VU Research Portal

Forecasting economic time series using score-driven dynamic models with mixed-data sampling

Gorgi, Paolo; Koopman, Siem Jan; Li, Mengheng

published in

International Journal of Forecasting
2019

DOI (link to publisher)

[10.1016/j.ijforecast.2018.11.005](https://doi.org/10.1016/j.ijforecast.2018.11.005)

document version

Publisher's PDF, also known as Version of record

document license

Article 25fa Dutch Copyright Act

[Link to publication in VU Research Portal](#)

citation for published version (APA)

Gorgi, P., Koopman, S. J., & Li, M. (2019). Forecasting economic time series using score-driven dynamic models with mixed-data sampling. *International Journal of Forecasting*, 35(4), 1735-1747.
<https://doi.org/10.1016/j.ijforecast.2018.11.005>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal ?

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

E-mail address:

vuresearchportal.ub@vu.nl

Forecasting economic time series using score-driven dynamic models with mixed-data sampling

Paolo Gorgi^{a,b}, Siem Jan Koopman^{a,b,c,*}, Mengheng Li^d

^a Vrije Universiteit Amsterdam School of Business and Economics, The Netherlands

^b Tinbergen Institute, The Netherlands

^c CREATES, Aarhus University, Denmark

^d University of Technology Sydney, Australia

ARTICLE INFO

Keywords:

Generalized autoregressive score models
Mixed frequency time series
Time-varying parameters
Gross domestic product
Inflation

ABSTRACT

We introduce a mixed-frequency score-driven dynamic model for multiple time series where the score contributions from high-frequency variables are transformed by means of a mixed-data sampling weighting scheme. The resulting dynamic model delivers a flexible and easy-to-implement framework for the forecasting of low-frequency time series variables through the use of timely information from high-frequency variables. We verify the in-sample and out-of-sample performances of the model in an empirical study on the forecasting of U.S. headline inflation and GDP growth. In particular, we forecast monthly headline inflation using daily oil prices and quarterly GDP growth using a measure of financial risk. The forecasting results and other findings are promising. Our proposed score-driven dynamic model with mixed-data sampling weighting outperforms competing models in terms of both point and density forecasts.

© 2018 International Institute of Forecasters. Published by Elsevier B.V. All rights reserved.

1. Introduction

In studies concerning the forecasting of economic time series with several variables, we often need to overcome complexities related to the different sampling frequencies at which we observe the variables over time. The challenges of mixed data frequency are reviewed in the context of econometric analysis by Ghysels and Marcellino (2016) and discussed in the context of forecasting by Andreou, Ghysels, and Kourtellis (2011) and Armesto, Engemann, and Owyang (2010). In cases of economic forecasting in particular, where both economic and financial variables are relevant, the distinction between low- and high-frequency data sampling can be substantial. Financial variables, such as stock prices, commodity prices and exchange rates, are typically available at the daily frequency, and increasingly

at the intraday level (ultra-high frequency), since it is relatively straightforward to record financial transactions electronically. On the other hand, it is both more complicated and more costly to collect data on economic variables, such as inflation and gross domestic product (GDP) growth. Hence, economic variables are available typically at the quarterly or monthly level. When the interest is in the forecasting of economic variables, the high-frequency financial variables may be relevant as predictors, and may be capable of improving the forecast accuracy. Recent research studies, such as those by Adrian, Boyarchenko, and Giannone (2018) and Carriero, Clark, and Marcellino (2018), among others, have highlighted the key role of variables measuring financial conditions in the context of forecasting macroeconomic variables.

A widely used method for incorporating high-frequency data into models for forecasting low-frequency variables is the mixed data sampling (MIDAS) method of Ghysels, Santa-Clara, and Valkanov (2004). MIDAS is a regression-based method that transforms high-frequency variables

* Correspondence to: Department of Econometrics, Vrije Universiteit Amsterdam School of Business and Economics, De Boelelaan 1105, 1081 HV Amsterdam, The Netherlands.

E-mail address: s.j.koopman@vu.nl (S.J. Koopman).

into low-frequency indicators via a parsimonious weighting scheme with possibly different weights for data sampled at a high frequency (within the low-frequency period). The weighting scheme can reflect the notion that more recent observations should be more informative for predicting future values of the low-frequency variable. The MIDAS approach (or touch) can be used easily within a (dynamic) regression model, but it can also be adopted within other models, including vector autoregressive and dynamic factor models. For instance, [Marcellino and Schumacher \(2010\)](#) considered a two-step approach that combined principal component analysis and MIDAS regressions.

The current study adopts a dynamic model with score-driven time-varying location and scale parameters. [Creal, Koopman, and Lucas \(2013\)](#) and [Harvey \(2013\)](#) have developed a general framework for specifying time-varying parameter models in an observation-driven setting. The resulting class of models is referred to as generalized autoregressive score (GAS) models. The defining feature of GAS models is that the time-varying parameters are driven by the score of the predictive log-likelihood function. The use of the score as an updating mechanism is intuitive: it can be viewed as a Newton–Raphson update that delivers a better fit, in terms of likelihood maximization, for the next period, conditional on past and current information. The score-driven updates have an optimality property. [Blasques, Koopman, and Lucas \(2015\)](#) show that the score update is optimal in that it minimizes the Kullback–Leibler divergence with respect to an unknown true distribution. Score-driven models provide an appealing forecasting method and have been employed successfully in a range of empirical applications for forecasting economic and financial variables, see for instance [Delle Monache and Petrella \(2017\)](#) on forecasting inflation, and [Blasques, Ji and Lucas \(2016\)](#) and [Lucas and Zhang \(2016\)](#) on forecasting exchange rates. In a more general context, the forecasting performances of GAS models are investigated in detail by [Koopman, Lucas, and Scharth \(2016\)](#). Finally, GAS models are appealing because they are flexible in terms of their specification, while retaining a simple practical implementation. The estimation of unknown parameters in GAS models can be based on standard likelihood inference that does not require computationally-intensive or simulation-based methods.

