0.714	0.706	0.703	0.715	0.704	0.707

This table reports the ratio between the RMSFE of model i and the RMSFE of the Random Walk (RW) model, computed as

$$RMSFE_i = \frac{\sqrt{\frac{1}{\bar{t}-\underline{t}+1} \sum_{\tau=\underline{t}}^{\bar{t}} e_{i,\tau}^2}}{\sqrt{\frac{1}{\bar{t}-\underline{t}+1} \sum_{\tau=\underline{t}}^{\bar{t}} e_{RW,\tau}^2}},$$

where $e_{i,\tau}^2$ and $e_{RW,\tau}^2$ are the squared forecast errors at time τ generated by model i and the RW model, respectively, and i denotes any of the models described in section 3. Values less than one for $RMSFE_i$ indicate that model i produces more accurate point forecasts than the RW model. The notation ‘AR’ refers to an autoregressive model, ‘AR-MIDAS’ refers to an augmented distributed lag MIDAS model, ‘FAR’ refers to a factor augmented autoregressive model, and ‘FAR-MIDAS’ refers to a factor augmented distributed lag MIDAS model. The suffixes ‘SV’ and ‘SV-MIDAS’ denote models with stochastic volatility and MIDAS stochastic volatility, respectively. For MIDAS models, the column headers denote the daily predictor used in the regressions, namely: the effective Federal Funds rate (Ffr), the interest rate spread between the 10-year government bond rate and the federal funds rate (Spr), value-weighted stock returns (Ret), the SML portfolio return (Smb), the HML portfolio return (Hml), the MOM portfolio return (Mom), the business cycle variable of [Aruoba et al. \(2009\)](#) (ADS), and the default spread between BAA and AAA rated bonds. All forecasts and forecast errors are produced with recursive estimates of the models. The out-of-sample period starts in 1982:1 and ends in 2011:12. Bold numbers indicate the lowest RMSFE across all predictors for a given model.

Table 5. Out-of-sample forecast performance for inflation rate - RMSFE

Model	No Midas	Ffr	Spr	Ret	Smb	Hml	Mom	ADS	Def
$H = 1$									
AR-MIDAS	0.918	0.912	0.905	0.921	0.921	0.921	0.928	0.919	0.930
FAR-MIDAS	0.918	0.920	0.914	0.916	0.922	0.921	0.925	0.916	0.918
AR-MIDAS SV	0.909	0.907	0.905	0.910	0.913	0.911	0.923	0.910	0.913
AR-MIDAS SV-MIDAS		0.905	0.901	0.912	0.913	0.913	0.922	0.909	0.917
FAR-MIDAS SV	0.911	0.915	0.916	0.910	0.915	0.913	0.923	0.913	0.908
FAR-MIDAS SV-MIDAS		0.914	0.910	0.910	0.914	0.914	0.921	0.908	0.908
$H = 3$									
AR-MIDAS	0.817	0.825	0.810	0.824	0.822	0.819	0.825	0.821	0.836
FAR-MIDAS	0.819	0.829	0.816	0.825	0.825	0.820	0.826	0.808	0.813
AR-MIDAS SV	0.815	0.822	0.812	0.817	0.814	0.815	0.817	0.814	0.823
AR-MIDAS SV-MIDAS		0.822	0.811	0.819	0.816	0.815	0.817	0.812	0.827
FAR-MIDAS SV	0.815	0.827	0.813	0.817	0.813	0.816	0.818	0.800	0.806
FAR-MIDAS SV-MIDAS		0.825	0.811	0.818	0.816	0.816	0.818	0.799	0.807
$H = 6$									
AR-MIDAS	0.795	0.792	0.800	0.804	0.796	0.801	0.804	0.809	0.820
FAR-MIDAS	0.802	0.803	0.809	0.813	0.806	0.806	0.812	0.813	0.816
AR-MIDAS SV	0.798	0.795	0.804	0.812	0.799	0.799	0.803	0.814	0.795
AR-MIDAS SV-MIDAS		0.786	0.791	0.795	0.788	0.794	0.796	0.796	0.799
FAR-MIDAS SV	0.806	0.806	0.811	0.816	0.800	0.806	0.811	0.810	0.791
FAR-MIDAS SV-MIDAS		0.790	0.796	0.798	0.790	0.797	0.800	0.794	0.791
$H = 9$									
AR-MIDAS	0.808	0.822	0.832	0.821	0.814	0.816	0.821	0.828	0.842
FAR-MIDAS	0.819	0.837	0.842	0.834	0.827	0.823	0.831	0.837	0.843
AR-MIDAS SV	0.798	0.799	0.824	0.791	0.787	0.790	0.790	0.800	0.790
AR-MIDAS SV-MIDAS		0.801	0.805	0.794	0.790	0.794	0.796	0.800	0.795
FAR-MIDAS SV	0.818	0.804	0.811	0.800	0.788	0.793	0.798	0.805	0.788
FAR-MIDAS SV-MIDAS		0.809	0.814	0.805	0.793	0.797	0.803	0.804	0.791
$H = 12$									
AR-MIDAS	0.786	0.785	0.802	0.794	0.792	0.801	0.793	0.798	0.810
FAR-MIDAS	0.782	0.796	0.798	0.787	0.788	0.790	0.790	0.797	0.797
AR-MIDAS SV	0.749	0.750	0.754	0.749	0.747	0.756	0.746	0.760	0.755
AR-MIDAS SV-MIDAS		0.757	0.765	0.755	0.753	0.765	0.752	0.759	0.760
FAR-MIDAS SV	0.750	0.755	0.756	0.752	0.747	0.753	0.748	0.753	0.747
FAR-MIDAS SV-MIDAS		0.762	0.764	0.759	0.754	0.762	0.755	0.756	0.754

This table reports the ratio between the RMSFE of model i and the RMSFE of the Random Walk (RW) model, computed as

$$RMSFE_i = \frac{\sqrt{\frac{1}{\bar{t}-t+1} \sum_{\tau=t}^{\bar{t}} e_{i,\tau}^2}}{\sqrt{\frac{1}{\bar{t}-t+1} \sum_{\tau=t}^{\bar{t}} e_{RW,\tau}^2}},$$

where $e_{i,\tau}^2$ and $e_{RW,\tau}^2$ are the squared forecast errors at time τ generated by model i and the RW model, respectively, and i denotes any of the models described in section 3. Values less than one for $RMSFE_i$ indicate that model i produces more accurate point forecasts than the RW model. The notation ‘AR’ refers to an autoregressive model, ‘AR-MIDAS’ refers to an augmented distributed lag MIDAS model, ‘FAR’ refers to a factor augmented autoregressive model, and ‘FAR-MIDAS’ refers to a factor augmented distributed lag MIDAS model. The suffixes ‘SV’ and ‘SV-MIDAS’ denote models with stochastic volatility and MIDAS stochastic volatility, respectively. For MIDAS models, the column headers denote the daily predictor used in the regressions, namely: the effective Federal Funds rate (Ffr), the interest rate spread between the 10-year government bond rate and the federal funds rate (Spr), value-weighted stock returns (Ret), the SML portfolio return (Smb), the HML portfolio return (Hml), the MOM portfolio return (Mom), the business cycle variable of [Aruoba et al. \(2009\)](#) (ADS), and the default spread between BAA and AAA rated bonds. All forecasts and forecast errors are produced with recursive estimates of the models. The out-of-sample period starts in 1982:1 and ends in 2011:12. Bold numbers indicate the lowest RMSFE across all predictors for a given model.

