



A Bayesian MIDAS Approach to Modeling First and Second Moment Dynamics

Davide Pettenuzzo, Economics Department, Brandeis University

Allan Timmermann, UCSD, CEPR, and CREATES

Roseen Valkanov, UCSD

Working Paper Series

A Bayesian MIDAS Approach to Modeling First and Second Moment Dynamics*

Davide Pettenuzzo

Brandeis University[†]

Allan Timmermann

UCSD, CEPR, and CREATES[‡]

Rossen Valkanov

UCSD[§]

July 24, 2014

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 and develop Gibbs sampling methods for Bayesian estimation in the presence of stochastic volatility dynamics. When applied to quarterly U.S. GDP growth data, we find strong evidence that models that feature MIDAS terms in the conditional volatility generate more accurate forecasts than conventional benchmarks. Finally, we find that forecast combination methods such as the optimal predictive pool of Geweke and Amisano (2011) produce consistent gains in out-of-sample predictive performance.

Key words: MIDAS regressions; Bayesian estimation; stochastic volatility; out-of-sample forecasts; GDP growth.

JEL classification: C53; C11; C32; E37

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

*We thank Elena Andreou, Eric Ghysels, and Sidney Ludvigson for making her macroeconomic variables available to us. We also thank Kenneth French for making his daily stock returns data available to us.

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

quarterly macroeconomic variables is appealing and has generated considerable academic interest in a rapidly expanding literature on mixed-data sampling (MIDAS) models.¹ 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 relative weighting on current and older values of predictors to be flexibly tailored to the data, the MIDAS approach makes it feasible to use the entire recent history of variables observed at a higher frequency than the outcome variable of interest.²

Empirical studies in the MIDAS literature have analyzed the dynamics in variables as diverse as GDP growth (Andreou et al. (2014), Carriero et al. (2012b), 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.

Benefits from using the MIDAS approach to introduce high frequency information in the conditional volatility specification, in addition to its level, have received much less attention. However, 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. First, studies such as Sims and Zha (2006) and Stock and Watson (2002) show that the volatility of macroeconomic variables varies over time. Second, accounting for dynamics in the volatility equation can lead to improved forecasts of levels as shown by Clark (2011), Carriero et al. (2012a), Carriero et al. (2012b), and Clark and Ravazzolo (2014). Third, an extensive literature summarized in Andersen et al. (2006) documents that high frequency variables can be used to produce better out-of-sample volatility forecasts.³

This paper proposes a new approach to forecasting macroeconomic variables such as GDP growth that uses the MIDAS approach to introduce daily variables in both the con-

¹Stock and Watson (2003) also use financial variables to forecast GDP growth, though measured at the monthly and quarterly frequencies.

²A 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.

³See also Ghysels and Valkanov (2012) for a review of the volatility MIDAS literature.

ditional mean and volatility equations while accounting for autoregressive and stochastic volatility dynamics. In doing so, our paper makes several contributions to the literature on forecasting with MIDAS models. First, we show how to cast the MIDAS lag polynomial as a linear regression model with transformed daily predictors and take advantage of this in a Bayesian estimation setting. In the case with constant volatility and normal innovations, Bayesian estimation can be undertaken using a two-block Gibbs sampler. Despite their growing popularity as models for forecasting and inference, relatively little work has been undertaken on Bayesian estimation and forecasting with MIDAS models, though notable exceptions include Carriero et al. (2012b), Ghysels (2012), Marcellino et al. (2013), and Rodriguez and Puggioni (2010).

Building on this insight we show, secondly, how to extend the MIDAS model to include stochastic volatility in the innovations of the outcome, along with autoregressive dynamics and lagged predictive factors. Third, we show how MIDAS effects can be extended from the first moment to the second moment in a parameterization that permits the log conditional volatility to depend linearly on a MIDAS term. Conditional on the sequence of log-volatilities and the parameters determining the stochastic volatility dynamics, our MIDAS specifications reduce to standard linear regression models and drawing the parameters of the mean equation from their posterior distributions becomes standard. 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. (2002) to allow for exogenous covariates in the volatility equation.⁴ Hence a four-block Gibbs sampler can be used to produce posterior estimates for the model parameters.

Our Bayesian MIDAS approach offers several advantages over standard classical methods. As the object of the Bayesian analysis is to obtain the full (posterior) predictive density given the data, as opposed to simply a point forecast, the predictive density forecasts account for the effect of parameter uncertainty. This can be important in empirical applications (including the one we consider here) with macroeconomic variables for which data samples are short and so parameters tend to be imprecisely estimated. Moreover, we can use the density forecasts to evaluate a whole range of measures of predictive accuracy including, as a special case, mean square forecast errors, log scores, and the continuously ranked probability score of Gneiting and Raftery (2007a). Finally, because we construct the predictive density for a range of different models, we can compute forecast combinations that optimally weigh the individual models. As new data arrive, the combination weights get updated and so models that start to perform better receive greater weight in the combination. Forecast combinations also provide a way to deal with model uncertainty

⁴Posterior simulation of the whole path of stochastic volatilities under an arbitrary second moment MIDAS lag polynomial would require the use of a particle filter.

since they do not depend on identifying a single best model.

We illustrate our approach in an empirical application to quarterly U.S. GDP growth, a variable that has been extensively studied in the literature. Andreou et al. (2014) use the MIDAS approach to predict quarterly U.S. GDP growth in the context of a large cross-section of daily predictors. They find that adding information on daily predictors improves the predictive accuracy of a model that also includes common factors and combines forecasts from a range of univariate prediction models. Compared with Andreou et al. (2014), our application uses a longer time series but a much smaller set of daily stock returns and interest rates. Even so, we find that the MIDAS models are capable of producing notable gains in predictive accuracy compared with models that use only quarterly information. The predictive accuracy of stochastic volatility models with MIDAS terms is particularly evident in turbulent periods because of efficiency gains in the parameters of the conditional mean. We also find that forecast combinations, particularly Bayesian Model Averaging and the optimal prediction pool of Geweke and Amisano (2011), lead to significant improvements in forecast accuracy. These gains are observed both for the point forecasts (through reduced RMSFE values) and for density forecasts (through higher log-scores). The results are established out-of-sample, using only predictors available in real time and so do not suffer from look-ahead biases.

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 application to quarterly U.S. GDP growth, while Section 5 covers different forecast combinations. Section 6 concludes.

2 MIDAS regression models

This section outlines how we generalize the conventional regression specification to account for MIDAS effects in the first and second conditional moment equations while 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, $x_{t+1}^{(m)}$, are observed m times between t and $t + 1$. For example, y_{t+1} could be a quarterly variable and $x_{t+1}^{(m)}$

could be a daily variable. In this case $m = (22 \times 3)$, assuming that the number of daily observations available within a month is constant and equals 22. 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 model

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

where

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

and $L^{k/m}$ is a lag operator such that $L^{1/m} x_{\tau}^{(m)} = x_{\tau-1/m}^{(m)}$, and $\varepsilon_{\tau+1}$ is i.i.d. with $E(\varepsilon_{\tau+1}) = 0$ and $Var(\varepsilon_{\tau+1}) = \sigma_{\varepsilon}^2$. The distinguishing feature of MIDAS models is that the lag coefficients in $B(k; \boldsymbol{\theta})$ are parametrized 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 quarterly series which gets affected by four quarters' worth of daily data, $x_t^{(m)}$. In this case, we would need $K = 264$ ($22 \times 3 \times 4$) lags of daily variables. Without any restrictions on the parameters in $B(L^{1/m}; \boldsymbol{\theta})$ there would be $264 + 2$ parameters to estimate in 1. By making $B(L^{1/m}; \boldsymbol{\theta})$ a function of a small set of parameters $p + 1 \ll K$ we can greatly reduce the number of parameters to estimate.

It is sometimes useful to cast the MIDAS model as

$$y_{\tau+1} = \beta_0 + \beta_1 B_1(L^{1/m}; \boldsymbol{\theta}_1) x_{\tau}^{(m)} + \varepsilon_{\tau+1}, \quad \tau = 1, \dots, t-1 \quad (2)$$

where $\beta_1 B_1(L^{1/m}; \boldsymbol{\theta}_1) = B(L^{1/m}; \boldsymbol{\theta})$ and β_1 is a scalar that captures the overall impact of lagged values of $x_{\tau}^{(m)}$ on $y_{\tau+1}$. Since β_1 enters multiplicatively in (2), it cannot be identified without imposing further restrictions on the polynomial $B_1(L^{1/m}; \boldsymbol{\theta}_1)$. One way to identify β_1 is to normalize the function $B_1(L^{1/m}; \boldsymbol{\theta}_1)$ to sum to unity. Normalization 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.

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

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

This regression requires the estimation of $(3 + p + p_y + p_z \times r)$ coefficients. The distributed lag term $\sum_{j=0}^{p_y-1} \rho_{j+1} y_{\tau-j}$ captures same-frequency dynamics in y_{t+1} , 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. The FAR-MIDAS model is called FADL-MIDAS (for factor augmented distributed lag MIDAS) in Andreou et al. (2014).

2.2 MIDAS weighting functions

The functional form of the MIDAS weights $B(L^{1/m}; \boldsymbol{\theta})$ 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. Since our principal interest lies in forecasting GDP growth, we adopt a simple unrestricted version of $B(L^{1/m}; \boldsymbol{\theta})$, 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 to be estimated. Under this parameterization, (3) takes the form

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

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

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

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

$$y_{\tau+1} = \alpha + \sum_{j=0}^{p_y-1} \rho_{j+1} y_{\tau-j} + \sum_{j=0}^{q_z-1} \gamma'_{j+1} \mathbf{z}_{\tau-j} + \boldsymbol{\theta}' \tilde{\mathbf{X}}_\tau^{(m)} + \varepsilon_{\tau+1}, \quad (8)$$

where $\tilde{\mathbf{X}}_\tau^{(m)} = \mathbf{Q} \mathbf{X}_\tau^{(m)}$ 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, however. 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 $\tilde{\mathbf{X}}_\tau^{(m)}$.

It is useful to briefly contrast the Almond weights in (4) with other parameterizations in the MIDAS literature. These include the exponential Almon lag

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

(Ghysels et al. (2005), Andreou et al. (2014)) 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 these thresholds 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 the second moment

The majority of MIDAS specifications assume constant volatility of the residuals in (8). Notable exceptions are Carriero et al. (2012b), and Marcellino et al. (2013) who allow for

stochastic volatility in $\varepsilon_{\tau+1}$. We extend this literature by introducing a MIDAS component not only in the level of $y_{\tau+1}$ but also in its conditional volatility. 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

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

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

$$h_{\tau+1} = h_{\tau} + \xi_{\tau+1}, \quad (11)$$

where $\xi_{\tau+1} \sim \mathcal{N}(0, \sigma_{\xi}^2)$ 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. (2012b), and Marcellino et al. (2013), but with a different parameterization of the MIDAS weights.⁵

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

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

The daily variables $v_{\tau}^{(m)}$ need not be the same as those entering the first moment in (10). 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.⁶ The stochastic volatility MIDAS specification is an analogue to the MIDAS specification in the mean equation (3).