Our main contribution is the introduction of a flexible and easy-to-implement forecasting method for mixed frequency variables that is based on a score-driven dynamic model. In particular, we consider a factor structure in which the score innovations from the high-frequency variables are transformed into a low-frequency score function via a MIDAS weighting scheme. We name the resulting approach MIDAS-GAS. The MIDAS-GAS model retains all of the appealing features of standard GAS models and elevates the MIDAS approach to a more general device for handling mixed frequencies. For example, we illustrate how the MIDAS-GAS framework can be used to specify mixed-frequency models with conditional heteroscedastic errors and parameter updates that are robust to outliers. Furthermore, we adopt the weighted likelihood approach of [Blasques, Koopman, Mallee and Zhang \(2016\)](#) for the likelihood-based estimation of parameters

in the MIDAS-GAS model. We discuss how the proposed weighted likelihood method can be reduced to the standard maximum likelihood method when considering only the likelihood contributions of the variable of interest. These developments deliver a computationally fast and easy-to-implement methodology for parameter estimation, analysis and forecasting. We illustrate the use of the MIDAS-GAS framework for producing forecasts of monthly U.S. headline inflation and GDP growth. In particular, we consider daily crude oil inflation as a predictor for forecasting monthly inflation, and take a daily measure of financial condition as a predictor for quarterly GDP. Furthermore, we also show how the MIDAS-GAS approach can be used for nowcasting, and apply it to the nowcasting of GDP growth. We present a detailed account of our forecasting and nowcasting study, including providing comparisons with many competing models, such as MIDAS regression models, autoregressive models and standard GAS models. The results show that the performance of the MIDAS-GAS model is promising in terms of both point and density predictions.

An alternative approach to MIDAS-based methods for dealing with mixed frequency data is provided by state space time series analyses that rely on the Kalman filter. This approach involves aligning the data at the highest data sampling frequency and introducing missing observations for the low frequency variables. The Kalman filter is then used to handle these artificial missing observations, see [Blasques, Koopman et al. \(2016\)](#), [Mariano and Mura-sawa \(2003\)](#) and [Schumacher and Breitung \(2008\)](#) for such solutions, including interesting illustrations. One limitation of this more rigorous approach relative to our MIDAS-GAS model is that the Kalman filter requires Gaussian and homoscedastic errors. There is a considerable body of empirical evidence that shows the importance of accounting for heteroscedastic errors and fat-tailed distributions in order to obtain more accurate forecasts for economic time series; see for example [Creal, Schwaab, Koopman, and Lucas \(2014\)](#).

We proceed as follows. Section 2 introduces our general modeling framework based on the MIDAS-GAS model and the weighted likelihood method for parameter estimation. Section 3 presents our MIDAS-GAS dynamic factor model with heteroscedastic errors and robust parameter updates. Section 4 illustrates the two empirical applications of the forecasting of monthly U.S. headline inflation and quarterly GDP growth. Section 5 concludes.

2. The MIDAS-GAS model

Assume that our aim is to forecast a key economic variable that is denoted by y_t^L . The variable is observed sequentially over time at a low data sampling frequency, as indicated by L . We assume that another related variable x_t^H can be observed at a high data sampling frequency, as indicated by H , where $L < H$. This second variable is not of interest in itself, but we assume that it can be exploited in order to help us obtain more accurate forecasts for the key variable y_t^L . Hence, at each time point t of the low frequency variable y_t^L we have the predictor $x_t^H = (x_{1,t}^H, \dots, x_{n_x,t}^H)'$, where $x_t^H \in \mathbb{R}^{n_x}$ is a vector-valued variable that contains all

available high-frequency observations within time period t , and where n_x is the number of observations of the high-frequency variable that is available in time period t . For example, when we forecast monthly headline inflation, that is y_t^L , using daily crude oil inflation, that is x_t^H , we have n_x being equal to the number of working days in a month. For notational convenience and simplicity of exposition, we assume that both variables y and x are univariate. However, all of the results discussed below can be extended to the multivariate case in a straightforward way.

2.1. The MIDAS touch

Of the range of forecasting methods that use mixed-frequency data, the MIDAS regression is regarded as a simple and direct forecasting method. We consider the multiplicative MIDAS specification as explored by Bai, Ghysels, and Wright (2013) and Chen and Ghysels (2010). Denote the h -step-ahead forecast of y_t^L by $\hat{y}_{T+h|T}^L$, where T denotes the sample size. This forecast can be constructed by considering the p -lag multiplicative MIDAS regression given by

$$y_{t+h}^L = c + D_p(B, b)y_t^L + D_p(B, a) \sum_{i=1}^{n_x} \omega_i(\varphi) x_{t,i}^H + \epsilon_t, \quad (1)$$

for $t = 1, \dots, T$, where c is an intercept, $D_p(B, z) = z_0 + z_1 B + \dots + z_p B^p$ is a lag polynomial function with backshift operator B , $b = (b_0, \dots, b_p)'$ and $a = (a_0, \dots, a_p)'$ are unknown parameter vectors, and $\omega_i(\varphi)$ for $i = 1, \dots, n_x$ are weighting coefficients with a parameter vector φ , and ϵ_t is an identically independently distributed (iid) error with mean zero and variance σ^2 . Ghysels et al. (2004) advocate a parsimonious weighting function for $\omega_i(\varphi)$, where $i = 1, \dots, n_x$, based on exponential Almon lag or Beta lag parameterizations. The q th order exponential Almon lag is specified as

$$\omega_i(\varphi) = \frac{\exp(\varphi_1 i + \varphi_2 i^2 + \dots + \varphi_q i^q)}{\sum_{i=1}^{n_x} \exp(\varphi_1 i + \varphi_2 i^2 + \dots + \varphi_q i^q)},$$

for some q -dimensional parameter vector $\varphi = (\varphi_1, \dots, \varphi_q)'$. In practice, q is set equal to two, which reduces the Almon lag to a normalized exponential quadratic weighting function. The Beta lag is specified as

$$\omega_i(\varphi_1, \varphi_2) = \frac{\text{Beta}(i/n_x; \varphi_1, \varphi_2)}{\sum_{i=1}^{n_x} \text{Beta}(i/n_x; \varphi_1, \varphi_2)},$$

where $\text{Beta}(\cdot; \varphi_1, \varphi_2)$ is the probability density function of a Beta distribution with parameter vector $\varphi = (\varphi_1, \varphi_2)'$. Fig. 1 presents some illustrations of weighting functions that are based on second-order exponential Almon and Beta lags, from which we can conclude that the shapes of these weighting functions are very flexible. The parameters of the MIDAS regression include c , b , a , φ and σ^2 , which can be estimated using either nonlinear least squares or maximum likelihood (ML). We obtain the h -step-ahead forecast $\hat{y}_{T+h|T}^L$ in the usual way. Andreou et al. (2011), and the references therein, present a more detailed discussion of MIDAS regressions.