Table 6. Diebold-Mariano p-values from tests of equal predictive ability - RMSFE

Comparison	Ffr	Spr	Ret	Smb	Hml	Mom	ADS	Def
<i>Panel A: IPI growth rate</i>								
<i>H = 1</i>								
AR-MIDAS vs. AR	0.353	0.969	0.018	0.236	0.846	0.210	0.013	0.264
AR-MIDAS SV vs. AR SV	0.509	0.945	0.012	0.326	0.631	0.220	0.003	0.133
FAR-MIDAS vs. FAR	0.704	0.384	0.260	0.847	0.777	0.691	0.015	0.285
FAR-MIDAS SV vs. FAR SV	0.871	0.300	0.146	0.857	0.546	0.808	0.004	0.299
<i>H = 3</i>								
AR-MIDAS vs. AR	0.782	0.999	0.104	0.689	0.762	0.154	0.777	0.610
AR-MIDAS SV vs. AR SV	0.895	0.995	0.065	0.702	0.640	0.301	0.266	0.592
FAR-MIDAS vs. FAR	0.949	0.896	0.153	0.903	0.691	0.417	0.408	0.740
FAR-MIDAS SV vs. FAR SV	0.951	0.917	0.156	0.882	0.659	0.489	0.501	0.579
<i>H = 6</i>								
AR-MIDAS vs. AR	0.922	0.991	0.441	0.849	0.635	0.653	0.337	0.820
AR-MIDAS SV vs. AR SV	0.898	0.970	0.278	0.796	0.632	0.960	0.087	0.672
FAR-MIDAS vs. FAR	0.826	0.982	0.393	0.896	0.745	0.976	0.328	0.902
FAR-MIDAS SV vs. FAR SV	0.895	0.977	0.354	0.772	0.645	0.973	0.418	0.774
<i>H = 9</i>								
AR-MIDAS vs. AR	0.898	0.913	0.580	0.935	0.546	0.969	0.253	0.897
AR-MIDAS SV vs. AR SV	0.894	0.903	0.427	0.701	0.667	0.992	0.135	0.862
FAR-MIDAS vs. FAR	0.893	0.973	0.595	0.975	0.754	0.984	0.609	0.762
FAR-MIDAS SV vs. FAR SV	0.931	0.989	0.633	0.795	0.869	0.990	0.258	0.730
<i>H = 12</i>								
AR-MIDAS vs. AR	0.900	0.675	0.535	0.996	0.451	0.976	0.586	0.961
AR-MIDAS SV vs. AR SV	0.880	0.406	0.850	0.968	0.559	0.953	0.827	0.936
FAR-MIDAS vs. FAR	0.967	0.937	0.997	0.993	0.627	0.971	0.952	0.891
FAR-MIDAS SV vs. FAR SV	0.991	0.655	0.999	0.980	0.706	0.976	0.927	0.850
<i>Panel B: Inflation rate</i>								
<i>H = 1</i>								
AR-MIDAS vs. AR	0.208	0.014	0.702	0.751	0.826	0.988	0.524	0.951
AR-MIDAS SV vs. AR SV	0.320	0.187	0.604	0.802	0.736	0.993	0.602	0.921
FAR-MIDAS vs. FAR	0.639	0.208	0.418	0.736	0.756	0.963	0.341	0.487
FAR-MIDAS SV vs. FAR SV	0.780	0.853	0.448	0.695	0.708	0.983	0.568	0.408
<i>H = 3</i>								
AR-MIDAS vs. AR	0.827	0.210	0.865	0.897	0.724	0.961	0.615	0.903
AR-MIDAS SV vs. AR SV	0.915	0.448	0.770	0.401	0.511	0.730	0.541	0.838
FAR-MIDAS vs. FAR	0.894	0.348	0.788	0.843	0.465	0.969	0.190	0.279
FAR-MIDAS SV vs. FAR SV	0.970	0.327	0.631	0.354	0.529	0.663	0.117	0.225
<i>H = 6</i>								
AR-MIDAS vs. AR	0.281	0.765	0.748	0.799	0.754	0.947	0.698	0.817
AR-MIDAS SV vs. AR SV	0.041	0.917	0.952	0.615	0.719	0.972	0.961	0.293
FAR-MIDAS vs. FAR	0.619	0.828	0.802	0.964	0.661	0.946	0.948	0.844
FAR-MIDAS SV vs. FAR SV	0.470	0.818	0.895	0.183	0.509	0.970	0.728	0.057
<i>H = 9</i>								
AR-MIDAS vs. AR	0.524	0.918	0.722	0.952	0.742	0.924	0.712	0.791
AR-MIDAS SV vs. AR SV	0.582	0.997	0.341	0.225	0.287	0.323	0.480	0.302
FAR-MIDAS vs. FAR	0.675	0.950	0.765	0.989	0.609	0.934	0.930	0.875
FAR-MIDAS SV vs. FAR SV	0.234	0.387	0.246	0.086	0.140	0.196	0.252	0.116
<i>H = 12</i>								
AR-MIDAS vs. AR	0.163	0.806	0.972	0.935	0.889	0.851	0.748	0.963
AR-MIDAS SV vs. AR SV	0.219	0.675	0.432	0.370	0.770	0.193	0.782	0.804
FAR-MIDAS vs. FAR	0.864	0.928	0.877	0.988	0.832	0.880	1.000	0.689
FAR-MIDAS SV vs. FAR SV	0.886	0.850	0.736	0.303	0.752	0.367	0.946	0.262

This table reports Diebold-Mariano p-values under the null that a given MIDAS-in-mean model has the same predictive ability as its nested non-MIDAS model, using mean squared error loss. The notation ‘AR’ refers to an autoregressive model, ‘AR-MIDAS’ refers to an augmented distributed lag MIDAS model, ‘FAR’ refers to a factor augmented autoregressive model, and ‘FAR-MIDAS’ refers to a factor augmented distributed lag MIDAS model. The suffixes ‘SV’ and ‘SV-MIDAS’ denote models with stochastic volatility and MIDAS stochastic volatility, respectively. For MIDAS models, the column headers denote the daily predictor used in the regressions, namely: the effective Federal Funds rate (Ffr), the interest rate spread between the 10-year government bond rate and the federal funds rate (Spr), value-weighted stock returns (Ret), the SML portfolio return (Smb), the HML portfolio return (Hml), the MOM portfolio return (Mom), the business cycle variable of [Aruoba et al. \(2009\)](#) (ADS), and the default spread between BAA and AAA rated bonds. The underlying p-values are based on one-sided t -test, where the t -statistics computed with a serial correlation robust variance, using the pre-whitened quadratic spectral estimation of Andrews and Monahan (1992). Bold numbers indicate significance at the 10 percent level.

Table 7. Out-of-sample forecast performance for IPI growth rate - Log score differentials (LSD)

Model	No Midas	Ffr	Spr	Ret	Smb	Hml	Mom	ADS	Def
$H = 1$									
AR-MIDAS	0.248	0.251	0.239	0.276	0.255	0.244	0.246	0.371	0.258
FAR-MIDAS	0.345	0.342	0.344	0.350	0.343	0.341	0.341	0.397	0.347
AR-MIDAS SV	0.339	0.329	0.310	0.352	0.341	0.326	0.332	0.435	0.343
AR-MIDAS SV-MIDAS		0.342	0.332	0.371	0.360	0.344	0.350	0.448	0.355
FAR-MIDAS SV	0.386	0.380	0.387	0.384	0.379	0.376	0.374	0.443	0.380
FAR-MIDAS SV-MIDAS		0.412	0.422	0.420	0.415	0.411	0.412	0.465	0.415
$H = 3$									
AR-MIDAS	0.235	0.233	0.202	0.274	0.234	0.230	0.245	0.240	0.227
FAR-MIDAS	0.281	0.271	0.272	0.291	0.272	0.276	0.277	0.277	0.270
AR-MIDAS SV	0.304	0.289	0.248	0.302	0.292	0.283	0.290	0.310	0.285
AR-MIDAS SV-MIDAS		0.316	0.288	0.345	0.325	0.315	0.333	0.340	0.321
FAR-MIDAS SV	0.296	0.284	0.280	0.295	0.277	0.275	0.275	0.281	0.274
FAR-MIDAS SV-MIDAS		0.321	0.332	0.343	0.324	0.318	0.325	0.331	0.311
$H = 6$									
AR-MIDAS	0.291	0.291	0.273	0.302	0.284	0.286	0.283	0.300	0.279
FAR-MIDAS	0.322	0.321	0.305	0.323	0.317	0.316	0.307	0.327	0.309
AR-MIDAS SV	0.339	0.312	0.307	0.337	0.322	0.296	0.284	0.351	0.292
AR-MIDAS SV-MIDAS		0.366	0.354	0.380	0.371	0.348	0.344	0.390	0.335
FAR-MIDAS SV	0.368	0.349	0.353	0.356	0.355	0.331	0.328	0.368	0.321
FAR-MIDAS SV-MIDAS		0.390	0.391	0.393	0.395	0.379	0.365	0.404	0.334
$H = 9$									
AR-MIDAS	0.302	0.306	0.294	0.304	0.294	0.298	0.287	0.308	0.289
FAR-MIDAS	0.319	0.317	0.302	0.316	0.311	0.313	0.306	0.318	0.314
AR-MIDAS SV	0.337	0.328	0.304	0.319	0.314	0.299	0.288	0.346	0.308
AR-MIDAS SV-MIDAS		0.386	0.337	0.371	0.364	0.359	0.342	0.384	0.345
FAR-MIDAS SV	0.338	0.334	0.302	0.327	0.325	0.304	0.310	0.343	0.315
FAR-MIDAS SV-MIDAS		0.394	0.351	0.380	0.375	0.368	0.357	0.387	0.368
$H = 12$									
AR-MIDAS	0.348	0.342	0.350	0.344	0.342	0.348	0.324	0.344	0.323
FAR-MIDAS	0.369	0.359	0.348	0.349	0.359	0.366	0.355	0.363	0.360
AR-MIDAS SV	0.360	0.321	0.352	0.347	0.335	0.329	0.318	0.350	0.332
AR-MIDAS SV-MIDAS		0.403	0.414	0.414	0.402	0.411	0.378	0.414	0.379
FAR-MIDAS SV	0.361	0.325	0.343	0.329	0.338	0.328	0.321	0.351	0.347
FAR-MIDAS SV-MIDAS		0.406	0.406	0.394	0.410	0.408	0.394	0.416	0.406