⁵The 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.

⁶We further restrict the AR(1) coefficient λ_1 to lie within the unit circle, i.e., $\lambda_1 \in [-1, 1]$. The addition of exogenous covariates in the log-volatility equation has been studied by Chib et al. (2002).

To complete the model in (12), we need to specify the SV-MIDAS weights, $B(k; \boldsymbol{\theta}_h)$, 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

$$h_{\tau+1} = \lambda_0 + \lambda_1 h_\tau + \boldsymbol{\theta}'_h \tilde{\mathbf{X}}_\tau^{(m)} + \xi_{\tau+1}, \quad \xi_{\tau+1} \sim \mathcal{N}(0, \sigma_\xi^2). \quad (13)$$

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

3 Bayesian estimation and forecasting

Most work on estimation of MIDAS regression models has used frequentist methods carried out using either OLS (when the MIDAS polynomials can be reparameterized as a linear model) or NLS. Relatively less work uses Bayesian methods to estimate the MIDAS polynomials, though notable exceptions include Carriero et al. (2012b), Ghysels (2012), Marcellino et al. (2013), and Rodriguez and Puggioni (2010). Carriero et al. (2012b) develop a Bayesian method for producing current-quarter forecasts of GDP growth that is closely related to the U-MIDAS approach proposed by Foroni 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. He develops approaches for both general MIDAS polynomial specifications and for the Almon lag polynomial specification, the step function polynomial specification of Ghysels et al. (2007), and the U-MIDAS approach of Foroni et al. (2013). 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. They cast their model in a Bayesian framework and derive a Gibbs sampler to estimate the model parameters. 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. To deal with the overparameterization arising from the need to estimate a potentially very large number of coefficients, they introduce a stochastic search variable selection (SSVS) step that allows the data to determine which of the high frequency lags should enter the model.

3.1 MIDAS models with constant volatility

Let $\boldsymbol{\Phi}$ denote the regression parameters in the constant volatility MIDAS model (3), excluding the MIDAS coefficients $\boldsymbol{\theta}$, i.e., $\boldsymbol{\Phi} = (\alpha, \rho_1, \dots, \rho_{p_y}, \gamma'_1, \dots, \gamma'_{q_z})'$. Conditional on $\boldsymbol{\theta}$,

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 $\varepsilon_{\tau+1}$, σ_ε^2 . Assuming standard independent Normal-inverted gamma priors on Φ and σ_ε^2 and normally distributed residuals, $\varepsilon_{\tau+1}$, 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 $\tilde{\mathbf{X}}_t^{(m)}$ 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 Forni et al. (2013). Assuming that $\varepsilon_{\tau+1}$ 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 .⁷

To see how this works, rewrite (8) as

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

where $\Psi = (\Phi', \theta)'$ and $\mathbf{Z}_\tau = (1, y_\tau, \dots, y_{\tau-p_y+1}, \mathbf{z}'_\tau, \dots, \mathbf{z}'_{\tau-q_z+1}, \tilde{\mathbf{X}}_\tau^{(m)'})'$. Following standard practice, suppose that the priors for the regression parameters Ψ in (14) are normally distributed and independent of σ_ε^2 ⁸

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

All elements of \underline{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{b} reflects the view that the best model for predicting real GDP growth is a random walk. We choose a data-based prior for \underline{V} :⁹

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

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 Ψ . Our analysis sets $\underline{\psi} = 25$, corresponding to relatively diffuse priors.

⁷Under the U-MIDAS approach of Forni et al. (2013), the matrix of transformed regressors, $\tilde{\mathbf{X}}_t^{(m)}$, is the same as the original matrix, $\mathbf{X}_t^{(m)}$.

⁸See for example Koop (2003), section 4.2.

⁹Priors 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.

For the constant volatility model we assume a standard gamma prior for the error precision of the return innovation, σ_ε^{-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 the degree of informativeness of this prior. $\underline{v}_0 \rightarrow 0$ corresponds to a diffuse prior on σ_ε^{-2} .¹⁰ Our baseline analysis sets $\underline{v}_0 = 0.005$, again representing an uninformative choice as it corresponds to a pre-sample of half of one percent of the data sample.

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

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

and

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

where

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

and

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

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

For more general MIDAS lag polynomials, obtaining posterior estimates for θ 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 $\theta = (\theta_1, \theta_2)'$, and suggests using a Gamma prior for both θ_1 and θ_2 since under the beta parametrization both of these parameters take on only positive values

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

¹⁰Following Koop (2003), we adopt the Gamma distribution parametrization of Poirier (1995). Namely, if the continuous random variable Y has a Gamma distribution with mean $\mu > 0$ and degrees of freedom $v > 0$, we write $Y \sim G(\mu, v)$. Then, in this case, $E(Y) = \mu$ and $Var(Y) = 2\mu^2/v$.

f_0 and F_0 are hyperparameters controlling the mean and degrees of freedom of the Gamma distribution. For example, setting $f_0 = 1$ and $F_0 = 10$ corresponds to a flat weighting scheme that puts equal weight on all high frequency lags. 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 $\boldsymbol{\theta}$. In the case of beta weights, Ghysels (2012) suggests a Gamma distribution as a suitable choice for $q(\boldsymbol{\theta}^* | \boldsymbol{\theta}^{[i]})$. Hence, 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), $\boldsymbol{\Psi} = (\boldsymbol{\Phi}', \boldsymbol{\theta}')'$, the sequence of log volatilities, $h^t = \{h_1, h_2, \dots, h_t\}$, the parameter vector $(\lambda_0, \lambda_1, \boldsymbol{\theta}_h)$, and the log-volatility variance σ_ξ^2 .

For the parameters in the mean equation, $\boldsymbol{\Psi}$, we follow the earlier choice of priors

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

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

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

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

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

with $h_{\tau+1} | \lambda_0, \lambda_1, \boldsymbol{\theta}_h, h_\tau, \sigma_\xi^{-2} \sim \mathcal{N}(\lambda_0 + \lambda_1 h_\tau + \boldsymbol{\theta}'_h \tilde{\mathbf{X}}_\tau^{(m)}, \sigma_\xi^2)$. Thus, to complete the prior elicitation for $p(h^t, \lambda_0, \lambda_1, \boldsymbol{\theta}_h, \sigma_\xi^{-2})$, 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 as follows

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

$$\begin{bmatrix} \lambda_0 \\ \lambda_1 \\ \boldsymbol{\theta}_h \end{bmatrix} \sim \mathcal{N}(\underline{m}_h, \underline{V}_h), \quad \lambda_1 \in (-1, 1), \quad (29)$$

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 $\underline{k}_h = 0.1$. Compared to the earlier choices of priors, these are more informative priors. The choice of $\underline{k}_\xi = 0.01$, in conjunction with $\underline{v}_\xi = 1$, 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.

As for the hyperparameters \underline{m}_h and \underline{V}_h in (29), following Clark and Ravazzolo (2014) we set the prior mean and standard deviation of the intercept and the MIDAS coefficients to 0 and 0.5, respectively, corresponding to uninformative priors on the intercept of the log volatility specification. Finally, we set the prior mean of the AR(1) coefficient, λ_1 , to 0.9 with a standard deviation of 0.01. This represents a more informative prior that matches persistent dynamics in the log volatility process.

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

1. $p(\boldsymbol{\Psi} | h^t, \lambda_0, \lambda_1, \boldsymbol{\theta}_h, \sigma_\xi^{-2}, \mathcal{D}^t)$.
2. $p(h^t | \boldsymbol{\Psi}, \lambda_0, \lambda_1, \boldsymbol{\theta}_h, \sigma_\xi^{-2}, \mathcal{D}^t)$.

$$3. p(\sigma_\xi^{-2} | \Psi, h^t, \lambda_0, \lambda_1, \boldsymbol{\theta}_h, \mathcal{D}^t)$$

$$4. p(\lambda_0, \lambda_1, \boldsymbol{\theta}_h | \Psi, h^t, \sigma_\xi^{-2}, \mathcal{D}^t)$$

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. As for the fourth step, note that the conditional posterior distribution of the SV-MIDAS parameters $p(\lambda_0, \lambda_1, \boldsymbol{\theta}_h | \Psi, h^t, \sigma_\xi^{-2}, \mathcal{D}^t)$ can be expressed as

$$\lambda_0, \lambda_1, \boldsymbol{\theta}_h | \Psi, h^t, \sigma_\xi^{-2}, \mathcal{D}^t \sim \mathcal{N}(\bar{m}_h, \bar{V}_h) \times \lambda_1 \in (-1, 1),$$

where

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

and

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

3.3 Forecasts from MIDAS models

The object of Bayesian estimation of the 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+1} is given by

$$p(y_{t+1} | \mathcal{D}^t) = \int_{\Psi, \sigma_\xi^{-2}} p(y_{t+1} | \Psi, \sigma_\xi^{-2}, \mathcal{D}^t) p(\Psi, \sigma_\xi^{-2} | \mathcal{D}^t) d\Psi d\sigma_\xi^{-2}, \quad (33)$$

where $p(\Psi, \sigma_\xi^{-2} | \mathcal{D}^t)$ 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+1} is given by

$$\begin{aligned} p(y_{t+1} | \mathcal{D}^t) &= \int_{\Psi, h^{t+1}, \lambda_0, \lambda_1, \boldsymbol{\theta}_h, \sigma_\xi^{-2}} p(y_{t+1} | \Psi, h_{t+1}, h^t, \lambda_0, \lambda_1, \boldsymbol{\theta}_h, \sigma_\xi^{-2}, \mathcal{D}^t) \\ &\quad \times p(h_{t+1} | \Psi, h^t, \lambda_0, \lambda_1, \boldsymbol{\theta}_h, \sigma_\xi^{-2}, \mathcal{D}^t) \\ &\quad \times p(\Psi, h^t, \lambda_0, \lambda_1, \boldsymbol{\theta}_h, \sigma_\xi^{-2} | \mathcal{D}^t) d\Psi dh^{t+1} d\lambda_0 d\lambda_1 d\boldsymbol{\theta}_h d\sigma_\xi^{-2}. \end{aligned} \quad (34)$$

We can use the Gibbs sampler to obtain draws from the predictive densities in (33) and (34). These draws, $y_{t+1|t}^{(j)}$, $j = 1, \dots, J$ can be used to compute objects such as point

forecasts, $\hat{y}_{t+1|t} = J^{-1} \sum_{j=1}^J y_{t+1|t}^{(j)}$ or the quantile of the realized value of the predicted variable, $J^{-1} \sum_{j=1}^J I(y_{t+1} \leq y_{t+1|t}^{(j)})$, where $I(y_{t+1} \leq y_{t+1|t}^{(j)})$ is an indicator function that equals one if the outcome, y_{t+1} , falls below the j th draw from the Gibbs sampler.