2.2. The MIDAS-GAS filter

We incorporate the mixed data sampling method using a parsimonious weighting function for the class of score-driven models that are proposed by Creal et al. (2013) and Harvey (2013), and refer to the resulting framework as MIDAS-GAS. The GAS filter provides a convenient framework for modeling time-varying parameters in an econometric model. The time-varying parameters are specified as autoregressive processes, where the innovations are the scores of the predictive log-likelihood function. This flexible approach leads to an observation-driven specification that facilitates inference in a straightforward manner because the likelihood function is available in closed form through the prediction error decomposition. We present the derivation of the MIDAS-GAS filter below.

We consider a multivariate conditional distribution for the observable variables y_t^L and x_t^H that is of the form

$$y_t^L, x_t^H | f_t \sim p(y_t^L, x_t^H | f_t; \psi), \quad t = 1, \dots, T, \quad (2)$$

where $p(\cdot | f_t; \psi)$ is a $(n_x + 1)$ -variate conditional density, f_t is the time-varying scalar or vector, and ψ is a vector of static parameters. Furthermore, we assume that (i) y_t^L is independent of x_t^H conditional on f_t , and (ii) all elements of the vector x_t^H are iid conditional on f_t . These assumptions have also been adopted for other dynamic multivariate time series models, including dynamic factor models and multivariate GAS models; see for example Creal et al. (2014) and Doz, Giannone, and Reichlin (2011). The conditional independence assumption does not imply that the variables are overall (or unconditionally) independent, since the dependencies between y_t^L and x_t^H , and between the elements of x_t^H , are determined by the common time-varying parameter f_t . Given these conditions, the joint conditional density of y_t^L and x_t^H can be factorized as

$$p(y_t^L, x_t^H | f_t; \psi) = p_y(y_t^L | f_t; \psi) \prod_{i=1}^{n_x} p_x(x_{t,i}^H | f_t; \psi), \quad (3)$$

where $p_y(\cdot | f_t; \psi)$ is the conditional density function of y_t^L and $p_x(\cdot | f_t; \psi)$ is the conditional density function of $x_{t,i}^H$, for $i = 1, \dots, n_x$ and $t = 1, \dots, T$.

The GAS model is defined through its time-varying parameter f_t , which is specified as an autoregressive process driven by the score of the predictive log-density in Eq. (3). Under standard differentiability conditions, the score function $\nabla_t = \partial \log p(y_t^L, x_t^H | f_t; \psi) / \partial f_t$ is given by

$$\nabla_t = \nabla_t^y + \sum_{i=1}^{n_x} \nabla_{t,i}^x,$$

where $\nabla_t^y = \partial \log p_y(y_t^L | f_t; \psi) / \partial f_t$ and $\nabla_{t,i}^x = \partial \log p_x(x_{t,i}^H | f_t; \psi) / \partial f_t$. In the GAS literature, the score in ∇_t is sometimes rescaled to account for the curvature of the likelihood; see Creal et al. (2013) for a more detailed discussion. By allowing some rescaling of the score, we define the score innovation as

$$s_t = s_t^y + \sum_{i=1}^{n_x} s_{t,i}^x,$$

Fig. 1. Weighting functions based on exponential Almon and Beta lags. Notes: The weights are assigned to daily observations within a month using different parameters of the exponential Almon lag (left) and Beta lag (right) functions.

where $s_t^y = S_t^y \nabla_t^y$ and $s_{i,t}^x = S_{i,t}^x \nabla_t^x$ for some given scaling factors S_t^y and $S_{i,t}^x$. For instance, these scaling factors can either be chosen to be some transformation of the Fisher information or be simply set equal to one. The score innovation s_t is easy to interpret: the score s_t^y can be viewed as the information from y_t^L that is used for updating the time-varying parameter f_t . Similarly, the score $s_{i,t}^x$ provides the information from $x_{i,t}^H$ for updating f_t . In the GAS framework, the score s_t is used as the innovation for updating the time-varying parameter f_t . It implies that the sources of information are weighted equally over the score contributions $s_{i,t}^x$, for $i = 1, \dots, n_x$. Since our objective is to forecast the variable y_t^L , it is a bit restrictive to assume that the predictive content of the elements $s_{i,t}^x$, for $i = 1, \dots, n_x$, is the same. For instance, in some situations it may be reasonable to assume that the score innovations at the end of time period t may be more informative for predicting y_{t+1}^L because they are closer to time period $t+1$. We therefore introduce a MIDAS weighting scheme for the score innovations $s_{i,t}^x$; that is, $\sum_{i=1}^{n_x} \omega_i(\varphi) s_{i,t}^x$. This allows the more recent score innovations to receive more weight. The resulting MIDAS-GAS filter takes the simple form

$$f_{t+1} = \delta + \beta f_t + \alpha_y s_t^y + \alpha_x \sum_{i=1}^{n_x} \omega_i(\varphi) s_{i,t}^x, \quad (4)$$

where δ , β , α_y , α_x and φ are static parameters that need to be estimated. The parameter δ is the constant and determines the mean of f_t , which is given by $\delta/(1 - \beta)$ when $|\beta| < 1$. The parameter β is the autoregressive coefficient and determines the persistence of process f_t . The parameters α_y and α_x determine the relative importance of y_t^L and x_t^H , respectively, for updating f_t and predicting future values of y_t^L . Given that the autoregressive dynamics of f_t in Eq. (4) are of order one, and only the contemporaneous score innovations s_t^y and $s_{i,t}^x$ are considered, we refer to this specification as the MIDAS-GAS(1,1). However, it is straightforward to extend Eq. (4) to a higher-order specification by adding lags of f_t and of the score innovations.

Furthermore, we currently treat f_t as a scalar, but, in general, it can be a vector of time-varying parameters.