This table reports the average log-score (LS) differential, $LSD_i = \sum_{\tau=\underline{t}}^{\bar{t}} (LS_{i,\tau} - LS_{RW,\tau})$, where $LS_{i,\tau}$ ($LS_{RW,\tau}$) denotes the log-score of model i (RW), computed at time τ , and i denotes any of the models described in section 3. Positive values of LSD_i indicate that model i produces more accurate density forecasts than the RW model. The notation ‘AR’ refers to an autoregressive model, ‘AR-MIDAS’ refers to an augmented distributed lag MIDAS model, ‘FAR’ refers to a factor augmented autoregressive model, and ‘FAR-MIDAS’ refers to a factor augmented distributed lag MIDAS model. The suffixes ‘SV’ and ‘SV-MIDAS’ denote models with stochastic volatility and MIDAS stochastic volatility, respectively. For MIDAS models, the column headers denote the daily predictor used in the regressions, namely: the effective Federal Funds rate (Ffr), the interest rate spread between the 10-year government bond rate and the federal funds rate (Spr), value-weighted stock returns (Ret), the SML portfolio return (Smb), the HML portfolio return (Hml), the MOM portfolio return (Mom), the business cycle variable of [Aruoba et al. \(2009\)](#) (ADS), and the default spread between BAA and AAA rated bonds. All forecasts and forecast errors are produced with recursive estimates of the models. The out-of-sample period starts in 1982:1 and ends in 2011:12. Bold numbers indicate the largest LSD across all predictors for a given model.

Table 8. Out-of-sample forecast performance for inflation rate - Log score differentials (LSD)

Model	No Midas	Ffr	Spr	Ret	Smb	Hml	Mom	ADS	Def
$H = 1$									
AR-MIDAS	0.141	0.159	0.158	0.136	0.140	0.134	0.139	0.139	0.137
FAR-MIDAS	0.164	0.161	0.162	0.165	0.167	0.158	0.159	0.167	0.162
AR-MIDAS SV	0.217	0.211	0.214	0.214	0.216	0.210	0.211	0.213	0.223
AR-MIDAS SV-MIDAS		0.242	0.254	0.239	0.241	0.225	0.230	0.235	0.241
FAR-MIDAS SV	0.221	0.211	0.217	0.215	0.222	0.211	0.217	0.228	0.236
FAR-MIDAS SV-MIDAS		0.248	0.259	0.255	0.259	0.238	0.247	0.259	0.264
$H = 3$									
AR-MIDAS	0.206	0.212	0.218	0.198	0.205	0.203	0.197	0.208	0.193
FAR-MIDAS	0.213	0.212	0.213	0.209	0.211	0.210	0.207	0.235	0.220
AR-MIDAS SV	0.267	0.254	0.261	0.258	0.269	0.258	0.264	0.257	0.262
AR-MIDAS SV-MIDAS		0.298	0.322	0.302	0.303	0.302	0.306	0.294	0.291
FAR-MIDAS SV	0.262	0.251	0.251	0.244	0.259	0.250	0.257	0.268	0.272
FAR-MIDAS SV-MIDAS		0.299	0.307	0.295	0.297	0.294	0.304	0.313	0.316
$H = 6$									
AR-MIDAS	0.202	0.199	0.193	0.195	0.195	0.198	0.191	0.191	0.181
FAR-MIDAS	0.198	0.196	0.192	0.190	0.188	0.193	0.187	0.184	0.185
AR-MIDAS SV	0.199	0.195	0.177	0.159	0.195	0.184	0.167	0.172	0.198
AR-MIDAS SV-MIDAS		0.276	0.268	0.267	0.274	0.272	0.257	0.251	0.255
FAR-MIDAS SV	0.189	0.184	0.166	0.148	0.204	0.183	0.146	0.179	0.193
FAR-MIDAS SV-MIDAS		0.274	0.268	0.255	0.273	0.270	0.253	0.245	0.263
$H = 9$									
AR-MIDAS	0.168	0.151	0.148	0.160	0.160	0.161	0.150	0.159	0.143
FAR-MIDAS	0.170	0.156	0.157	0.160	0.161	0.165	0.156	0.148	0.150
AR-MIDAS SV	0.212	0.208	0.156	0.207	0.214	0.205	0.215	0.194	0.203
AR-MIDAS SV-MIDAS		0.240	0.205	0.248	0.251	0.250	0.246	0.223	0.225
FAR-MIDAS SV	0.182	0.196	0.185	0.190	0.199	0.187	0.199	0.172	0.188
FAR-MIDAS SV-MIDAS		0.238	0.215	0.235	0.245	0.232	0.231	0.217	0.233
$H = 12$									
AR-MIDAS	0.212	0.219	0.208	0.206	0.204	0.199	0.203	0.208	0.193
FAR-MIDAS	0.225	0.217	0.223	0.219	0.216	0.219	0.217	0.212	0.209
AR-MIDAS SV	0.281	0.277	0.245	0.270	0.273	0.265	0.278	0.255	0.279
AR-MIDAS SV-MIDAS		0.315	0.268	0.308	0.310	0.303	0.316	0.286	0.279
FAR-MIDAS SV	0.270	0.267	0.260	0.250	0.258	0.258	0.268	0.254	0.268
FAR-MIDAS SV-MIDAS		0.306	0.289	0.296	0.303	0.295	0.310	0.291	0.302

This table reports the average log-score (LS) differential, $LSD_i = \sum_{\tau=\underline{t}}^{\bar{t}} (LS_{i,\tau} - LS_{RW,\tau})$, where $LS_{i,\tau}$ ($LS_{RW,\tau}$) denotes the log-score of model i (RW), computed at time τ , and i denotes any of the models described in section 3. Positive values of LSD_i indicate that model i produces more accurate density forecasts than the RW model. The notation ‘AR’ refers to an autoregressive model, ‘AR-MIDAS’ refers to an augmented distributed lag MIDAS model, ‘FAR’ refers to a factor augmented autoregressive model, and ‘FAR-MIDAS’ refers to a factor augmented distributed lag MIDAS model. The suffixes ‘SV’ and ‘SV-MIDAS’ denote models with stochastic volatility and MIDAS stochastic volatility, respectively. For MIDAS models, the column headers denote the daily predictor used in the regressions, namely: the effective Federal Funds rate (Ffr), the interest rate spread between the 10-year government bond rate and the federal funds rate (Spr), value-weighted stock returns (Ret), the SML portfolio return (Smb), the HML portfolio return (Hml), the MOM portfolio return (Mom), the business cycle variable of [Aruoba et al. \(2009\)](#) (ADS), and the default spread between BAA and AAA rated bonds. All forecasts and forecast errors are produced with recursive estimates of the models. The out-of-sample period starts in 1982:1 and ends in 2011:12. Bold numbers indicate the largest LSD across all predictors for a given model.