4 Empirical Results

This section introduces the quarterly data of U.S. real GDP growth, a set of macroeconomic factors, and the daily predictors. We then analyze the full-sample and out-of-sample predictive accuracy of GDP growth forecasts for the model specifications described in sections 2 and 3. We apply a range of measures to evaluate the predictive accuracy of our GDP growth 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 quarterly data of U.S. real GDP growth along with a set of monthly macro variables and daily financial series. Real GDP growth (denoted $y_{\tau+1}$) is measured as the annualized log change in real GDP and is obtained from IHS Global Insight. The monthly variables are an updated version of the 132 macroeconomic series used in Ludvigson and Ng (2009), 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.¹¹ We follow Stock and Watson (2012) and Andreou et al. (2014) and extract two common factors (\mathbf{z}_τ) from the 132 macroeconomic series using principal components, after first taking quarterly averages of the monthly series. Hence, the real GDP growth rate and the macro factors are both observable at quarterly frequency.¹²

The daily financial series considered in this study, $x_\tau^{(m)}$, are (i) the effective Federal Funds rate (Ffr); (ii) the interest rate spread between the 10-year government bond rate

¹¹We thank Sydney Ludvigson for making this data available on her website, at <http://www.econ.nyu.edu/user/ludvigsons/jlndata.zip>

¹²Entering the macro factors in the MIDAS specification at the monthly frequency would complicate the estimation as it would involve three different frequencies (quarterly, monthly, and daily) and two different MIDAS polynomials. We therefore use quarterly macro factors.

and the federal fund rate (Spr); (iii) value-weighted returns on all 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); and (vi) returns on a winner minus loser momentum spread portfolio (Mom).¹³ The interest rate series are from the Federal Reserve of St. Louis database FRED and are transformed to eliminate trends. 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.¹⁴

All our variables span the period from 1962:I to 2011:IV. This is a considerably longer time period than the one covered in Andreou et al. (2014), but we have a much smaller set of daily predictors. To facilitate the comparison of our results with theirs, we also show the key results for their sample, 2001:I to 2008:IV.

We test whether a better full-sample fit and out-of-sample forecasts can be obtained by allowing the daily stock market return series and interest rates to impact real GDP growth forecasts through MIDAS polynomials. To this end, we estimate several versions of the MIDAS specifications discussed above. These fall into one of three categories: (i) MIDAS in the 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) of quarterly GDP growth.

For each of these specifications, we consider two versions which allow us to test the contribution of the MIDAS components. Specifically, we estimate (8) with AR lags and a MIDAS component (AR-MIDAS) as well as with AR lags, factors, and MIDAS components (FAR-MIDAS). We estimate the same two versions for model (10) and (11), labeled the AR-MIDAS SV and FAR-MIDAS SV, respectively, and for (10) and (13), labeled the AR-MIDAS SV-MIDAS and FAR-MIDAS SV-MIDAS models, respectively. In total we consider six MIDAS specifications.

Every MIDAS specification is estimated with one of the six daily variables, Ffr, Spr, Ret, Smb, Hml, and Mom, as defined above. Therefore, we have a total of 36 MIDAS forecasting models. In addition, we estimate four non-MIDAS models: a purely autoregressive model of quarterly GDP growth (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).

Our analysis assumes four AR lags of quarterly GDP growth ($p_y = 4$), two lags of

¹³Using 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.

¹⁴We thank Kenneth French for making this data available on his website, at http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html

the macro factors ($p_z = 2$) and uses four quarters of past daily observations. The full estimation sample covers the period 1962:I to 2011:IV.

4.2 Full sample estimates and model comparisons

We first compare the fit of the different model specifications over the full sample, 1962-2011. 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 in favor of M_1 is not compelling. 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 evidence”, respectively, in support of M_1 versus the null, M_0 .

Table 1 shows pairwise model comparisons based on the transformed Bayes factors, $2\ln(B_{1,0})$. Panel A displays results for the AR-MIDAS and FAR-MIDAS models (8) relative to AR or FAR models, respectively. The MIDAS models are estimated by including a single daily predictor at a time, each shown in separate columns. In essence, this is the comparison that Andreou et al. (2014) conduct using a frequentist approach, a different set of factors and MIDAS predictors, and different MIDAS polynomial parameterization. In the first row of Panel A, MIDAS models with Ffr and Ret as predictors produce strong evidence that daily stock returns and interest rates improve the prediction model. The second row of Panel A provides weak evidence that a MIDAS model that includes quarterly information on factors and daily information on Ffr produces a better fit than the benchmark FAR model. Conversely, Panel B suggests that once SV dynamics is allowed in both the baseline and alternative models, there is little evidence suggesting that the MIDAS-in-mean models provide a better fit than models without such MIDAS effects.

Panel C of Table 1 reports Bayes factors for the specifications with MIDAS effects in the volatility of GDP growth relative to non-MIDAS stochastic volatility 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). Both rows in the panel suggest that adding MIDAS effects to the stochastic volatility equation leads to significant improvements in the model fit. Indeed, for all daily predictors the values of $2\ln(B_{1,0})$ exceed ten. This is true relative to both the AR-MIDAS SV and the FAR-MIDAS SV benchmarks. Finally, Panel D shows that adding MIDAS effects to both the conditional

mean and conditional volatility equations improves the performance relative to models with no MIDAS effects, but reduces the Bayes factors compared with models that only include MIDAS effects in the conditional volatility equation (Panel C).

Overall, the time series evidence supports the claim that the fit of the GDP forecasting models improve when we use MIDAS polynomials to introduce information from daily financial variables. Moreover, our results suggest two takeaways. First, contrasting the results in panels A and B, it is clear that the importance of MIDAS-in-mean effects depend on how dynamics in second moments are modeled. Second, we find strong benefits from including MIDAS effects in the conditional volatility equation even after accounting for stochastic volatility dynamics. We next address whether these results carry over to an out-of-sample setting.

4.3 Out-of-sample forecasts with a single MIDAS predictor

In our 1962:I to 2011:IV sample, we use the first twenty years of data as an initial training sample, i.e., we estimate our regression models over the period 1962:I-1981:IV and use the resulting estimates to predict real GDP growth for 1982:I. Next, we include 1982:I in the estimation sample, which thus becomes 1962:I-1982:I, and use the corresponding estimates to predict GDP growth for 1982:II. We proceed progressively 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:I to 2011:IV. To allow for direct comparisons with the results in Andreou et al. (2014), we also consider a shorter forecast evaluation period, 2001:I–2008:IV.

4.3.1 Point forecasts

First we consider point forecasts. For each of the MIDAS models we obtain point forecasts by repeatedly drawing from the predictive densities, $p(y_\tau | M_i, \mathcal{D}^{\tau-1})$, 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 1982:I to 2011:IV. Following Stock and Watson (2003) and Andreou et al. (2014), we measure the predictive performance of the MIDAS models relative to the random walk (RW) model of GDP growth. 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

$$RMSFE_i = \frac{\sqrt{\frac{1}{\bar{t}-\underline{t}+1} \sum_{\tau=\underline{t}}^{\bar{t}} e_{i,\tau}^2}}{\sqrt{\frac{1}{\bar{t}-\underline{t}+1} \sum_{\tau=\underline{t}}^{\bar{t}} e_{RW,\tau}^2}}, \quad (35)$$

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.

In the top panel of Figure 1 we plot the sequence of point forecasts generated by AR-MIDAS, AR-MIDAS SV, and AR-MIDAS SV-MIDAS models as well as RW forecasts. The RW forecasts are quite different and notably more volatile from those of the three other models, which in turn are very similar. The volatility forecasts of the MIDAS models, in the bottom panel of Figure 1, are quite different, however. Not surprisingly, the AR-MIDAS SV and AR-MIDAS SV-MIDAS models are better able to capture not only the highly volatile periods of the early 1980s and the 2007-2009 financial crisis, but also periods of moderate volatility.

Table 2 displays results for the RMSFE ratio in (35). The smallest RMSFE value across the financial variables is displayed in bold. Panel A displays the results for the MIDAS models with constant volatility. The simple AR model produces an RMSFE ratio of 0.868, thus bettering the RW model's forecasting performance by 13%. The AR-MIDAS model with Ret as a daily predictor does slightly better, generating an RMSFE ratio of 0.841. The other daily variables do not improve on the predictive accuracy, however. Adding information on the quarterly factors greatly improves the results. Indeed, the FAR and FAR-MIDAS models obtain RMSFE ratios of 0.775 and 0.766, respectively, the latter being generated with a model that includes Ffr. For MIDAS-in-mean models without stochastic volatility, there is only weak evidence that daily stock returns or interest rates lead to better forecasts.

Panel B of Table 2 features results for the models that allow for stochastic volatility and MIDAS effects in the conditional volatility. The RMSFE value for the AR-SV model is 0.875 which is only slightly lower than the corresponding value for the simple AR model. Somewhat lower RMSFE values of about 0.849 are obtained by the AR-MIDAS SV and AR-MIDAS SV-MIDAS models when Ret is used as the daily predictor. As in Panel A, we obtain better results by including factors in the predictive regressions. For example, the RMSFE is 0.772 for the Mom predictor included in the FAR-MIDAS SV model, as compared with 0.784 for the FAR-SV model.

To see how the precision of the point forecasts evolves over time, we compute the Cumulative Sum of Squared prediction Error Differences (CSSED) introduced by Welch and Goyal (2008)

$$CSSED_{i,t} = \sum_{\tau=\underline{t}}^t (e_{RW,\tau}^2 - e_{i,\tau}^2). \quad (36)$$

Increases in the value of $CSSED_{i,t}$ indicate that model i generates more accurate point forecasts than the RW model, while decreasing values suggest the opposite. If the two models are equally good, the $CSSED$ measure should hover close to zero.

Figure 2 plots the evolution in the cumulative sum of squared prediction error differences for various models and their MIDAS counterparts. The figure suggests a fairly stable pattern in the gains in predictive accuracy obtained by the MIDAS models and their simpler counterparts, measured relative to the RW model.