The MIDAS-GAS model specified by Eqs. (2)–(4) is very general: it allows a wide class of observation densities to be considered. For instance, the MIDAS-GAS filter can be employed when the observed variables y_t^L and x_t^H are ordinal or categorical, or originate from densities such as the ordered logit. We refer the reader to the paper by Creal et al. (2014) for a review of possible applications in this context. We present some MIDAS-GAS specifications with dynamic means and variances in Section 3. These specifications are well suited for obtaining point and density forecasts of economic variables. In addition, Section 3 shows that the MIDAS-GAS approach nests the multiplicative MIDAS model in Eq. (1) when a Gaussian distribution for the error term is considered.

2.3. Weighted likelihood estimation

One of the appealing features of GAS models is that the likelihood function is available in closed form through the prediction error decomposition, meaning that ML estimation is easy to implement and computationally fast. This feature also applies to the case of our MIDAS-GAS model in Eqs. (2)–(4), for which we need to estimate the parameter vector $\theta = (\psi', \delta, \beta, \alpha_y, \alpha_x, \varphi')$ that contains all of the static parameters in the model. As an alternative to ML, we consider the weighted maximum likelihood (WML) approach, as proposed by Blasques, Koopman et al. (2016) for a Gaussian dynamic factor model with mixed frequency data. They show that forecasting can be improved by using parameters that are estimated by maximizing a likelihood function that gives different weights to different dependent variables, in our case y_t^L and x_t^H . The novelty is in introducing variable-specific rather than observation-specific weights in the likelihood function. This method is particularly appealing in our setting because we are only interested in forecasting the low-frequency variable y_t^L , with the high-frequency variable x_t^H being used only as a predictor.

By following Blasques, Koopman et al. (2016), the weighted likelihood of the MIDAS-GAS model can be expressed as

$$L_T^W(\theta) = \sum_{t=1}^T \log p_y(y_t^L | f_t; \psi) + W \sum_{t=1}^T \sum_{i=1}^{n_x} \log p_x(x_{i,t}^H | f_t; \psi),$$

for a predetermined weight $W \in [0, 1]$. When the weight W is set to one, the weighted likelihood function reduces to the usual likelihood function; when the weight W is set to zero, only the likelihood of the variable y_t^L is considered. We highlight the fact that setting $W = 0$ may lead to a lack of identifiability of some parameters; this needs to be accounted for on a case-by-case basis, when we have a specific model at hand. Some parameter restrictions for ensuring identifiability are discussed in Section 4. The maximization of the weighted likelihood function is carried out via standard numerical optimization routines. In general, the weight W cannot be estimated along with the other parameters. A cross-validation approach can be considered for selecting an appropriate value for the weight W .

In the empirical applications of Section 4, we set the weight to zero, and therefore consider only the likelihood contribution of the univariate time series y_t^L . Although Blasques, Koopman et al. (2016) show that a different choice of the weight may provide better out-of-sample forecast results, we consider $W = 0$ in order to obtain a simple form of the likelihood function. Furthermore, it facilitates a more realistic forecasting comparison with other models in the empirical study. In the case where we set W to zero, the MIDAS-GAS model can be regarded as a univariate model of y_t^L of the form

$$y_t^L \sim p_y(y_t^L | f_t; \psi),$$

$$f_{t+1} = \delta + \beta f_t + \alpha_y s_t^y + \alpha_x \sum_{i=1}^{n_x} \omega_i(\varphi) s_{i,t}^x, \quad t = 1, \dots, T.$$

The high-frequency variable x_t^H still enters the MIDAS-GAS updating equation for the time-varying parameter f_t through the score innovations $s_{i,t}^x$. Hence, this model can be viewed as a univariate GAS model with an exogenous predictor x_t^H . Therefore, the standard asymptotic theory for the ML estimation of GAS models of Blasques, Koopman, and Lucas (2014) applies in this case.

3. A MIDAS-GAS factor model

This section develops a selection of stylized dynamic factor model specifications for mixed frequency data based on the general MIDAS-GAS modelling framework in Eqs. (2)–(4). We formulate a parsimonious factor forecasting model that can be regarded as a multivariate time-varying parameter model for mixed frequency data. The time-varying parameters are for both the location (or mean) and the scale (or variance) of the observations. A basic formulation of our proposed factor model is given by

$$\begin{bmatrix} y_t^L \\ x_t^H \end{bmatrix} = \begin{bmatrix} 1 \\ \lambda_\mu \mathbf{1}_{n_x} \end{bmatrix} \mu_t + \sigma_t \begin{bmatrix} \epsilon_t^y \\ \epsilon_t^x \end{bmatrix}, \quad t = 1, \dots, T, \quad (5)$$

where the scalar λ_μ is the factor loading for the high-frequency variable, $\mathbf{1}_n$ is a $n \times 1$ vector of ones, the scalar

μ_t is the time-varying location variable, and the error ϵ_t^y and the errors in the vector $\epsilon_t^x = (\epsilon_{1,t}^x, \dots, \epsilon_{n_x,t}^x)'$ are independent of each other, are serially independent (over time index t), and come from some parametric distributions. In particular, the error ϵ_t^y has a mean zero and unit variance, and the error $\epsilon_{i,t}^x$ has a mean zero and variance λ_{σ_i} , for $i = 1, \dots, n_x$ and $t = 1, \dots, T$. The specification in Eq. (5) entails a factor structure for μ_t . The inclusion of the dynamic scale variable σ_t^2 enables the model to capture conditional heteroscedasticity in the data. Conditional heteroscedasticity is analyzed widely in financial time series using the generalized autoregressive conditional heteroscedasticity (GARCH) model of Bollerslev (1986) and Engle (1982). Our treatment for σ_t^2 below is similar to GARCH, and is even equivalent when normal errors are assumed; see the discussion by Creal et al. (2013). Allowing for conditional heteroscedasticity can be particularly useful for improving density forecasts.