Table 9. Diebold-Mariano p-values from tests of equal predictive ability - Log score differentials

Comparison	Ffr	Spr	Ret	Smb	Hml	Mom	ADS	Def
<i>Panel A: IPI growth rate</i>								
<i>H = 1</i>								
AR-MIDAS SV-MIDAS vs. AR-MIDAS SV	0.156	0.065	0.102	0.095	0.089	0.093	0.175	0.164
FAR-MIDAS SV-MIDAS vs. FAR-MIDAS SV	0.007	0.004	0.008	0.004	0.013	0.002	0.052	0.008
<i>H = 3</i>								
AR-MIDAS SV-MIDAS vs. AR-MIDAS SV	0.060	0.010	0.019	0.022	0.021	0.012	0.024	0.014
FAR-MIDAS SV-MIDAS vs. FAR-MIDAS SV	0.008	0.001	0.004	0.003	0.005	0.002	0.001	0.020
<i>H = 6</i>								
AR-MIDAS SV-MIDAS vs. AR-MIDAS SV	0.003	0.003	0.002	0.006	0.002	0.016	0.016	0.024
FAR-MIDAS SV-MIDAS vs. FAR-MIDAS SV	0.005	0.005	0.005	0.009	0.001	0.027	0.011	0.191
<i>H = 9</i>								
AR-MIDAS SV-MIDAS vs. AR-MIDAS SV	0.004	0.067	0.003	0.016	0.009	0.026	0.056	0.069
FAR-MIDAS SV-MIDAS vs. FAR-MIDAS SV	0.003	0.015	0.002	0.009	0.003	0.017	0.016	0.011
<i>H = 12</i>								
AR-MIDAS SV-MIDAS vs. AR-MIDAS SV	0.003	0.006	0.004	0.005	0.001	0.026	0.009	0.072
FAR-MIDAS SV-MIDAS vs. FAR-MIDAS SV	0.004	0.007	0.005	0.005	0.003	0.008	0.012	0.023
<i>Panel B: Inflation rate</i>								
<i>H = 1</i>								
AR-MIDAS SV-MIDAS vs. AR-MIDAS SV	0.005	0.001	0.016	0.028	0.103	0.079	0.043	0.043
FAR-MIDAS SV-MIDAS vs. FAR-MIDAS SV	0.001	0.000	0.000	0.002	0.015	0.014	0.009	0.008
<i>H = 3</i>								
AR-MIDAS SV-MIDAS vs. AR-MIDAS SV	0.004	0.000	0.004	0.019	0.004	0.005	0.022	0.031
FAR-MIDAS SV-MIDAS vs. FAR-MIDAS SV	0.002	0.001	0.001	0.013	0.003	0.002	0.005	0.004
<i>H = 6</i>								
AR-MIDAS SV-MIDAS vs. AR-MIDAS SV	0.015	0.007	0.009	0.015	0.012	0.014	0.009	0.035
FAR-MIDAS SV-MIDAS vs. FAR-MIDAS SV	0.014	0.006	0.008	0.023	0.012	0.009	0.019	0.022
<i>H = 9</i>								
AR-MIDAS SV-MIDAS vs. AR-MIDAS SV	0.085	0.035	0.017	0.039	0.017	0.055	0.092	0.108
FAR-MIDAS SV-MIDAS vs. FAR-MIDAS SV	0.018	0.072	0.009	0.018	0.019	0.049	0.023	0.011
<i>H = 12</i>								
AR-MIDAS SV-MIDAS vs. AR-MIDAS SV	0.019	0.073	0.022	0.015	0.021	0.016	0.064	0.458
FAR-MIDAS SV-MIDAS vs. FAR-MIDAS SV	0.014	0.061	0.009	0.004	0.025	0.011	0.032	0.040

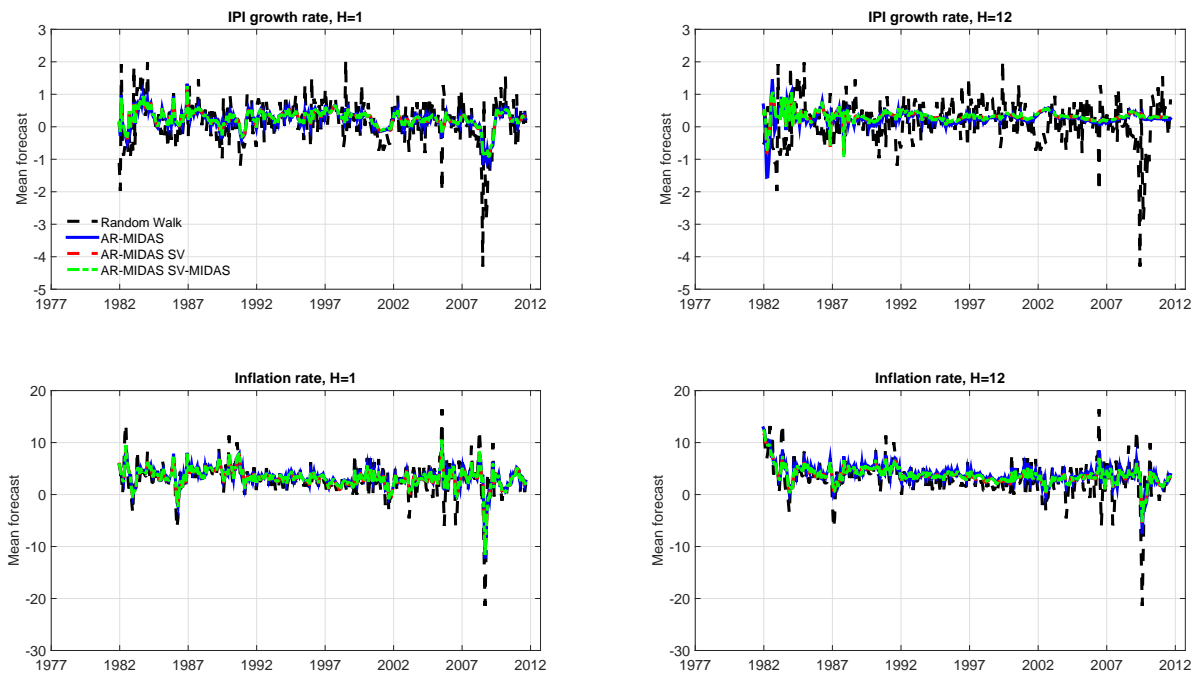
This table reports Diebold-Mariano p-values under the null that a given (F)AR-MIDAS SV-MIDAS model has the same predictive ability as the (F)AR-MIDAS SV model. P-values are computed for the LSD statistic. The notation ‘AR’ refers to an autoregressive model, ‘AR-MIDAS’ refers to an augmented distributed lag MIDAS model, ‘FAR’ refers to a factor augmented autoregressive model, and ‘FAR-MIDAS’ refers to a factor augmented distributed lag MIDAS model. The suffixes ‘SV’ and ‘SV-MIDAS’ denote models with stochastic volatility and MIDAS stochastic volatility, respectively. For MIDAS models, the column headers denote the daily predictor used in the regressions, namely: the effective Federal Funds rate (Ffr), the interest rate spread between the 10-year government bond rate and the federal funds rate (Spr), value-weighted stock returns (Ret), the SML portfolio return (Smb), the HML portfolio return (Hml), the MOM portfolio return (Mom), the business cycle variable of [Aruoba et al. \(2009\)](#) (ADS), and the default spread between BAA and AAA rated bonds. The underlying p-values are based on one-sided t -test, where the t -statistics computed with a serial correlation robust variance, using the pre-whitened quadratic spectral estimation of [Andrews and Monahan \(1992\)](#). Bold numbers indicate significance at the 10 percent level.