4.3.2 Density forecasts

One limitation to 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 predictive density $p(y_{t+1} | M_i, \mathcal{D}^t)$. Indeed, comparing the plots of the point forecasts and the volatility forecasts in Figure 1, it is clear that there are much greater differences between the the volatility forecasts generated by the different models. Figure 3 shows that such differences produce very different density forecasts, because of the efficiency gains in the estimation of the conditional mean parameters. Compared to AR-MIDAS or FAR-MIDAS models, for the snapshots shown in this figure, SV dynamics tends to compress the predictive distribution. SV-MIDAS has a similar, but weaker effect, preserving some of the greater uncertainty associated with the constant volatility forecasts.

To address this issue, we consider two measures of predictive performance. First, following Amisano and Giacomini (2007), Geweke and Amisano (2010), and Hall and Mitchell (2007), we consider the average log-score (LS) differential

$$LSD_i = \sum_{\tau=\bar{t}}^{\bar{t}} (LS_{i,\tau} - LS_{RW,\tau}), \quad (37)$$

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.

Panels A and B in Table 3 report results for the log-score measure for the constant and stochastic volatility models, respectively. In Panel A, the AR model obtains an LSD of 0.164. Adding the MIDAS forecasting variables one by one increases the precision for the best models. In particular, the LSD for the MIDAS model based on daily stock returns (Ret) is 0.204. Introducing the macroeconomic factors leads to an even greater improvement with an LSD of 0.310. The FAR-MIDAS model that uses the daily federal funds rate (Ffr) generates an LSD value of 0.317.

The results in Panel B are even more encouraging. Using MIDAS polynomials to allow stock returns (Ret) to affect the first and second moments of the AR SV model increases

the LSD value from 0.312 to 0.364. Similarly, using daily information on the federal funds rate (Ffr) increases the LSD value of the FAR SV model from 0.347 to 0.434. Comparing these numbers to the LSD values for the models that only allow for MIDAS effects in the conditional mean, we see that the addition of MIDAS terms in the conditional volatility of GDP growth generates large improvements in predictive accuracy.

To see how the log score differential evolves over time, as in (36) we compute the cumulative log score differential for model i versus the RW model

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

Increasing values in $CLSD_{i,t}$ suggest that model i produces more accurate density forecasts than the RW model.

Figure 4 shows these cumulative log score differentials for a variety of models with and without MIDAS terms. This type of plot can help diagnose patterns in relative predictive accuracy, i.e., if a single episode is responsible for most of the forecast gains or losses or if we see more continual improvements in the forecasts. In most cases, and notably for the FAR-SV model compared to the FAR-MIDAS SV-MIDAS model, it is clear that the MIDAS density forecasts dominate on a consistent basis.

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

$$CRPSD_i = \frac{\frac{1}{\bar{t}-\underline{t}+1} \sum_{\tau=\underline{t}}^{\bar{t}} CRPS_{i,\tau}}{\frac{1}{\bar{t}-\underline{t}+1} \sum_{\tau=\underline{t}}^{\bar{t}} CRPS_{RW,\tau}}. \quad (39)$$

$CRPS_{i,\tau}$ ($CRPS_{RW,\tau}$) 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 $CRPSD_m$ suggest that model i performs better than the benchmark RW model. Gneiting 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, but not equal, to the realization.

Table 3 shows result for the CRPSD statistic in Panels C (constant volatility models) and D (time varying volatility models). If a model's CRPSD is smaller than one, its empirical CDF is closer to that of the data than is achieved by the RW model. The CRPSD value of the AR model in Panel C is 0.845, which implies an improvement of 15% over the RW model. The addition of MIDAS terms in the mean increases the predictive

accuracy somewhat. The CRPSD value for the AR-MIDAS model with Ret is 0.822. Greater improvements relative to the RW benchmark are observed for the FAR and FAR-MIDAS models. The FAR-MIDAS model with Ffr as a predictor lowers the CRPSD to 0.740. Including macroeconomic factors in the constant volatility models greatly increases the predictive accuracy of the GDP growth models, while including MIDAS effects from either Ffr or Ret reduce the CRPSD values by an additional 1-2 percent.

Turning to the models with time varying volatility, the CRPSD of the AR-SV model is 21% lower than the RW model (Panel D). The AR-MIDAS SV model produces a CRPSD value of 0.779 when based on the Ret variable. A larger improvement is observed for the AR-MIDAS SV-MIDAS model that uses Ret and has a CRPSD value of 0.767. The FAR models continue to produce better out-of-sample forecasts than the AR models. The FAR-SV model has a CRPSD value of 0.730. This compares with a CRPSD value of 0.706 for the FAR-MIDAS SV-MIDAS model that includes daily information on the federal funds rate.

Figure 5 illustrates how the ranking of models varies across the RMSFE, LSD, and CRPSD criteria. 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 the three classes of models according to the RMSFE criterion. Thus, the benefits from using MIDAS models to incorporate daily information show up much more strongly in the density forecast measures than in point forecasts. The right window in Figure 5 shows that there is 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.

To test more formally if the MIDAS forecasts are more accurate than various competitors, Table 4 reports Diebold-Mariano p -values under the null that a given model has the same predictive ability as the benchmark. Panels A-C show p -values for the RMSFE, LSD, and CRPSD measures, respectively. Each panel reports results for a range of models against the null of the RW or AR specifications which Stock and Watson (2003) have shown are hard to beat in the case of GDP growth forecasts. Panel A shows that all models perform better than the RW model with p -values between 0.085 and 0.000. The AR model is much harder to beat according to the RMSFE metric. Only the MIDAS specifications that also include the macroeconomic factors as predictors produce p -values below 0.10 for most of the daily variables.

The following picture emerges from Tables 3 and 4. Consistent with Andreou et al. (2014) we find that information on macroeconomic factors leads to sizable improvements in the precision of out-of-sample GDP growth forecasts. For some of the daily predictors such as the federal funds rate and stock market returns, MIDAS-in-mean effects can further improve the out-of-sample forecasts of GDP growth. The addition of stochastic volatility dynamics further increases the out-of-sample predictive performance of the models, supporting the results of Clark (2011). Finally, the second moment SV-MIDAS models enhance the forecasting performance improvements relative to both the constant volatility MIDAS models and to the simple SV models.

4.3.3 Subsample performance

Andreou et al. (2014) use MIDAS-in-mean models to predict quarterly U.S. GDP growth using a large cross-section of daily financial variables. Due to limitations on data availability, they use a shorter out-of-sample period, 2001-2008. To make our results directly comparable to theirs, Table 5 shows results for the same subsample. First consider the RMSFE results reported in Panels A and B. Like Andreou et al. (2014), we find that the AR model improves marginally on the performance of the RW model, while the FAR model provides a notably larger improvement with an RMSFE ratio, measured relative to the RW model, of 0.83. The best MIDAS model which uses daily stock market returns (Ret) as a predictor generates reductions of around 6% in the RMSFE ratio relative to the corresponding models that assume constant volatility.

The best MIDAS models continue to generate more accurate point forecasts for the models that allow for stochastic volatility (Panel B). For example, relative to an AR-SV model, the MIDAS model that includes daily stock returns reduce the RMSFE ratio by 6%, while relative to the FAR-SV model we see improvements of about 3%. Interestingly, for the 2001-2008 subsample the first-moment MIDAS effects capture most of the improvements in forecast accuracy, while the second-moment MIDAS effects account for much less of the improvements than we found in the longer out-of-sample experiment which used the sample 1982-2008.

Our finding that first-moment MIDAS effects help improve the predictive accuracy of (point) forecasts during the sample 2001-2008 is consistent with the findings reported by Andreou et al. (2014). The reason second-moment MIDAS effects are not so important during the 2001-2008 period is that the volatility was quite low and stable during this period, with exception of the last two quarters of 2008 (see Figure 1).

5 Forecast combinations

So far our analysis covers multiple specifications that differ in terms of the identity of the daily MIDAS variables, inclusion of macro factors, as well as the 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

$$p(y_{t+1}|\mathcal{D}^t) = \frac{1}{N} \sum_{i=1}^N p(y_{t+1}|M_i, \mathcal{D}^t). \quad (40)$$

Next, consider the optimal predictive pool proposed by Geweke and Amisano (2011),

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

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 that relies only on information available up to time t ,

$$\mathbf{w}_t^* = \arg \max_{\mathbf{w}} \sum_{\tau=1}^{t-1} \log \left[\sum_{i=1}^N w_i \times S_{\tau+1,i} \right], \quad (42)$$

subject to $\mathbf{w}_t^* \in [0, 1]^N$. $S_{\tau+1,i} = \exp(LS_{\tau+1,i})$ is the recursively computed log-score for model i at time $\tau + 1$. As $t \rightarrow \infty$ the weights chosen according to (42) minimize the Kullback-Leibler distance between the data generating process D and the combined density. Finally, we consider Bayesian model averaging (BMA) which weighs the individual models by their posterior probabilities

$$p(y_{t+1}|\mathcal{D}^t) = \sum_{i=1}^N \Pr(M_i|\mathcal{D}^t) p(y_{t+1}|M_i, \mathcal{D}^t). \quad (43)$$

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

$$\Pr(M_i|\mathcal{D}^t) = \frac{\Pr(\mathcal{D}^t|M_i) \Pr(M_i)}{\sum_{j=1}^N \Pr(\mathcal{D}^t|M_j) \Pr(M_j)}. \quad (44)$$

$\Pr(\mathcal{D}^t | M_i)$ 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 $\Pr(M_i) = 1/N$. We follow Geweke and Amisano (2010) and compute the marginal likelihoods by cumulating the predictive log scores of each model, after conditioning on the initial warm-up estimation sample, as

$$\Pr(\mathcal{D}^t | M_i) = \exp\left(\sum_{\tau=t}^t LS_{\tau,i}\right). \quad (45)$$

The idea of using forecast combination methods in MIDAS regressions is also explored by Andreou et al. (2014) in the context of predicting quarterly GDP growth using a host of daily financial variables. We note, however, that Andreou et al. (2014) combine point forecasts from the individual MIDAS whereas our approach combines density forecasts.

5.2 Empirical results for forecast combinations

Table 6 reports out-of-sample values of our three measures of predictive accuracy for the full sample (1982-2011, Panel A) and the shorter subsample (2001-2008, Panel B). All combination schemes produce large improvements to the RMSFE results for the benchmark RW model, reducing the RMSFE value by 22% or more. The predictive pool of Geweke and Amisano (2011) produces the best results with an RMSFE ratio of 0.742. This ratio is lower than that obtained by any of the individual models shown in Table 2. For the shorter subsample, 2001-2008, the combinations again perform well. The equal-weighted combination is worst with an RMSFE ratio of 0.867, while the BMA approach is best, producing an RMSFE ratio of 0.778, comparable to that of the best individual models shown in Table 5.