We specify the dynamic location μ_t and the scale σ_t^2 according to the MIDAS-GAS updating equation in Eq. (4), that is,

$$\mu_{t+1} = \delta_\mu + \beta_\mu \mu_t + \alpha_\mu^y s_t^y + \alpha_\mu^x \sum_{i=1}^{n_x} \omega_i(\varphi) s_{i,t}^x,$$

$$\sigma_{t+1}^2 = \delta_\sigma + \beta_\sigma \sigma_t^2 + \alpha_\sigma^y v_t^y + \alpha_\sigma^x \sum_{i=1}^{n_x} \omega_i(\varphi) v_{i,t}^x, \quad (6)$$

where s_t^y and $s_{i,t}^x$ are the score contributions for the time-varying mean μ_t from the variables y_t^L and $x_{i,t}^H$, respectively, and v_t^y and $v_{i,t}^x$ are the score contributions for the time-varying variance σ_t^2 from y_t^L and $x_{i,t}^H$, respectively. The functional forms of the score innovations s_t^y , $s_{i,t}^x$, v_t^y and $v_{i,t}^x$ depend on the choice of the distributions of the error terms ϵ_t^y and ϵ_t^x . In what follows, we consider two different specifications for the error terms: (i) the normal distribution and (ii) the Student- t distribution. The use of a Student- t distribution leads to a more robust updating of the time-varying mean and variance.

In the first specification, we consider the errors to be distributed normally: $\epsilon_t^y \sim N(0, 1)$ and $\epsilon_{i,t}^x \sim N(0, \lambda_{\sigma_i})$. For this specification, and by scaling the score using the Fisher information, we obtain the score innovations, which are, up to a constant, given by

$$s_t^y = y_t^L - \mu_t, \quad s_{i,t}^x = x_{i,t}^H - \lambda_{\mu_i} \mu_t,$$

$$v_t^y = (y_t^L - \mu_t)^2 - \sigma_t^2, \quad v_{i,t}^x = (x_{i,t}^H - \lambda_{\mu_i} \mu_t)^2 - \lambda_{\sigma_i} \sigma_t^2. \quad (7)$$

The specification in Eqs. (6) and (7) nests the multiplicative MIDAS model in Eq. (1). In particular, by setting $\beta_\mu = \alpha_\mu^y + \alpha_\mu^x \lambda_\mu$ and $\alpha_\sigma^y = \alpha_\sigma^x = \beta_\sigma = 0$, the model for y_t becomes

$$y_t = \mu_t + \sqrt{\delta_\sigma} \epsilon_t^y,$$

$$\mu_t = \delta_\mu + \alpha_\mu^y y_{t-1} + \alpha_\mu^x \sum_{i=1}^{n_x} \omega_i(\varphi) x_{i,t-1}^H, \quad t = 1, \dots, T.$$

This corresponds to the multiplicative MIDAS models of order one that are developed by Bai et al. (2013) and Chen and Ghysels (2010). The MIDAS-GAS can be extended to higher orders by including more lags of μ_t and of the

score innovations. In a similar way, the factor MIDAS-GAS can nest multiplicative MIDAS of any order, as long as the normal distribution is considered for the errors.

In the second specification, we assume that ϵ_t^y come from a Student- t distribution with zero mean and unit variance: $\epsilon_t^y \sim t_\nu(0, 1)$. Hence, the conditional density function of y_t^L is given by

$$p_y(y_t^L | \mu_t, \sigma_t^2; \nu) = \frac{\Gamma(\frac{\nu+1}{2})}{\Gamma(\frac{\nu}{2})\sqrt{\pi(\nu-2)\sigma_t^2}} \left(1 + \frac{(y_t - \mu_t)^2}{(\nu-2)\sigma_t^2}\right)^{-\frac{\nu+1}{2}},$$

where $\nu > 2$ represents the degrees of freedom of the Student- t distribution for ϵ_t^y . The use of a Student- t distribution can be particularly important for density forecasts when the time series of interest y_t^H exhibits fat tails. Furthermore, as we will illustrate below, the Student- t model specification delivers a robust update for the time-varying parameter in the MIDAS-GAS model that may lead to more accurate point forecasts. The Student- t distribution for the error terms in ϵ_t^x can also be considered. However, we keep the exposition simple by considering a normal distribution for the terms in ϵ_t^x . We stress that the choice of the distribution of ϵ_t^x is less relevant than the choice of that of ϵ_t^y , because we are not interested in forecasting x_t^H . Furthermore, we consider mainly the estimation of the parameters using the WML method, with $W = 0$. For this second specification, the score innovations (up to a scaling constant) are given by

$$\begin{aligned} s_t^y &= \frac{(\nu+1)(y_t^L - \mu_t)}{(\nu-2) + (y_t^L - \mu_t)^2 \sigma_t^{-2}}, \\ s_{i,t}^x &= x_{i,t}^H - \lambda_\mu \mu_t, \\ v_t^y &= \frac{(\nu+1)(y_t^L - \mu_t)^2}{(\nu-2) + (y_t^L - \mu_t)^2 \sigma_t^{-2}} - \sigma_t^2, \\ v_{i,t}^x &= (x_{i,t}^H - \lambda_\mu \mu_t)^2 - \lambda_\sigma \sigma_t^2. \end{aligned} \quad (8)$$

These score innovations are robust against outliers. This is a convenient feature of the GAS framework when it relies on fat-tailed distributions; see [Harvey and Luati \(2014\)](#) and [Harvey \(2013\)](#) for a discussion of robust score updates with the Student- t distribution.

We can identify three differences between the MIDAS-GAS model and the widely-used Gaussian linear dynamic factor model. First, we do not restrict ϵ_t^y and ϵ_t^x to come from Gaussian distributions. Second, the dynamic factor process is non-linear due to the score updating mechanism, which minimizes the KL divergence locally between the true measure of the data and that implied by the factor model; see the discussion by [Blasques et al. \(2015\)](#). Third, the MIDAS touch incorporated in the dynamic factor process in Eq. (6) balances the predictive information from y_t^L and x_t^H in a different way from the MIDAS factor model of [Marcellino and Schumacher \(2010\)](#). The MIDAS factor model extracts the factors directly from x_t^H , and the resulting high-frequency factors are then included as regressors in a standard MIDAS regression. It is also different from the MIDAS dynamic factor model of [Frail and Monteforte \(2011\)](#), which is basically a bivariate Gaussian model with observation vector $(y_t^L, \sum_{i=1}^{n_x} \omega_i(\varphi) x_{i,t}^H)'$ and a single factor.

Both models rely on both Gaussian and linear assumptions, such that methods of principal component analysis (PCA) and the Kalman filter can be used for extracting factors and producing forecasts. The GAS filter in our model has the appealing feature that it is not limited to Gaussian and linear assumptions. As a result, density forecasts can be improved by incorporating stochastic volatility and distributions with fat tails. This has been documented extensively in studies where the forecasting performances of dynamic models with stochastic volatility are considered; see for example [Chib, Nardari, and Shephard \(2002\)](#), [Kim, Shephard, and Chib \(1998\)](#) and [Tse and Tsui \(2002\)](#). However, the analysis for such models is demanding computationally, and Bayesian methods are often used. On the other hand, the estimation of MIDAS-GAS models is straightforward, relatively fast and based on standard maximum likelihood methods.