Table 10. Out-of-sample forecast performance - model combinations

<i>IPI growth rate</i>										
Model	<i>Panel A: RMSFE</i>					<i>Panel B: LSD</i>				
	<i>H = 1</i>	<i>H = 3</i>	<i>H = 6</i>	<i>H = 9</i>	<i>H = 12</i>	<i>H = 1</i>	<i>H = 3</i>	<i>H = 6</i>	<i>H = 9</i>	<i>H = 12</i>
Equal-weighted combination	0.675	0.780	0.766	0.743	0.697	0.407	0.338	0.386	0.379	0.419
Bayesian model averaging	0.649	0.787	0.770	0.761	0.723	0.453	0.325	0.382	0.369	0.366
Optimal prediction pool	0.657	0.780	0.774	0.758	0.717	0.446	0.331	0.376	0.368	0.374
<i>Inflation rate</i>										
Model	<i>Panel A: RMSFE</i>					<i>Panel B: LSD</i>				
	<i>H = 1</i>	<i>H = 3</i>	<i>H = 6</i>	<i>H = 9</i>	<i>H = 12</i>	<i>H = 1</i>	<i>H = 3</i>	<i>H = 6</i>	<i>H = 9</i>	<i>H = 12</i>
Equal-weighted combination	0.000	0.016	0.089	0.321	0.214	0.000	0.000	0.000	0.001	0.005
Bayesian model averaging	0.001	0.110	0.240	0.823	0.838	0.000	0.001	0.001	0.027	0.360
Optimal prediction pool	0.000	0.027	0.228	0.818	0.772	0.000	0.001	0.002	0.020	0.295
<i>IPI growth rate</i>										
Model	<i>Panel A: RMSFE</i>					<i>Panel B: LSD</i>				
	<i>H = 1</i>	<i>H = 3</i>	<i>H = 6</i>	<i>H = 9</i>	<i>H = 12</i>	<i>H = 1</i>	<i>H = 3</i>	<i>H = 6</i>	<i>H = 9</i>	<i>H = 12</i>
Equal-weighted combination	0.000	0.010	0.032	0.141	0.026	0.001	0.096	0.051	0.073	0.031
Bayesian model averaging	0.002	0.089	0.219	0.796	0.740	0.000	0.179	0.072	0.158	0.447
Optimal prediction pool	0.001	0.015	0.212	0.709	0.611	0.000	0.098	0.071	0.136	0.358
<i>Inflation rate</i>										
Model	<i>Panel A: RMSFE</i>					<i>Panel B: LSD</i>				
	<i>H = 1</i>	<i>H = 3</i>	<i>H = 6</i>	<i>H = 9</i>	<i>H = 12</i>	<i>H = 1</i>	<i>H = 3</i>	<i>H = 6</i>	<i>H = 9</i>	<i>H = 12</i>
Equal-weighted combination	0.908	0.812	0.778	0.797	0.759	0.229	0.283	0.254	0.227	0.286
Bayesian model averaging	0.910	0.805	0.783	0.799	0.760	0.252	0.309	0.261	0.228	0.292
Optimal prediction pool	0.912	0.797	0.773	0.787	0.764	0.241	0.308	0.266	0.233	0.282
<i>Inflation rate</i>										
Model	<i>Panel A: RMSFE</i>					<i>Panel B: LSD</i>				
	<i>H = 1</i>	<i>H = 3</i>	<i>H = 6</i>	<i>H = 9</i>	<i>H = 12</i>	<i>H = 1</i>	<i>H = 3</i>	<i>H = 6</i>	<i>H = 9</i>	<i>H = 12</i>
Equal-weighted combination	0.033	0.347	0.247	0.175	0.004	0.000	0.000	0.009	0.003	0.000
Bayesian model averaging	0.129	0.096	0.375	0.239	0.011	0.000	0.000	0.022	0.019	0.001
Optimal prediction pool	0.182	0.049	0.193	0.066	0.036	0.000	0.000	0.012	0.020	0.006
<i>Inflation rate</i>										
Model	<i>Panel A: RMSFE</i>					<i>Panel B: LSD</i>				
	<i>H = 1</i>	<i>H = 3</i>	<i>H = 6</i>	<i>H = 9</i>	<i>H = 12</i>	<i>H = 1</i>	<i>H = 3</i>	<i>H = 6</i>	<i>H = 9</i>	<i>H = 12</i>
Equal-weighted combination	0.428	0.427	0.141	0.450	0.868	0.154	0.145	0.018	0.222	0.378
Bayesian model averaging	0.528	0.147	0.167	0.477	0.870	0.004	0.011	0.012	0.214	0.294
Optimal prediction pool	0.621	0.061	0.061	0.252	0.927	0.035	0.009	0.005	0.112	0.463

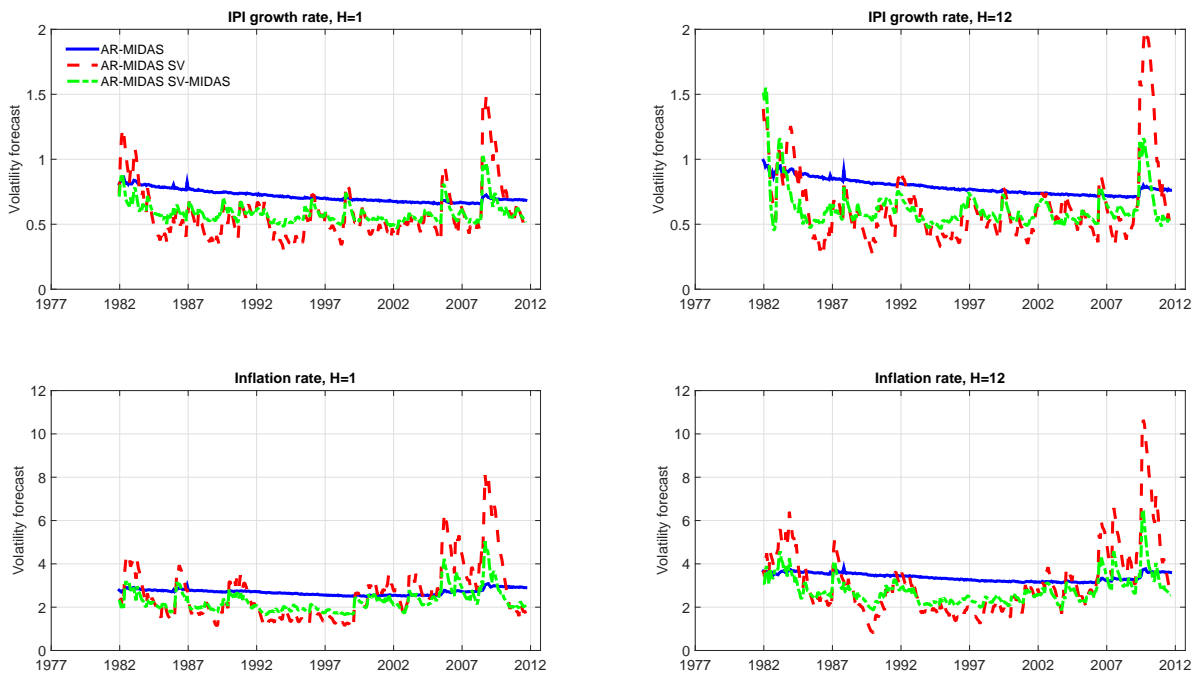
This table reports out-of-sample results for the optimal predictive pool of Geweke and Amisano (2011), an equal-weighted model combination scheme, and Bayesian Model Averaging applied to 36 MIDAS forecasting models that use different daily predictors and make different choices regarding the inclusion or exclusion of additional macro factors and the volatility dynamics of US quarterly GDP growth. In each case the models and combination weights are estimated recursively using only data up to the point of the forecast. Panels A and B report various measures of point and density forecast performance over two out-of-sample periods, namely 1982:1 to 2011:12 (panel A), and 2001:I to 2008:IV (panel B). The statistics reported are: $RMSFE_i$, the ratio between the RMSFE of model combination i and the RMSFE of the Random Walk (RW) model; LSD_i , the average log-score differential between model combination i and the RW model; $CRPSD_i$, the average continuously ranked probability score differential of model combination i relative to the RW model. Values less than one for $RMSFE_i$ indicate that model combination i produces more accurate point forecasts than the RW model; positive values of LSD_i indicate that model combination i produces more accurate density forecasts than the RW model; values less than one for $CRPSD_i$ suggest that model combination i performs better than the benchmark RW model. Panels C and D report the Diebold-Mariano p-values under the null that a given model combination has the same predictive ability than the benchmark model, over the same two out-of-sample periods, 1982:1 to 2011:12 (panel C), and 2001:I to 2008:IV (panel D). P-values are computed for the RMSFE, LSD, and the CRPSD statistics, both against the Random Walk (RW) benchmark and the autoregressive (AR) benchmark, and are based on one-sided t -tests, where the t -statistics are computed with a serial correlation robust variance, using the pre-whitened quadratic spectral estimation of Andrews and Monahan (1992). Bold numbers in Panels C-F indicate significance at the 10 percent level.