The log-score differential values obtained by the combination schemes are also impressively large. In the long sample, the predictive pool obtains an LSD value of 0.512, considerably higher than that obtained by all of the individual models in Table 3, the largest of which is 0.434. In the shorter sample, the BMA approach obtains an LSD value of 0.413, again higher than that of any of the individual models in Table 5.

The results for the CRPSD values are equally good. At 0.664, the CRPSD value obtained by the predictive pool is lower than that of any of the individual models in Table 3. A similar conclusion holds for the shorter subsample, although the lowest CRPSD value (0.701) is now obtained for the BMA approach.

Figure 6 plots the recursive performance of the forecast combinations, measured relative to the RW model. The top figure plots the cumulative sum of squared prediction error difference, while the bottom plot shows the cumulative sum of log-score differentials.

Both plots show that the dominance of the forecast combinations over the RW benchmark is not confined to specific points in time, but occurs throughout the entire sample.

To gain insights into how this improved performance is obtained, Figure 7 provides a recursive plot of the weights assigned to different classes of models. Between 1984 and 1986 the FAR-MIDAS SV model receives full weight, only to receive no weight subsequently. The FAR-MIDAS SV-MIDAS model receives by far the greatest weight during most of the sample, close to 70% for the last 20 years of the sample. The remaining weight goes mainly to the AR-MIDAS SV-MIDAS model, although the AR-MIDAS SV model also receives some weight, mainly during the 1990s.

Turning to the probability mass assigned to models containing the six different daily predictor variables in the MIDAS models, Figure 8 shows that the federal funds rate (Ffr) models receive all the weight between 1985 and 1992, but that the weight on MIDAS models with this predictor drifts down to around 35 percent at the end of the sample. The remaining probability weight is assigned to MIDAS models that include daily stock returns (10-30%), the momentum variable (up to 30%, with a particularly large weight at the end of the sample), and the Spr and Smb variables.

In summary, these results show clear advantages, both in the full sample and in the shorter subsample, from using forecast combinations as a way to deal with model uncertainty. 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 manage to adaptively adjust to such changes and produce better out-of-sample forecasts.

6 Conclusions

We develop a Bayesian approach to estimation of forecasting models that allow 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.

In an empirical application to quarterly U.S. GDP growth we find that the approach produces notably more accurate out-of-sample forecasts than both random walk and autoregressive models. The benefits from considering MIDAS effects in the second moment of GDP growth are, unsurprisingly, larger in the sample that includes periods of particularly high volatility and become more modest during the Great Moderation. Conversely, MIDAS effects in the first moment help improve the forecasting performance during the Great Moderation. The gains from adding MIDAS effects in volatility are larger for measures

of predictive accuracy that consider the full predictive density as opposed to conventional measures, such as mean squared forecast error, that only look at point forecasts. Finally, when we exploit the posterior probabilities to optimally combine forecasting models and handle model uncertainty, we find sizeable gains in predictive accuracy.

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 (2014). Should macroeconomic forecasters use daily financial data and how? *Journal of Business Economics and Statistics*, forthcoming.
- Carriero, A., T. E. Clark, and M. Marcellino (2012a). Common drifting volatility in large bayesian vars. CEPR Discussion Papers 8894.
- Carriero, A., T. E. Clark, and M. Marcellino (2012b). Real-time nowcasting with a bayesian mixed frequency model with stochastic volatility. Federal Reserve Bank of Cleveland Working Paper 12-27.
- 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. Galvao (2009). Forecasting us output growth using leading indicators: an appraisal using midas models. *Journal of Applied Econometrics* 24(7), 1187–1206.
- Corsi, F. (2009). A simple approximate long-memory model of realized volatility. *Journal of Financial Econometrics* 7(2), 174–196.

- 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.
- Foroni, 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)*, forthcoming.
- 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. New York University Working Paper.
- 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.
- Koop, G. (2003). *Bayesian Econometrics*. John Wiley & Sons, Ltd.
- 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.
- Poirier, D. J. (1995). *Intermediate Statistics and Econometrics: A Comparative Approach*. MIT Press.
- 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 (2012). Generalized shrinkage methods for forecasting using many predictors. *Journal of Business and Economic Statistics* 30(4), 481–493.
- Welch, I. and A. Goyal (2008). A comprehensive look at the empirical performance of equity premium prediction. *Review of Financial Studies* 21(4), 1455–1508.

Table 1. Bayes factors

M_1 vs. M_0	Ffr	Spr	Ret	Smb	Hml	Mom
<i>Panel A: linear models, MIDAS in mean</i>						
AR-MIDAS vs. AR	8.227	1.235	9.705	-0.240	-1.128	-0.865
FAR-MIDAS vs. FAR	1.277	-2.640	-2.353	-3.908	-2.005	-2.550
<i>Panel B: SV models, MIDAS in mean</i>						
AR-MIDAS SV vs. AR SV	-0.837	-15.724	-4.029	-8.064	-13.720	-7.703
FAR-MIDAS SV vs. FAR SV	-0.157	-8.117	-14.992	-10.556	-15.788	-7.532
<i>Panel C: SV models, MIDAS in volatility</i>						
AR-MIDAS SV-MIDAS vs. AR-MIDAS SV	12.392	17.346	17.638	15.264	23.382	19.265
FAR-MIDAS SV-MIDAS vs. FAR-MIDAS SV	20.762	21.642	25.402	24.252	25.605	26.417
<i>Panel D: SV models, MIDAS in both mean and volatility</i>						
AR-MIDAS SV-MIDAS vs. AR SV	11.556	1.622	13.609	7.201	9.662	11.563
FAR-MIDAS SV-MIDAS vs. FAR SV	20.605	13.524	10.409	13.695	9.816	18.885

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

$$B_{1,0} = \frac{Pr(\mathcal{D}^t | M_1)}{Pr(\mathcal{D}^t | M_0)}$$

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), and the MOM portfolio return (Mom). Kass and Raftery (1995) suggest interpreting the results as follows: when $2 \times (\ln B_{1,0})$ is lower than 0, the evidence in favor of M_1 is not compelling. For $2 \times (\ln B_{1,0})$ between 0 and 2, there is “weak evidence” that M_1 is a more likely characterization of the data than M_0 . For $2 \times (\ln B_{1,0})$ between 2 and 6, 6 and 10, and higher than 10, we view that as “some evidence,” “strong evidence,” and “very strong evidence” in the data supporting M_1 relative to M_0 .

Table 2. Out-of-sample forecast performance - RMSFE

Model	–	Ffr	Spr	Ret	Smb	Hml	Mom
<i>Panel A: linear models</i>							
AR	0.868						
AR-MIDAS		0.869	0.926	0.841	0.878	0.876	0.877
FAR	0.775						
FAR-MIDAS		0.766	0.783	0.778	0.785	0.787	0.768
<i>Panel B: SV models</i>							
AR SV	0.875						
AR-MIDAS SV		0.878	0.906	0.849	0.887	0.891	0.900
AR-MIDAS SV-MIDAS		0.881	0.911	0.848	0.883	0.887	0.888
FAR SV	0.784						
FAR-MIDAS SV		0.784	0.809	0.804	0.803	0.790	0.771
FAR-MIDAS SV-MIDAS		0.780	0.801	0.795	0.797	0.790	0.772

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

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

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), and the MOM portfolio return (Mom). All forecasts and forecast errors are produced with recursive estimates of the models. The out-of-sample period starts in 1982:I and ends in 2011:IV.

Table 3. Out-of-sample forecast performance - Log score differentials (LSD) and CRPS differentials (CRPSD)

Model	–	Ffr	Spr	Ret	Smb	Hml	Mom
<i>Panel A: LSD, linear models</i>							
AR	0.164						
AR-MIDAS		0.193	0.167	0.204	0.163	0.161	0.159
FAR	0.310						
FAR-MIDAS		0.317	0.298	0.302	0.294	0.303	0.302
<i>Panel B: LSD, SV models</i>							
AR SV	0.312						
AR-MIDAS SV		0.299	0.241	0.293	0.277	0.256	0.276
AR-MIDAS SV-MIDAS		0.348	0.310	0.364	0.339	0.352	0.354
FAR SV	0.347						
FAR-MIDAS SV		0.350	0.310	0.279	0.300	0.281	0.318
FAR-MIDAS SV-MIDAS		0.434	0.399	0.384	0.397	0.384	0.423
<i>Panel C: CRPSD, linear models</i>							
AR	0.845						
AR-MIDAS		0.826	0.871	0.822	0.849	0.849	0.848
FAR	0.751						
FAR-MIDAS		0.740	0.759	0.757	0.760	0.758	0.748
<i>Panel D: CRPSD, SV models</i>							
AR SV	0.790						
AR-MIDAS SV		0.786	0.822	0.779	0.802	0.810	0.801
AR-MIDAS SV-MIDAS		0.775	0.812	0.767	0.787	0.788	0.779
FAR SV	0.730						
FAR-MIDAS SV		0.720	0.749	0.754	0.748	0.744	0.722
FAR-MIDAS SV-MIDAS		0.706	0.728	0.736	0.730	0.731	0.709

Panels A and B report the average log-score (LS) differential, $LSD_i = \frac{1}{\bar{t}-\underline{t}+1} \sum_{\tau=\underline{t}}^{\bar{t}} (LS_{i,\tau} - LS_{RW,\tau})$, 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. Panels C and D report the average continuously ranked probability score differential (CRPSD) of model i relative to the RW model,

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

where $CRPS_{i,\tau}$ ($CRPS_{RW,\tau}$) measures the average distance between the empirical cumulative distribution function (CDF) of 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 perform better than the benchmark 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), and the MOM portfolio return (Mom). All forecasts and forecast errors are produced with recursive estimates of the models. The out-of-sample period starts in 1982:I and ends in 2011:IV.