4. Empirical applications

Next, we employ the MIDAS-GAS factor model for forecasting U.S. headline inflation and gross domestic product (GDP) growth. In our first application, we forecast the key variable of monthly headline inflation using daily crude oil inflation as the high-frequency variable. In our second application, we forecast the key variable of quarterly GDP growth using a daily measure of financial conditions, as obtained from the S&P 500 stock index. This allows us to assess the performance of the MIDAS-GAS model for different frequencies and different variables.

4.1. Forecasting monthly inflation with daily oil prices

4.1.1. The dataset and in-sample results

There is some evidence in the economics literature that oil prices have relevant predictive content for U.S. inflation; see for example [Clark and Terry \(2010\)](#) and [Stock and Watson \(2003\)](#). We consider time series of monthly U.S. headline inflation and crude oil inflation from January 1986 to August 2018. [Fig. 2](#) presents the time series plots.

We consider the MIDAS-GAS factor model presented in Section 3, where the parameters are estimated by the WML method with $W = 0$. We impose the following parameter restrictions. First, we consider a random walk process for the time-varying mean (or location) by imposing $\beta_\mu = 1$ and $\delta_\mu = 0$. Furthermore, we set $\lambda_\sigma = 1$ in order to ensure identifiability; the variance λ_σ can be set to any positive value, and the resulting model will be equivalent, up to a reparameterization. Finally, we consider a second-order exponential Almon lag function for the MIDAS weights, with $\varphi = (\varphi_1, \varphi_2)'$. The factor MIDAS-GAS model with the restrictions above can be written as a univariate model for y_t^L of the form

$$\begin{aligned} y_t^L &= \mu_t + \sigma_t \epsilon_t^y, \quad \mu_{t+1} = \mu_t + \alpha_\mu^y s_t^y + \alpha_\mu^x \sum_{i=1}^{n_x} \omega_i(\varphi) s_{i,t}^x, \\ \sigma_{t+1}^2 &= \delta_\sigma + \beta_\sigma \sigma_t^2 + \alpha_\sigma^y v_t^y + \alpha_\sigma^x \sum_{i=1}^{n_x} \omega_i(\varphi) v_{i,t}^x, \end{aligned} \quad (9)$$

Fig. 2. The headline consumer price index (CPI) inflation and oil price inflation. Notes: Left: The monthly headline inflation, computed as the first difference of the logarithm of monthly consumer price indices (CPI). Right: The daily oil price inflation, computed as the first difference of the logarithm of daily WTI crude oil prices.

Table 1

Full-sample parameter estimates of MIDAS-GAS factor models using monthly inflation with daily oil prices.

	ν	λ_{μ}^x	α_{μ}^y	α_{μ}^x	α_{σ}^y	α_{σ}^x	δ_{σ}	β_{σ}	llik	AIC
t-MIDAS-GASg	4.86	0.13	0.14	0.11	0.34	0.01	3.48	0.83	−892.81	4.61
t-MIDAS-GAS	4.24	0.16	0.30	0.18	–	–	7.62	–	−926.23	4.76
MIDAS-GASg	–	0.14	0.18	0.10	0.24	0.00	4.63	0.91	−906.13	4.67
MIDAS-GAS	–	0.15	0.28	0.15	–	–	8.57	–	−977.24	5.02

The last two columns report the log-likelihood and the average AIC value, respectively.

where the score innovations have the functional form as in either Eq. (7) or Eq. (8), depending on the choice of the distribution for the errors.

We analyze four different specifications of the MIDAS-GAS model in Eq. (9). The t-MIDAS-GASg has a Student- t error for y_t^L , and the score innovations are given by Eq. (8); the t-MIDAS-GAS is the same as the t-MIDAS-GASg but without conditional heteroscedasticity, that is, $\beta_{\sigma} = \alpha_{\sigma}^x = \alpha_{\sigma}^y = 0$; the MIDAS-GASg has a normal error for y_t^L and the score innovations are given by Eq. (7); and the MIDAS-GAS is the same as the MIDAS-GASg but with a constant variance, that is, $\beta_{\sigma} = \alpha_{\sigma}^x = \alpha_{\sigma}^y = 0$. Table 1 presents the estimation results. We find that the Student- t distribution improves the in-sample fit, as the models with Student- t errors perform best in terms of the Akaike information criterion (AIC). In addition, the degrees of freedom parameters ν are estimated as relatively small values, around four, which indicates the presence of fat tails. The better fit of the Student- t model is not surprising, since the headline inflation series exhibits several extreme observations; see Fig. 2. Furthermore, the conditional heteroscedasticity of the error term delivers a clear in-sample improvement in terms of AIC. This finding is also coherent with the volatility clustering of the inflation series that we can observe in Fig. 2.

Fig. 3 presents the graphs of the estimated MIDAS weighting functions for the four different model specifications. The graphs show that the estimated functions give more weight to more recent observations of the high frequency variable within the month. This result is coherent

with the notion that observations further away in time should be less relevant for forecasting the future values of the variable of interest. Furthermore, the results are rather consistent across the different models, as the shapes of the four weighting functions are very similar.

4.1.2. Out-of-sample exercise

We consider two forecasting tasks: point forecasting and density forecasting. Point forecasting is a core task of central banks and is carried out on a daily basis. For example, inflation forecasts facilitate the use of forward-looking monetary policy, which supports the computation of the ex ante real interest rate for determining the aggregate demand or IS curve for an economy. Density forecasting is important because it provides a risk metric for the accuracy of the point forecast. We evaluate the performances of point forecasts using the forecast mean squared error (FMSE), and those of density forecasts using the log score criterion. The log score criterion is a standard method for evaluating density forecasts that is based on Kullback–Leibler divergence; see for instance Geweke and Amisano (2011).