Figure 1. Mean forecasts of MIDAS models



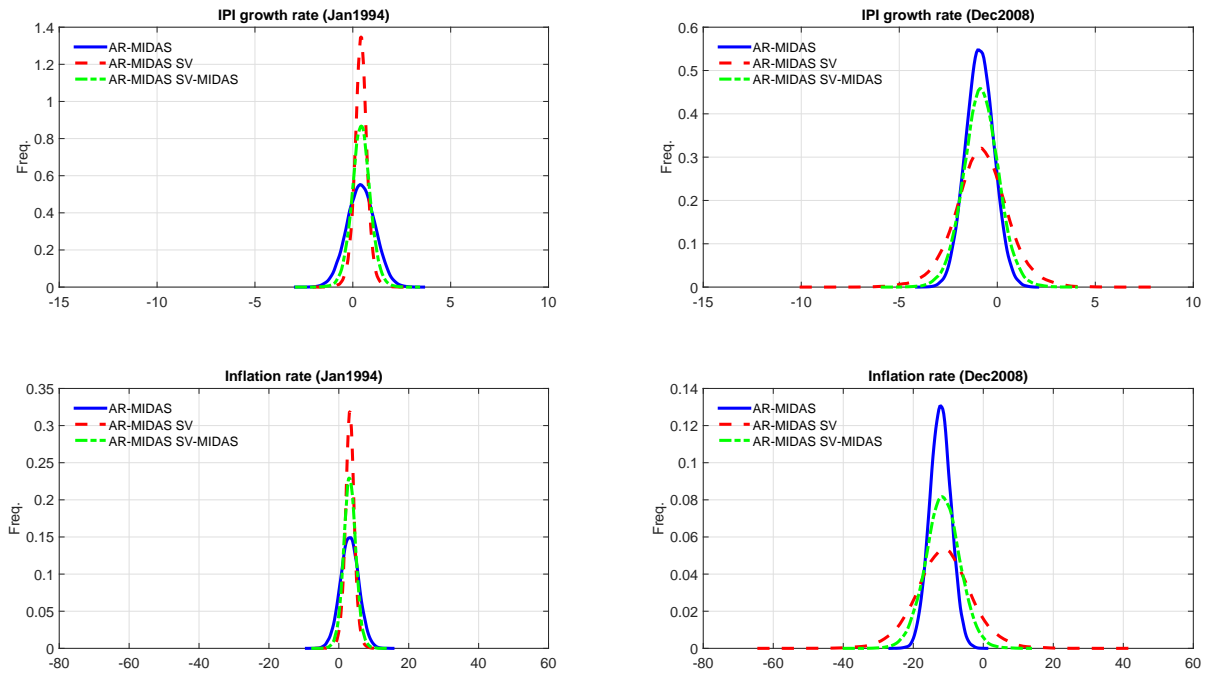
This figure shows the recursive conditional mean forecasts computed using the Random Walk model (dashed black line) as well as the predictive distributions of the AR-MIDAS (solid blue line), the AR-MIDAS SV (dashed red line), and the AR-MIDAS SV-MIDAS (green dashed-dotted line) models. All MIDAS models displayed in the two panels use the Ffr daily series as predictor. All models are estimated recursively over the out-of-sample period, which starts in 1982:1 and ends in 2011:12.

Figure 2. Volatility forecasts of MIDAS models



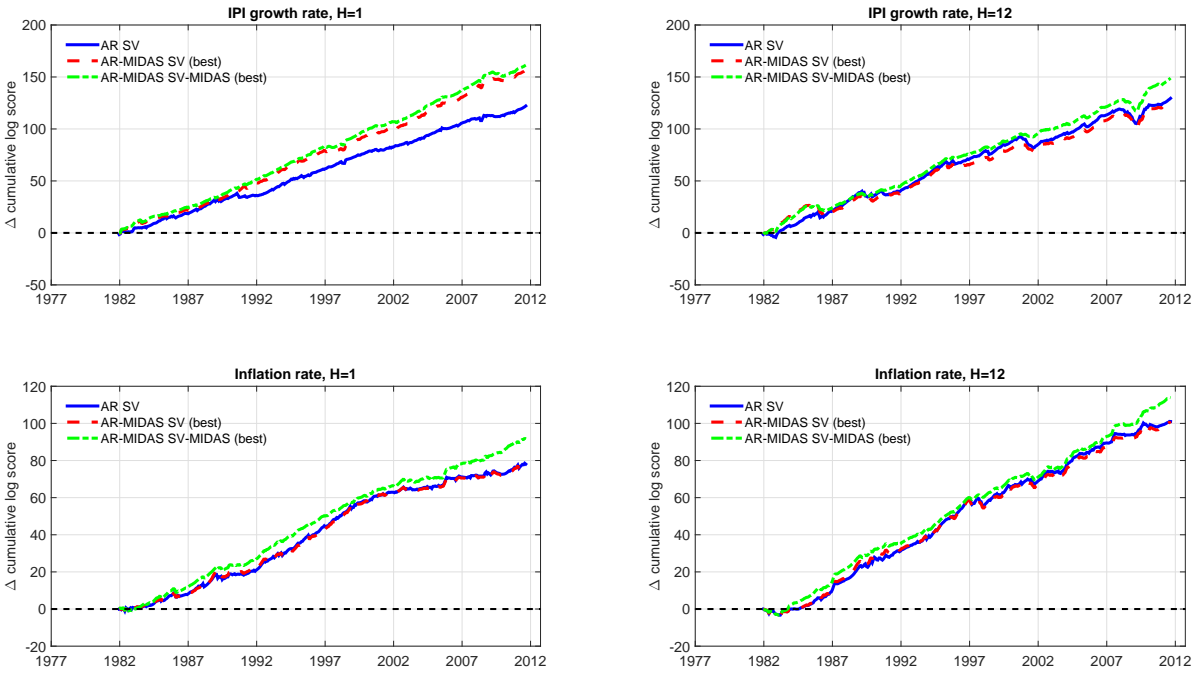
This figure shows the recursive conditional volatility forecasts computed using the predictive distributions of the AR-MIDAS (solid blue line), the AR-MIDAS SV (dashed red line), and the AR-MIDAS SV-MIDAS (green dashed-dotted line) models. All MIDAS models displayed in the two panels use the Ffr daily series as predictor. All models are estimated recursively over the out-of-sample period, which starts in 1982:1 and ends in 2011:12.