Table 4. Diebold-Mariano p-values from tests of equal predictive ability

Comparison	Ffr	Spr	Ret	Smb	Hml	Mom
<i>Panel A: RMSFE</i>						
AR-MIDAS vs. RW	0.008	0.085	0.000	0.005	0.006	0.001
FAR-MIDAS vs. RW	0.000	0.000	0.001	0.000	0.001	0.000
AR-MIDAS SV vs. RW	0.008	0.025	0.000	0.007	0.009	0.010
AR-MIDAS SV-MIDAS vs. RW	0.009	0.034	0.000	0.006	0.010	0.005
FAR-MIDAS SV vs. RW	0.001	0.000	0.001	0.000	0.001	0.000
FAR-MIDAS SV-MIDAS vs. RW	0.000	0.000	0.001	0.000	0.001	0.000
AR-MIDAS vs. AR	0.294	0.929	0.220	0.726	0.788	0.605
FAR-MIDAS vs. AR	0.043	0.067	0.110	0.078	0.124	0.061
AR-MIDAS SV vs. AR	0.296	0.821	0.133	0.755	0.930	0.642
AR-MIDAS SV-MIDAS vs. AR	0.330	0.844	0.164	0.677	0.868	0.545
FAR-MIDAS SV vs. AR	0.067	0.082	0.136	0.092	0.095	0.065
FAR-MIDAS SV-MIDAS vs. AR	0.057	0.067	0.127	0.081	0.108	0.063
<i>Panel B: LSD</i>						
AR-MIDAS vs. RW	0.000	0.000	0.000	0.000	0.000	0.000
FAR-MIDAS vs. RW	0.000	0.000	0.000	0.000	0.000	0.000
AR-MIDAS SV vs. RW	0.000	0.003	0.000	0.000	0.002	0.001
AR-MIDAS SV-MIDAS vs. RW	0.000	0.000	0.000	0.000	0.000	0.000
FAR-MIDAS SV vs. RW	0.000	0.000	0.000	0.000	0.000	0.000
FAR-MIDAS SV-MIDAS vs. RW	0.000	0.000	0.000	0.000	0.000	0.000
AR-MIDAS vs. AR	0.049	0.400	0.010	0.541	0.722	0.713
FAR-MIDAS vs. AR	0.000	0.000	0.000	0.000	0.000	0.000
AR-MIDAS SV vs. AR	0.011	0.120	0.026	0.053	0.120	0.047
AR-MIDAS SV-MIDAS vs. AR	0.001	0.003	0.000	0.000	0.000	0.000
FAR-MIDAS SV vs. AR	0.008	0.013	0.061	0.025	0.054	0.033
FAR-MIDAS SV-MIDAS vs. AR	0.000	0.000	0.000	0.000	0.000	0.000
<i>Panel C: CRPSD</i>						
AR-MIDAS vs. RW	0.000	0.001	0.000	0.000	0.000	0.000
FAR-MIDAS vs. RW	0.000	0.000	0.000	0.000	0.000	0.000
AR-MIDAS SV vs. RW	0.000	0.000	0.000	0.000	0.000	0.000
AR-MIDAS SV-MIDAS vs. RW	0.000	0.000	0.000	0.000	0.000	0.000
FAR-MIDAS SV vs. RW	0.000	0.000	0.000	0.000	0.000	0.000
FAR-MIDAS SV-MIDAS vs. RW	0.000	0.000	0.000	0.000	0.000	0.000
AR-MIDAS vs. AR	0.088	0.823	0.124	0.630	0.711	0.510
FAR-MIDAS vs. AR	0.002	0.006	0.019	0.007	0.018	0.004
AR-MIDAS SV vs. AR	0.004	0.133	0.003	0.031	0.092	0.014
AR-MIDAS SV-MIDAS vs. AR	0.001	0.061	0.001	0.003	0.004	0.000
FAR-MIDAS SV vs. AR	0.002	0.004	0.016	0.005	0.009	0.004
FAR-MIDAS SV-MIDAS vs. AR	0.001	0.001	0.007	0.001	0.004	0.001

This table reports the Diebold-Mariano p-values under the null that a given model has the same predictive ability than the benchmark model. P-values are computed for the RMSFE statistic (panel A), the LSD statistic (panel B), and the CRPSD statistic (panel C), both against the Random Walk (RW) benchmark and the first-order autoregressive (AR) benchmark. 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), the value-weighted return (Ret), the SML portfolio return (Smb), the HML portfolio return (Hml), and the MOM portfolio return (Mom). 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).

Table 5. Out-of-sample forecast performance - subsample (2001-2008)

Model	–	Ffr	Spr	Ret	Smb	Hml	Mom
<i>Panel A: RMSFE, linear models</i>							
AR	0.987						
AR-MIDAS		1.038	1.068	0.918	1.015	0.947	0.941
FAR	0.830						
FAR-MIDAS		0.828	0.792	0.779	0.860	0.810	0.799
<i>Panel B: RMSFE, SV models</i>							
AR SV	1.003						
AR-MIDAS SV		1.045	1.029	0.941	1.029	0.990	0.966
AR-MIDAS SV-MIDAS		1.048	1.038	0.938	1.029	0.983	0.958
FAR SV	0.800						
FAR-MIDAS SV		0.827	0.789	0.773	0.829	0.782	0.780
FAR-MIDAS SV-MIDAS		0.821	0.787	0.769	0.836	0.785	0.778
<i>Panel C: LSD, linear models</i>							
AR	0.098						
AR-MIDAS		0.083	0.060	0.160	0.085	0.122	0.121
FAR	0.272						
FAR-MIDAS		0.279	0.296	0.308	0.246	0.286	0.286
<i>Panel D: LSD, SV models</i>							
AR SV	0.175						
AR-MIDAS SV		0.144	0.119	0.180	0.137	0.174	0.175
AR-MIDAS SV-MIDAS		0.142	0.206	0.238	0.177	0.283	0.267
FAR SV	0.296						
FAR-MIDAS SV		0.280	0.289	0.299	0.249	0.277	0.278
FAR-MIDAS SV-MIDAS		0.334	0.369	0.383	0.292	0.372	0.370
<i>Panel E: CRPSD, linear models</i>							
AR	0.901						
AR-MIDAS		0.914	0.950	0.863	0.917	0.865	0.860
FAR	0.782						
FAR-MIDAS		0.771	0.757	0.746	0.808	0.762	0.749
<i>Panel F: CRPSD, SV models</i>							
AR SV	0.876						
AR-MIDAS SV		0.895	0.903	0.844	0.903	0.857	0.827
AR-MIDAS SV-MIDAS		0.889	0.891	0.839	0.893	0.831	0.806
FAR SV	0.755						
FAR-MIDAS SV		0.761	0.752	0.732	0.787	0.739	0.723
FAR-MIDAS SV-MIDAS		0.749	0.736	0.718	0.784	0.725	0.706

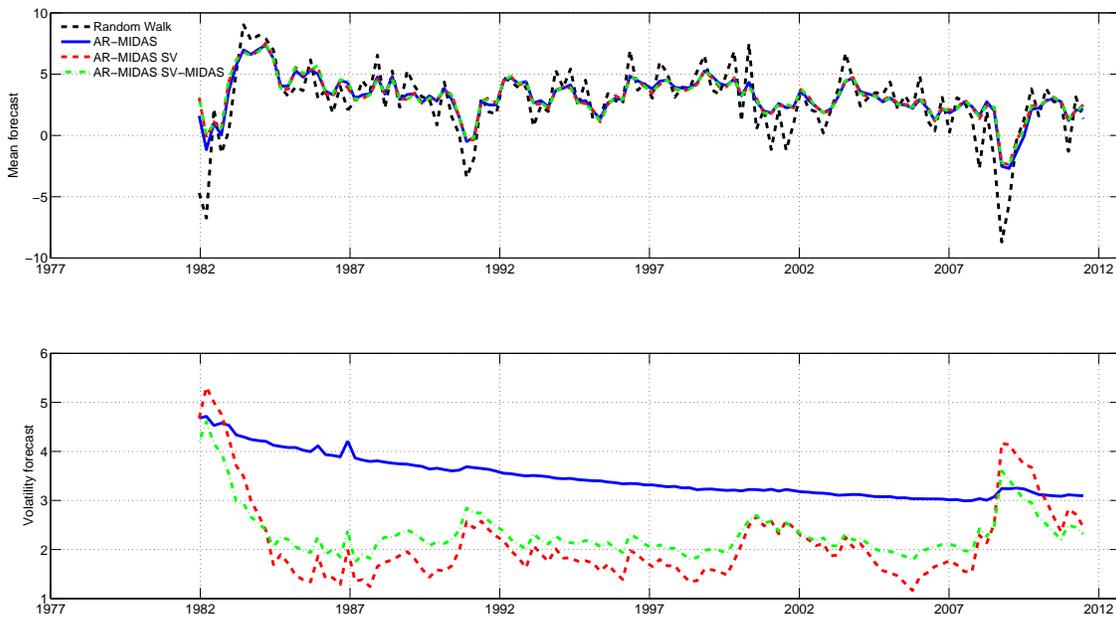
Panels A and B report $RMSFE_i$, the ratio between the RMSFE of model i and the RMSFE of the Random Walk (RW) model, where i denotes any of the models described in section 3. Panels C and D report LSD_i , the average log-score differential between model i and the RW model, while panels E and F report $CRPSD_i$, the average continuously ranked probability score differential of model i relative to the RW model. Values less than one for $RMSFE_i$ indicate that model i produces more accurate point forecasts than the RW model; positive values of LSD_i indicate that model i produces more accurate density forecasts than the RW model; values less than one for $CRPSD_i$ suggest that model i performs better than the benchmark 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), and the MOM portfolio return (Mom). All forecasts and forecast errors are produced with recursive estimation of the models. The out-of-sample period starts in 2001:I and ends in 2008:IV.