The full data sample consists of 392 months, and runs from April 1986 to August 2018. We split this sample into two subsamples: the first sample of 187 months is used as an in-sample training period, and the second sample of 200 months is used for out-of-sample evaluation. We consider a rolling window forecasting exercise, meaning that the length of the in-sample estimation period is equal to 187 months for all forecasts. We further consider multi-step

Fig. 3. Estimated MIDAS-GAS weighting functions for the four model specifications using monthly inflation with daily oil prices. Notes: The horizontal axis indicates the day in the month (there are about 21 working days in a month). More weight on the last days of the month indicates that more recent observations are more relevant.

forecasts, from one to six months ahead. We highlight the fact that the forecasts have the same horizon for the high- and low-frequency variables; for instance, the 1-month-ahead forecast of January CPI is obtained using (low- and high-frequency) data up to the end of December only. In addition to our four MIDAS-GAS models, we also include a set of competing models for the purpose of forecast comparison, including MIDAS regression models, autoregressive models, standard GAS models and the MIDAS factor model of [Frale and Monteforte \(2011\)](#). For these models, we consider Student-*t* error distributions and conditional heteroscedasticity. [Table 2](#) displays the specifications of the competing models that are included in the comparison. The MIDAS regression models and the autoregressive models are estimated as direct forecasting methods for each forecasting horizon. The statistical significance of the difference in performance of each model, compared to that of a benchmark model (t-MIDAS-GASg), is tested using the Diebold–Mariano (DM) test of [Diebold and Mariano \(1995\)](#). In general, the DM test is problematic when testing nested models. However, [Giacomini and White \(2006\)](#) showed that the DM test remains valid when a rolling window approach is considered. We therefore adopt a rolling window forecasting exercise.

[Table 3](#) reports the results of the forecasting study. It shows that MIDAS-GAS models tend to have the best performances in terms of point forecasts, except for one-step-ahead forecasts, where the variants and extensions of multiplicative MIDAS regressions perform equally well. Furthermore, we notice that the inclusion of conditional heteroscedasticity and Student-*t* errors plays a major role in the success of our framework. The t-MIDAS-GASg model tends to have the best performance among the set of MIDAS-GAS models. We obtain a similar result for density forecasts. Here, the t-MIDAS-GASg has the best performance for several forecasting horizons. Overall, we can conclude that MIDAS-GAS models can deliver accurate forecasts relative to a wide pool of competing models.

Table 2

Specification of the competing models used in the out-of-sample exercise.

	Model description
t-MIDASg(<i>p</i>)	The <i>p</i> -lag MIDAS regression in Eq. (1) with Student- <i>t</i> and conditionally heteroscedastic errors.
t-MIDAS(<i>p</i>)	The <i>p</i> -lag MIDAS regression in Eq. (1) with Student- <i>t</i> errors.
MIDASg(<i>p</i>)	The <i>p</i> -lag MIDAS regression in Eq. (1) with normal and conditionally heteroscedastic errors.
MIDAS(<i>p</i>)	The <i>p</i> -lag MIDAS regression in Eq. (1) with normal errors.
t-ARg(<i>p</i>)	Autoregressive model of order <i>p</i> with Student- <i>t</i> and conditionally heteroscedastic errors.
t-AR(<i>p</i>)	Autoregressive model of order <i>p</i> with Student- <i>t</i> errors.
ARg(<i>p</i>)	Autoregressive model of order <i>p</i> with normal and conditionally heteroscedastic errors.
AR(<i>p</i>)	Autoregressive model of order <i>p</i> with normal errors.
t-GASg	Standard GAS model with Student- <i>t</i> and conditionally heteroscedastic errors.
t-GAS	Standard GAS model with Student- <i>t</i> errors.
GASg	Standard GAS model with normal and conditionally heteroscedastic errors.
GAS	Standard GAS model with normal errors.
fMIDAS	The MIDAS factor model of Frale and Monteforte (2011) .

Given the importance of inflation forecasting, we regard these findings as strongly supporting our MIDAS-GAS framework.

4.2. Forecasting GDP growth with financial condition measures

4.2.1. The dataset and in-sample results

Our second illustration concerns the forecasting and nowcasting of quarterly U.S. GDP growth, which is another key economic variable. The high frequency predictor variable for our MIDAS-GAS model is a basic daily measure

Table 3

Relative MSE and log score criterion values for the different model specifications using monthly inflation with daily oil prices.