Figure 3. Predictive density of IPI growth and inflation under different AR-MIDAS models



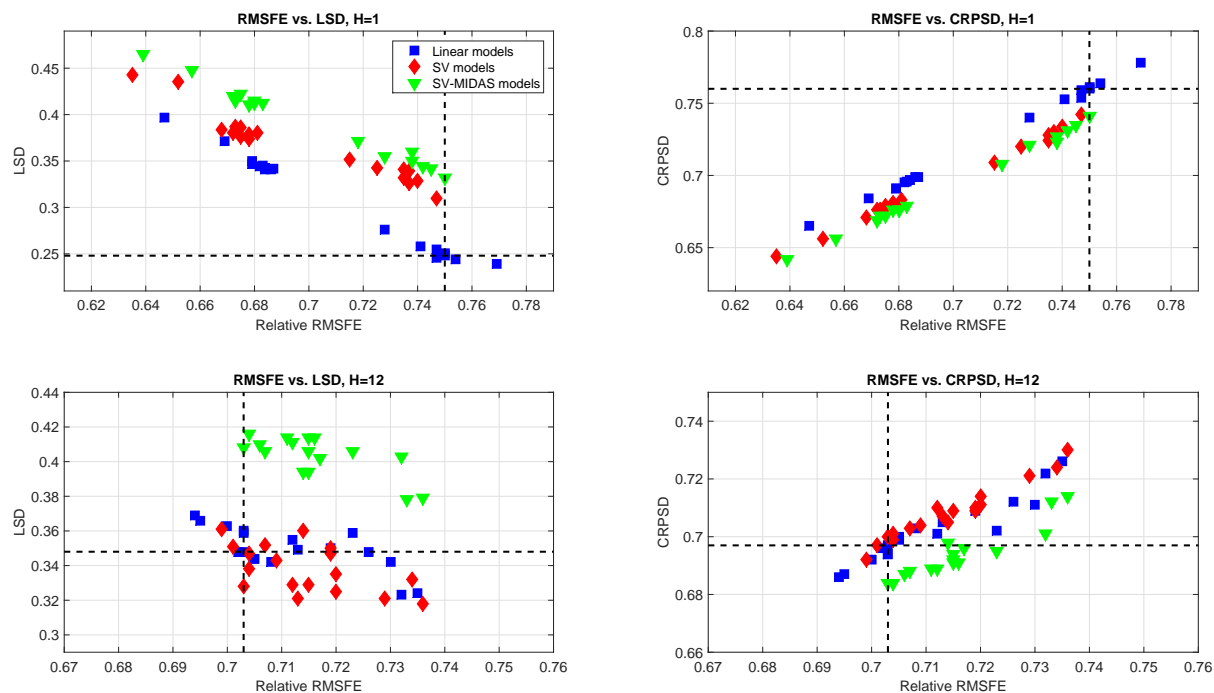
This figure shows the predictive density for the IPI growth rate (top panels) and the inflation rate (bottom panels) under alternative MIDAS models, produced in pseudo real-time for 1994:1 (left panels) and 2008:12 (right panels). The top panels display the predictive density for IPI growth under three alternative MIDAS models, where the blue solid line corresponds to the AR-MIDAS model, the red dashed line refers to the AR-MIDAS SV model, and the green dashed-dotted line refers to the AR-MIDAS SV-MIDAS model. The bottom panels display the predictive density for inflation using the same three MIDAS models. All MIDAS models use the daily Ffr series as predictor, and all predictive densities are produced using recursive model estimates.

Figure 4. Cumulative sum of log-score differentials (CLSD) for AR-MIDAS SV models



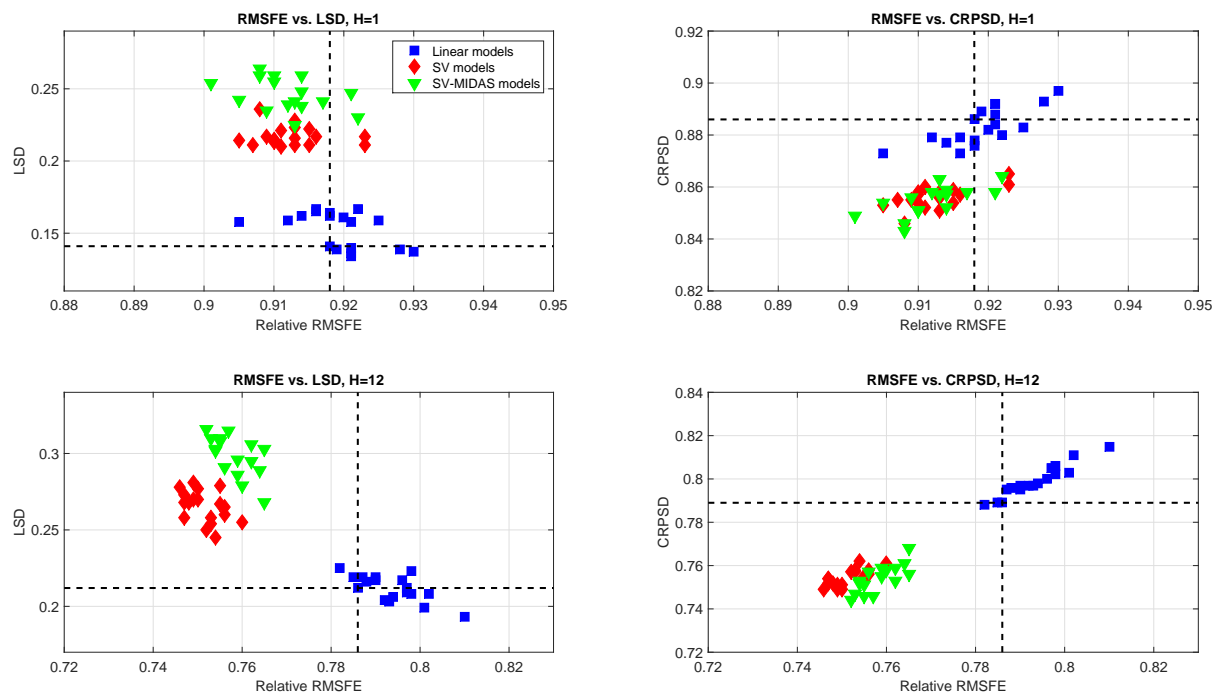
This figure shows the sum of log predictive scores from AR-MIDAS SV and AR SV models, computed relative the sum of log predictive scores of the Random Walk (RW) model. Each quarter we estimate the parameters of the forecast models recursively and generate one-step-ahead density forecasts of real GDP growth which are in turn used to compute log-predictive scores. This procedure is applied to the benchmark RW model as well as to all the alternative forecasting models. We then plot the cumulative sum of log-predictive scores (LS_t) for the alternative models computed relative to the cumulative sum of log-predictive scores of the RW model, $LS_t - LS_t^{RW}$. Values above zero indicate that a forecast model generates better performance than the RW benchmark, while negative values suggest the opposite. The top two panels compare the forecasting performance of different autoregressive models for IPI growth with time-varying volatility and different forecast horizons, while the bottom two panels focuses on Inflation.

Figure 5. Relation between RMSFE and other measures of predictive performance for IPI growth rate