Table 6. Out-of-sample forecast performance - model combinations

Model	<i>Panel A: Forecast accuracy, 1982 - 2011</i>			<i>Panel B: Forecast accuracy, 2001 - 2008</i>		
	RMSFE	LSD	CRPSD	RMSFE	LSD	CRPSD
Equal-weighted combination	0.782	0.402	0.714	0.867	0.328	0.763
Bayesian model averaging	0.758	0.477	0.682	0.778	0.413	0.701
Optimal prediction pool	0.742	0.519	0.664	0.793	0.393	0.710
Comparison	<i>Panel C: DM p-values, 1982 - 2011</i>			<i>Panel D: DM p-values, 2001 - 2008</i>		
	RMSFE	LSD	CRPSD	RMSFE	LSD	CRPSD
Equal-weighted combination vs. RW	0.000	0.000	0.000	0.097	0.021	0.005
Bayesian model averaging vs. RW	0.000	0.000	0.000	0.082	0.014	0.008
Optimal prediction pool vs. RW	0.000	0.000	0.000	0.052	0.020	0.004
Equal-weighted combination vs. AR	0.006	0.000	0.000	0.070	0.003	0.006
Bayesian model averaging vs. AR	0.028	0.000	0.000	0.084	0.003	0.016
Optimal prediction pool vs. AR	0.011	0.000	0.000	0.072	0.004	0.009

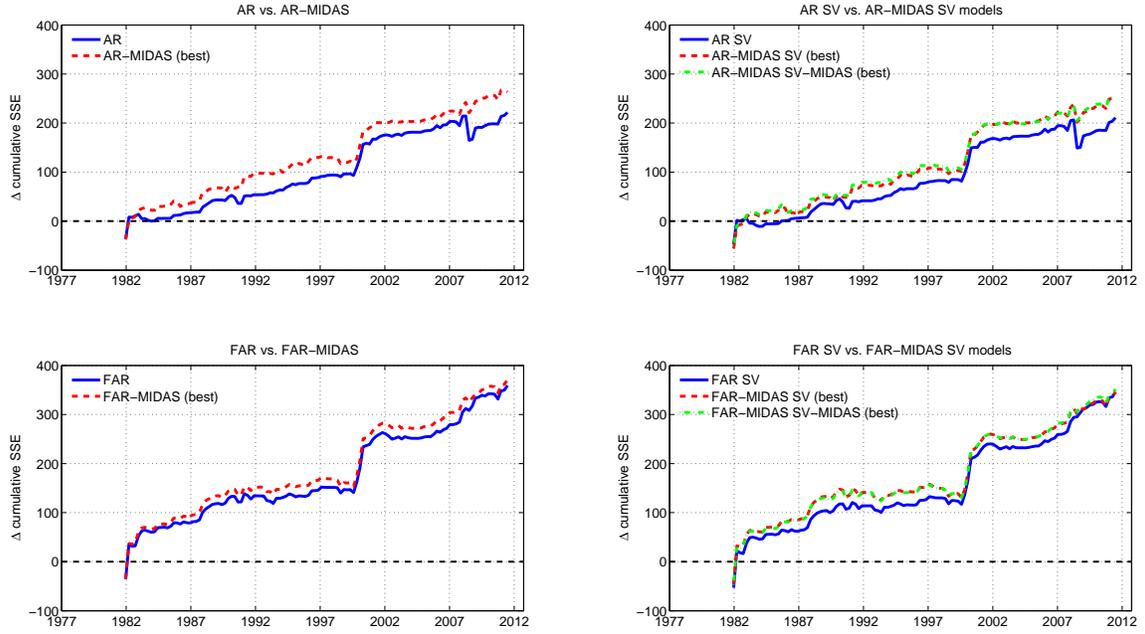
This table reports out-of-sample results for the optimal predictive pool of Geweke and Amisano (2011), an equal-weighted model combination scheme, and Bayesian Model Averaging applied to 36 MIDAS forecasting models that use different daily predictors and make different choices regarding the inclusion 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:I to 2011:IV (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:I to 2011:IV (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).

Figure 1. Mean and volatility forecasts of MIDAS models



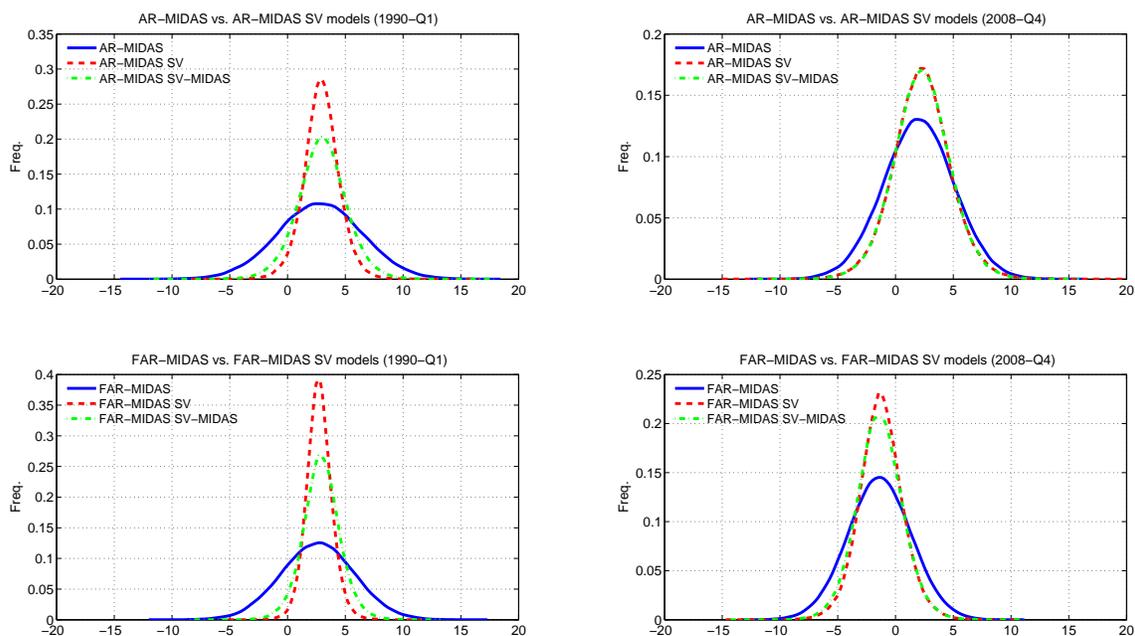
This figure shows the recursive conditional mean forecasts (top panel) and volatility forecasts (bottom panel) computed using the Random Walk model (dashed black line) as well as the predictive distributions of the AR-MIDAS (solid blue line), the AR-MIDAS SV (dashed red line), and the AR-MIDAS SV-MIDAS (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:I and ends in 2011:IV.

Figure 2. Cumulative Sum of Squared prediction Error Difference (CSSED)



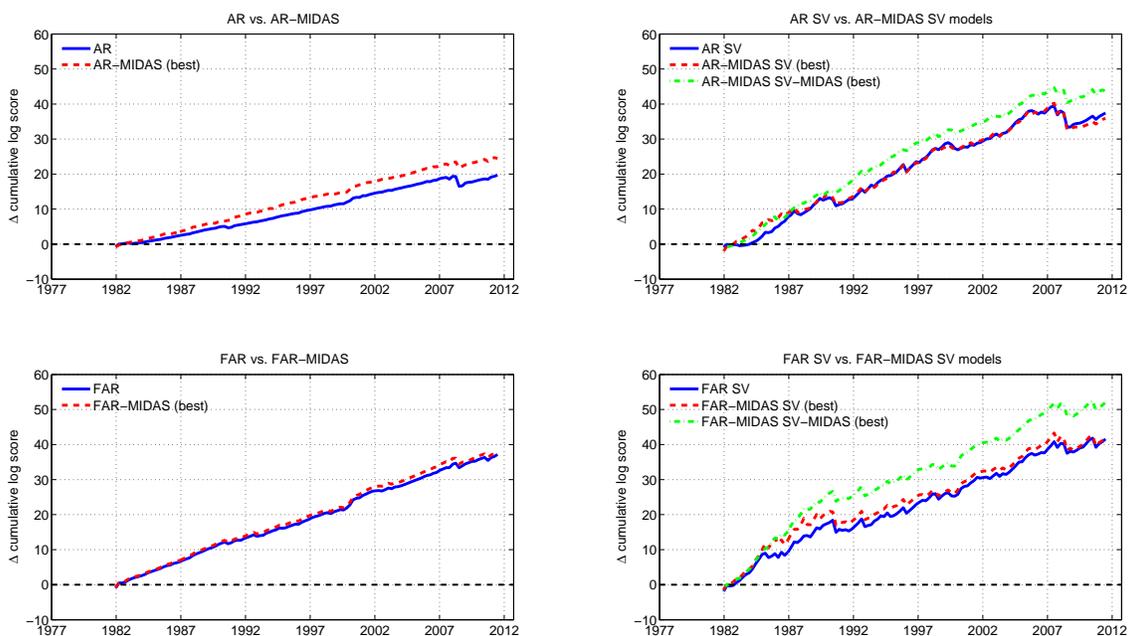
This figure shows the sum of squared forecast errors of the Random Walk model (RW) model minus the sum of squared forecast errors of alternative MIDAS models. Each month we estimate the parameters of the forecasting models recursively and generate one-step-ahead forecasts of real GDP growth rate which are in turn used to compute out-of-sample forecast errors. This procedure is applied to the RW model, which is our benchmark, as well as to all the alternative forecasting models. We then plot the cumulative sum of squared forecast errors (SSE_t) of the RW forecasts (SSE_t^{RW}) relative to the alternative forecasts, $SSE_t^{RW} - SSE_t$. Values above zero indicate that a forecasting 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 with constant volatility (first panel) or time-varying volatility (second panel), while the bottom two panels display the forecasting performance of different factor-augmented autoregressive models with constant volatility (third panel) or time-varying volatility (fourth panel).

Figure 3. Predictive density of real GDP growth under different MIDAS models



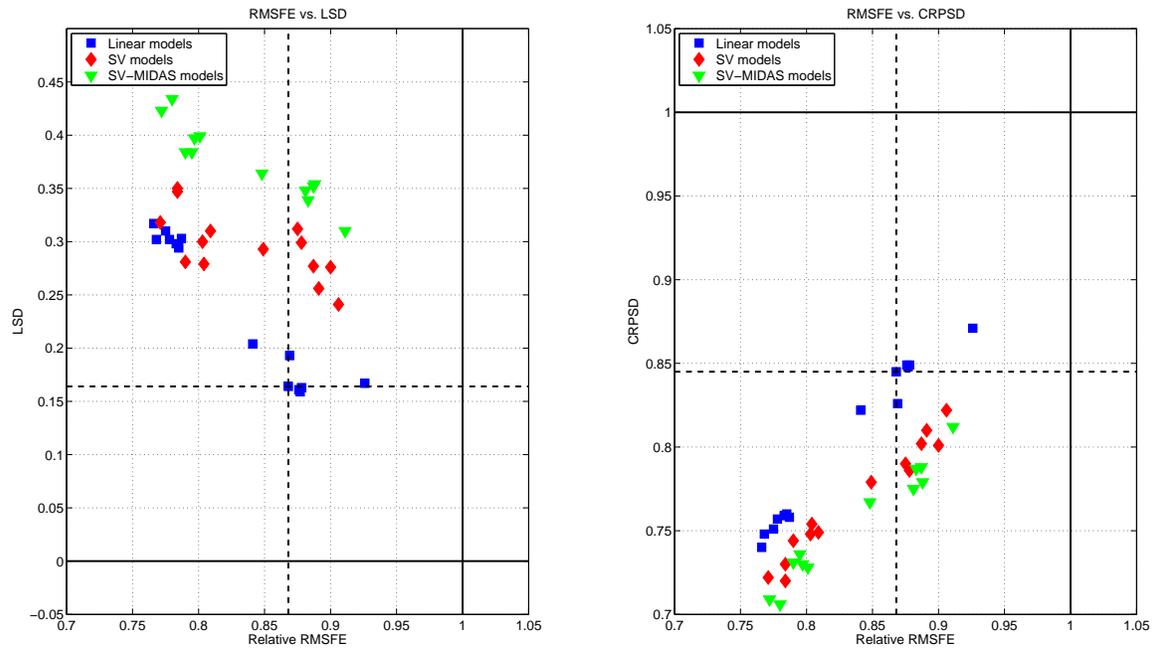
This figure shows the predictive density of the real GDP growth under alternative MIDAS models, produced in pseudo real-time for 1990:1 (left panels) and 2008:4 (right panels). The top panels display the real GDP growth predictive density under three alternative MIDAS models, where the blue solid line 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 real GDP growth predictive density under three alternative factor augmented MIDAS models, where the blue solid line corresponds to the FAR-MIDAS model, the red dashed line refers to the FAR-MIDAS SV model, and the green dashed-dotted line refers to the FAR-MIDAS SV-MIDAS model. All MIDAS models displayed use the Ffr daily series as predictor. All predictive densities are produced with recursive estimation of the models.

Figure 4. Cumulative sum of log-score differentials (CLSD)



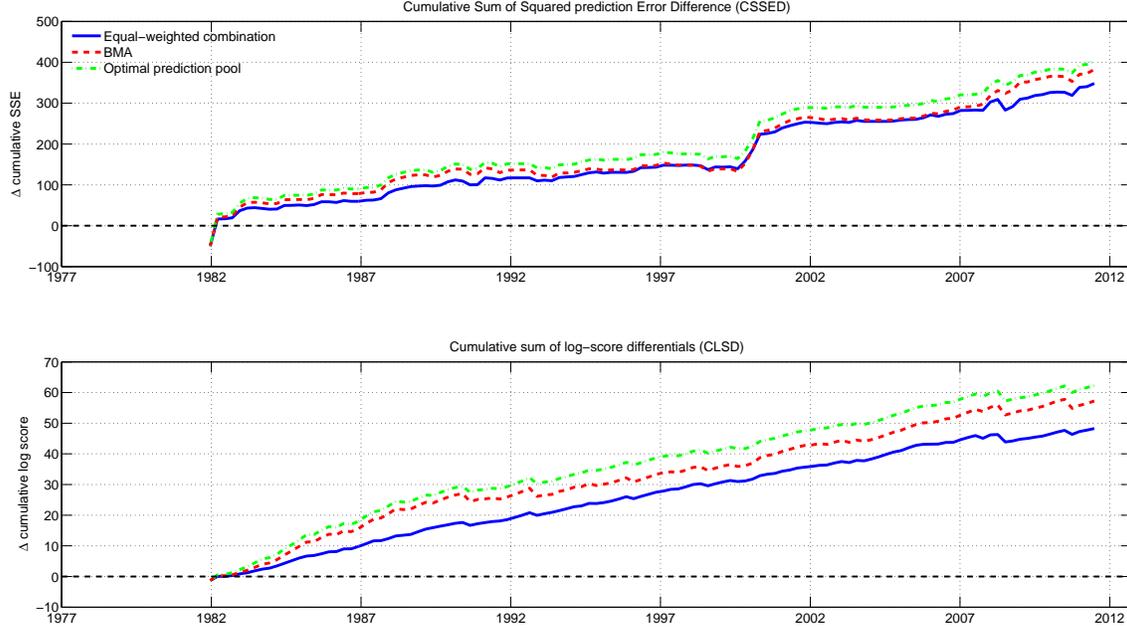
This figure shows the sum of log predictive scores from a variety of alternative forecasting models, computed relative to the sum of log predictive scores of the Random Walk (RW) model. Each quarter we estimate the parameters of the forecast models recursively and generate one-step-ahead density forecasts 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 with constant volatility (top left panel) or time-varying volatility (top right panel), while the bottom two panels display the forecasting performance of different factor-augmented autoregressive models with constant volatility (bottom left panel) or time-varying volatility (bottom right panel).

Figure 5. Relation between RMSFE and other measures of predictive performance



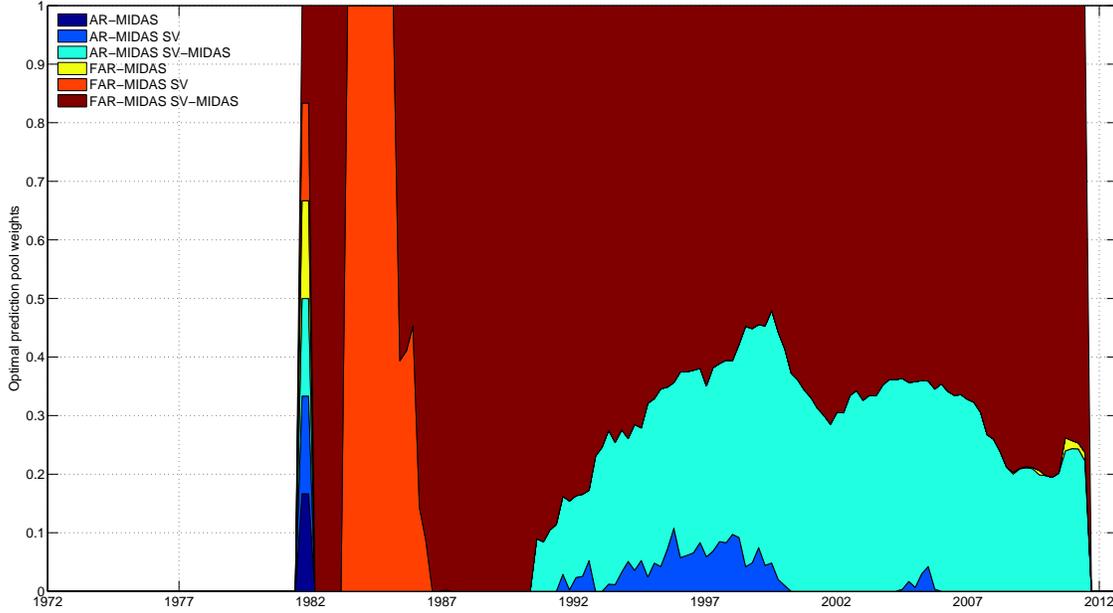
This figure presents scatter plots of out-of-sample predictive performance measures against the RMSFE (root mean squared forecast error) measure. All estimates of predictive performance are measured relative to the Random Walk (RW) model, so that the origin in each figure, shown as the intersection of 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:I and ends in 2011:IV.

Figure 6. Out-of-sample forecast performance of model combinations



The top panel of this figure shows the sum of squared forecast errors of the Random Walk model (RW) model minus the sum of squared forecast errors for three alternative forecasting models using model combinations. Three forecast combination schemes are considered, namely a simple equal-weighted combination, the optimal prediction pool of Geweke and Amisano (2011) and Bayesian Model Averaging (BMA). Each quarter we estimate the parameters of the forecasting models recursively and generate one-step-ahead forecasts of real GDP growth which are in turn used to compute out-of-sample forecast errors. This procedure is applied to the benchmark RW model, which is our benchmark, as well as to all forecasting models entering the model combinations. We then plot the cumulative sum of squared forecast errors of the RW forecasts (SSE_t^{RW}) relative to the alternative forecasts, $SSE_t^{RW} - SSE_t^{COMB}$. Values above zero indicate that a model combination generates better performance than the RW benchmark, while negative values suggest the opposite. The bottom panel of this figure shows the sum of log predictive scores for the same three alternative forecasting models using model combinations, relative to the sum of log predictive scores of the Random Walk (RW) model. Each quarter we estimate the parameters of the forecasting models recursively and generate one-step-ahead density forecasts of real GDP growth which are used to compute log-predictive scores. This procedure is applied to the benchmark RW model as well as to all forecasting models entering the model combinations. We then plot the cumulative sum of log-predictive scores for the model combinations (LS_t^{COMB}) minus the cumulative sum of log-predictive scores of the RW model, $LS_t^{COMB} - LS_t^{RW}$. Values above zero indicate that a model combination generates better performance than the RW benchmark, while negative values suggest the opposite. The blue solid line represents the equal-weighted model combination, the red dotted line tracks the optimal prediction pool, and the green dashed-dotted line depicts results for the BMA combination.

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

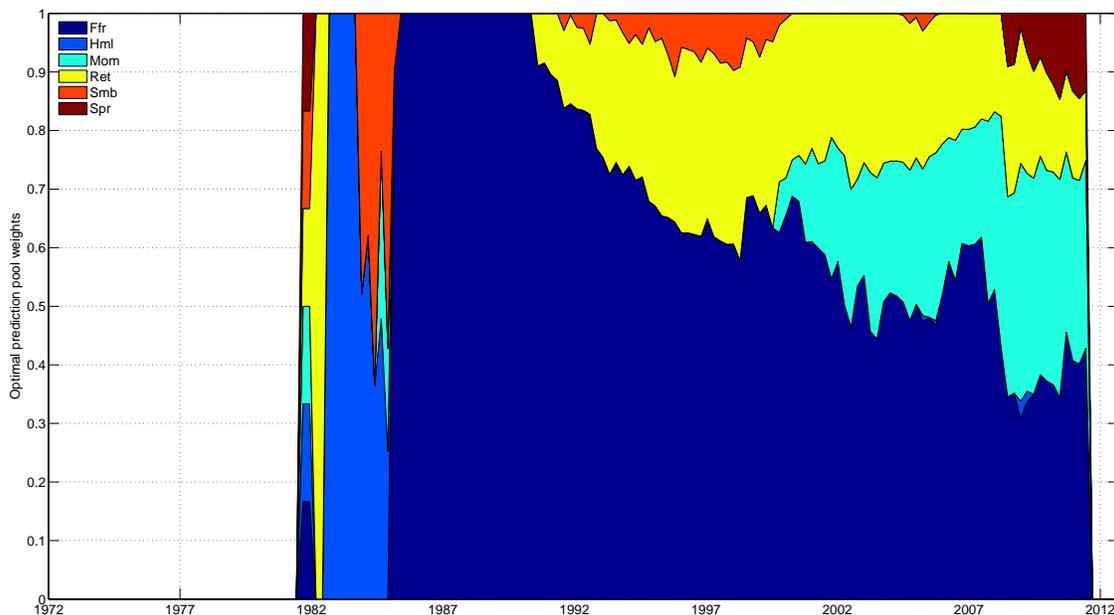


This figure plots the optimal weights on different models in the predictive pool, computed in real time by solving the minimization problem

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

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; yellow 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 maroon 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 the HML portfolio return (Hml), and light blue bars show the weights assigned to the the MOM portfolio return (Mom); yellow bars show the weights associated with the value-weighted stock returns, orange bars show the weights assigned to the SML portfolio return (Smb), and maroon bars show the weights assigned to the interest rate spread between the 10-year government bond rate and the federal fund rate (Spr).