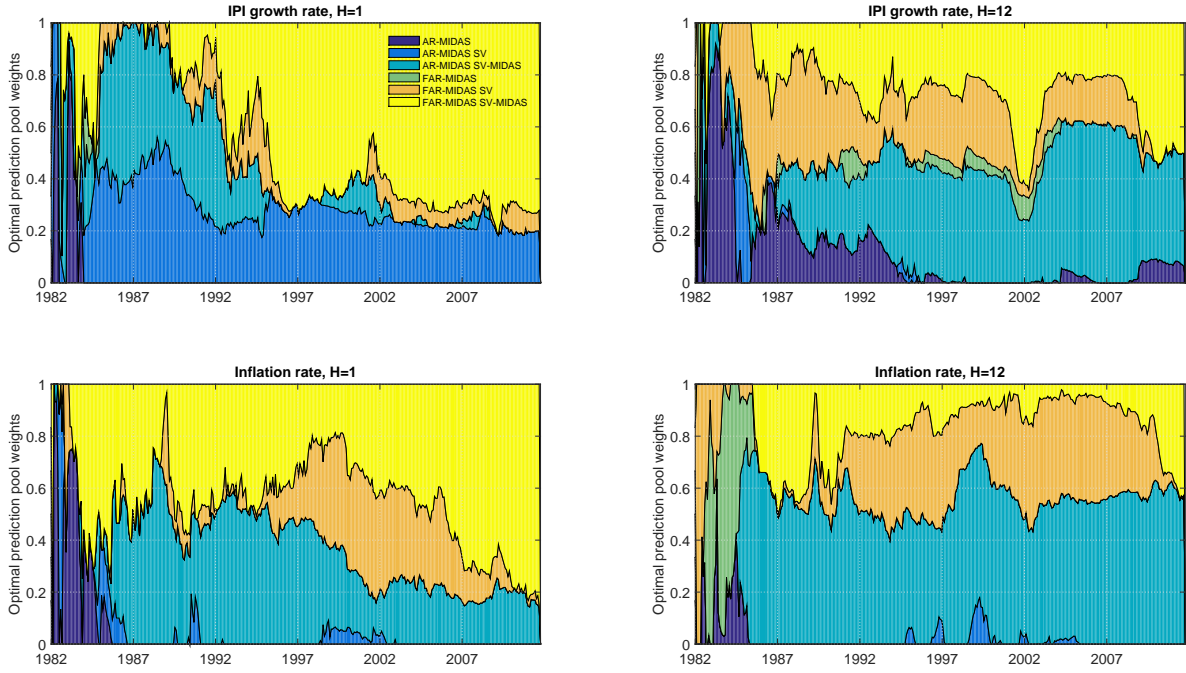
This figure presents scatter plots of out-of-sample predictive performance measures against the RMSFE (root mean squared forecast error) measure. All estimates of predictive performance are measured relative to the Random Walk (RW) model, so that the origin in each figure, shown as the intersection of the solid black lines, corresponds to the RW model. In addition, the intersection of the black dashed lines displays the relative out-of-sample performance of the AR model. The left panel displays the relation between RMSFE and the average log-score differential (LSD), while the right panel shows the relation between RMSFE and the continuously ranked probability score differential (CRPSD). Blue squares correspond to linear models (including AR, AR-MIDAS, FAR, and FAR-MIDAS models), red diamonds represent the stochastic volatility models (including AR SV, AR-MIDAS SV, FAR SV, and FAR-MIDAS SV), and the green triangles depict the MIDAS in volatility models (including AR-MIDAS SV-MIDAS and FAR-MIDAS SV-MIDAS). All measures of predictive performance are produced with recursive estimates of the models. The out-of-sample period starts in 1982:1 and ends in 2011:12.

Figure 6. Relation between RMSFE and other measures of predictive performance for Inflation rate



This figure presents scatter plots of out-of-sample predictive performance measures against the RMSFE (root mean squared forecast error) measure. All estimates of predictive performance are measured relative to the Random Walk (RW) model, so that the origin in each figure, shown as the intersection of the solid black lines, corresponds to the RW model. In addition, the intersection of the black dashed lines displays the relative out-of-sample performance of the AR model. The left panel displays the relation between RMSFE and the average log-score differential (LSD), while the right panel shows the relation between RMSFE and the continuously ranked probability score differential (CRPSD). Blue squares correspond to linear models (including AR, AR-MIDAS, FAR, and FAR-MIDAS models), red diamonds represent the stochastic volatility models (including AR SV, AR-MIDAS SV, FAR SV, and FAR-MIDAS SV), and the green triangles depict the MIDAS in volatility models (including AR-MIDAS SV-MIDAS and FAR-MIDAS SV-MIDAS). All measures of predictive performance are produced with recursive estimates of the models. The out-of-sample period starts in 1982:1 and ends in 2011:12.

Figure 7. Weights on different model classes in the optimal prediction pool

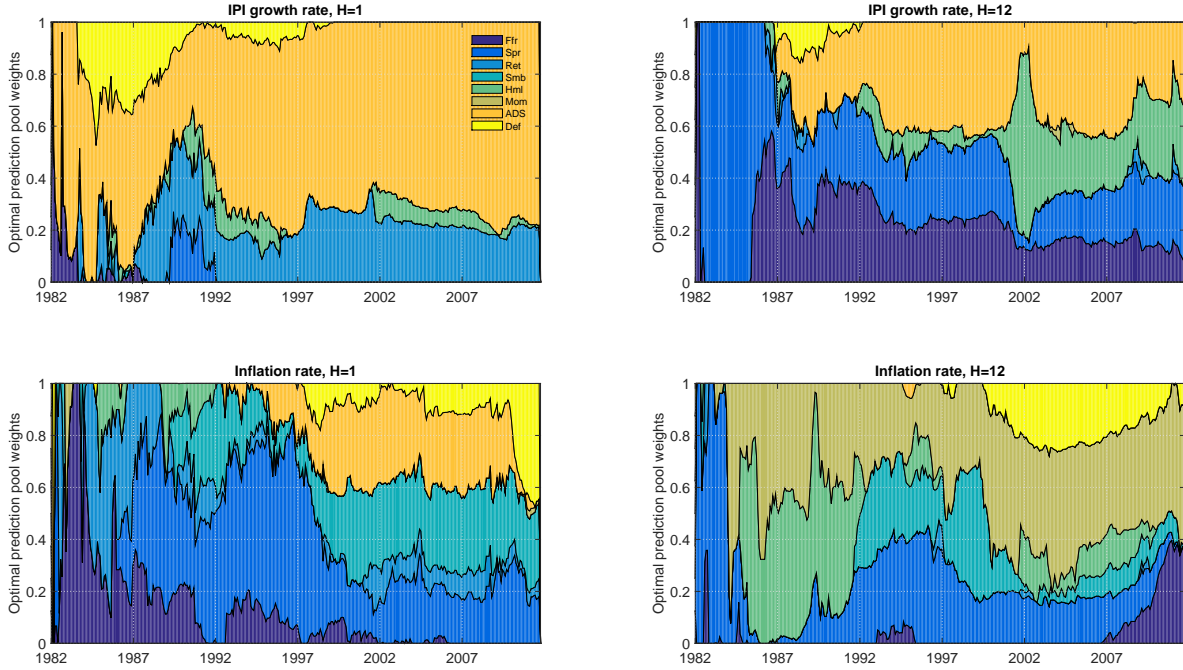


This figure plots the optimal weights on different models in the predictive pool, computed in real time by solving the minimization problem

$$\mathbf{w}_t^* = \arg \max_{\mathbf{w}} \sum_{\tau=1}^{t-1} \log \left[\sum_{i=1}^N w_i \times S_{\tau+1,i} \right]$$

where $N = 36$ is the number of models considered, and the solution is found subject to \mathbf{w}_t^* belonging to the N -dimensional unit simplex. $S_{\tau+1,i}$ denotes the time $\tau + 1$ recursively computed log score for model i , i.e. $S_{\tau+1,i} = \exp(LS_{\tau+1,i})$. Dark blue bars show the weights on the AR-MIDAS models in the optimal prediction pool, blue bars show the weights assigned to the AR-MIDAS SV models, and light blue bars show the weights assigned to the AR-MIDAS SV-MIDAS models; gray bars show the weights on the FAR-MIDAS models in the optimal prediction pool, orange bars show the weights assigned to the FAR-MIDAS SV models, and yellow bars show the weights assigned to the FAR-MIDAS SV-MIDAS models.

Figure 8. Weights on individual daily predictors in the optimal prediction pool



This figure plots the optimal weights on different daily predictors in the predictive pool, computed in real time by solving the minimization problem

$$\mathbf{w}_t^* = \arg \max_{\mathbf{w}} \sum_{\tau=1}^{t-1} \log \left[\sum_{i=1}^N w_i \times S_{\tau+1,i} \right]$$

where $N = 36$ is the number of models considered, and the solution is found subject to \mathbf{w}_t^* belonging to the N -dimensional unit simplex. $S_{\tau+1,i}$ denotes the time $\tau + 1$ recursively computed log score for model i , i.e. $S_{\tau+1,i} = \exp(LS_{\tau+1,i})$. Dark blue bars show the weights associated with the effective Federal Funds rate (Ffr) predictor in the optimal prediction pool, blue bars show the weights associated with the 10-year government bond rate and the federal fund rate (Spr), and light blue bars show the weights assigned to the value-weighted stock returns (Ret); celeste bars show the weights associated with the SML portfolio return (Smb), and green bars show the weights assigned to the HML portfolio return (Hml); gray, orange, and yellow bars show the weights associated with the MOM portfolio return (Mom), the Aruoba-Diebold-Scotti index (ADS), and the Default spread (Def)