	Forecast mean squared error						Log score criterion					
	<i>h</i> = 1	<i>h</i> = 2	<i>h</i> = 3	<i>h</i> = 4	<i>h</i> = 5	<i>h</i> = 6	<i>h</i> = 1	<i>h</i> = 2	<i>h</i> = 3	<i>h</i> = 4	<i>h</i> = 5	<i>h</i> = 6
t-MIDAS-GASg	1.00	1.00	1.00	1.00	1.00	1.00	−2.45	−2.63	−2.65	−2.63	−2.62	−2.63
t-MIDAS-GAS	0.99	1.00	1.00	1.00	1.00	1.00	−2.48	−2.68	−2.68	−2.63	−2.61	−2.61
MIDAS-GASg	0.99	1.01	1.00	1.00	1.00	1.00	−2.51**	−2.80**	−2.86*	−2.87*	−2.90**	−2.96*
MIDAS-GAS	0.99	1.01	1.00	1.00	1.00	1.00	−2.68*	−2.97*	−2.97*	−2.94	−2.94	−2.96
t-MIDASg(2)	0.98	1.02	1.03*	1.03*	1.02**	1.02*	−2.46	−2.67	−2.74**	−2.73***	−2.71***	−2.71***
t-MIDAS(2)	0.97	1.05	1.03	1.05**	1.03**	1.03*	−2.57	−2.78*	−2.75	−2.75*	−2.72	−2.71
MIDASg(2)	0.95	1.01	1.02*	1.03*	1.02	1.03*	−2.52	−2.84**	−2.93**	−2.99***	−3.00**	−3.13*
MIDAS(2)	0.95	1.02*	1.03*	1.06*	1.04**	1.04*	−2.79*	−3.07*	−3.08*	−3.09*	−3.05*	−3.09*
t-MIDASg(4)	1.00	1.04	1.05*	1.06**	1.04**	1.03**	−2.49	−2.69	−2.74**	−2.72***	−2.72***	−2.72***
t-MIDAS(4)	1.01	1.03*	1.03**	1.05**	1.04***	1.03***	−2.61**	−2.80*	−2.79*	−2.79*	−2.79**	−2.80*
MIDASg(4)	1.00	1.05*	1.07*	1.07**	1.05**	1.03**	−2.55	−2.80*	−2.90**	−2.87**	−2.89**	−2.93*
MIDAS(4)	1.00	1.05*	1.06*	1.08**	1.06**	1.03*	−2.79*	−3.09*	−3.14*	−3.12*	−3.08*	−3.07*
t-ARg(2)	1.07**	1.03	1.01	1.01***	1.01***	1.01**	−2.56***	−2.67	−2.67	−2.66*	−2.66*	−2.65
t-AR(2)	1.08**	1.04	1.02	1.03**	1.02**	1.01**	−2.59**	−2.70	2.64	−2.64	−2.63	−2.62
ARg(2)	1.07**	1.01	1.00	1.01*	1.01*	1.01	−2.63**	−2.83*	−2.91*	−2.92**	−2.96**	−3.02*
AR(2)	1.06	1.03	1.02	1.03**	1.02**	1.01**	−2.82*	−3.03*	−3.00	−2.98*	−2.98*	−2.97*
t-ARg(4)	1.08**	1.04	1.02*	1.01*	1.01	1.01*	−2.58***	−2.69*	−2.69*	−2.66*	−2.66*	−2.65
t-AR(4)	1.08**	1.05	1.02	1.03*	1.02*	1.01**	−2.62***	−2.70	−2.66	−2.64	−2.62	−2.62
ARg(4)	1.09***	1.04*	1.02*	1.02	1.01*	1.01*	−2.65***	−2.85**	−2.93**	−2.92**	−2.95**	−3.02*
AR(4)	1.08*	1.06*	1.04*	1.04**	1.02**	1.01**	−2.86**	−3.13*	−3.12*	−3.08*	−3.04*	−3.03
t-GASg	1.11***	1.05*	1.05**	1.03	1.04	1.02	−2.59***	−2.74**	−2.77***	−2.72*	−2.67	−2.72**
t-GAS	1.15***	1.03	1.04	1.04	1.03	1.02	−2.68***	−2.72*	−2.70	−2.66	−2.63	−2.64
GASg	1.08***	1.02	1.01*	1.01*	1.01*	1.01	−2.63***	−2.86*	−2.95*	−2.94*	−2.99**	−3.05*
GAS	1.06*	1.01	1.01*	1.01*	1.01*	1.01*	−2.83**	−3.03*	−3.05*	−3.00	−3.00	−3.02
fMIDAS	1.13***	1.03	1.03	1.03	1.03	1.02	−2.88**	−2.97*	−2.97	−2.94	−2.83	−2.91*

Notes: The first six columns of the table report the ratio of the MSE of each model to that of the benchmark model (t-MIDAS-GASg) for several forecasting horizons (from one to six steps ahead). A value greater than one indicates that the model underperforms the benchmark model, while the opposite is true when the MSE ratio is smaller than one. The last six columns report the log score criterion for several forecasting horizons (from one to six steps ahead). Bold values indicate the best performing model for a given horizon.

***Indicates the significance level of the DM test at 0.1%.

**Indicates the significance level of the DM test at 1%.

*Indicates the significance level of the DM test at 5%.

of financial condition. In particular, we consider the log-squares of the S&P 500 log-returns as a predictor: we have $x_{it} = -\log(r_{it}^2)$, where r_{it} denotes the log-returns of the S&P 500 index, on day i in month t . The use of this type of financial risk measure for forecasting GDP growth has gained some attention in recent times. [Johannes, Lochstoer, and Mou \(2016\)](#) and [Orlik and Veldkamp \(2014\)](#) show that agents form beliefs regarding economic growth using signals of financial risk. Furthermore, the recent literature has also found that the vulnerability and downside risk of economic growth are preceded by volatility increases; see for example [D'Agostino, Gambetti, and Giannone \(2013\)](#) and [Gourio \(2012\)](#). These interactions between economic growth and financial risk are also discussed in the context of intermediary asset pricing; see for example [Brunnermeier and Sannikov \(2014\)](#).

We adopt x_{it} as our basic indicator of (the inverse of) the financial volatility for illustrating the implementation of MIDAS-GAS models because it is available on a daily basis with a sufficiently long series of historical data. More comprehensive financial condition indicators such as the weekly National Financial Conditions Index (NFCI), constructed from 105 measures of financial risks including equity volatility, credit spreads and the term spread by the Federal Reserve Bank of Chicago, are also available, but with a much shorter history. The time series that we

consider are from the first quarter of 1950 to the second quarter of 2018. [Fig. 4](#) shows the quarterly GDP growth series and the daily measure of financial risk.

[Table 4](#) reports the estimates of the MIDAS-GAS models. The results reveal that the use of the Student- t distribution leads to a small improvement in the in-sample fit. The in-sample fit improves even further when conditional heteroscedasticity is included in the models. [Fig. 5](#) presents the estimated MIDAS weighting functions. Interestingly, we observe that the estimated functions tend to give more weight to days in the first and second months of the quarter. This confirms the common wisdom that financial variables are “fast” and macroeconomic variables are “slow” in their evolutions over time. A similar discussion applies to structural vector autoregression analysis for macroeconomic and financial risk variables; see for example [Carriero et al. \(2018\)](#) and [Galí and Gambetti \(2015\)](#). [Brunnermeier and Sannikov \(2014\)](#) also suggest that an increase in volatility after a low-volatility period may precede a downward move in output growth. Hence, the optimal timing of the effect is not necessarily at the end of the intraday period (most recent high-frequency observations), as we witness in [Fig. 5](#).

4.2.2. Out-of-sample exercise

We perform a similar rolling-window forecasting exercise as in the empirical application for inflation. The

Table 4

Full-sample parameter estimates of MIDAS-GAS factor models using the quarterly GDP growth rate and the daily measure of financial risk.

	ν	λ_{μ}^x	α_{μ}^y	α_{μ}^x	α_{σ}^y	α_{σ}^x	δ_{σ}	β_{σ}	llik	AIC
t-MIDAS-GASg	9.12	1.94	0.24	0.17	0.13	0.00	9.11	0.98	−700.51	5.19
t-MIDAS-GAS	7.12	1.79	0.28	0.21	–	–	11.66	–	−719.81	5.31
MIDAS-GASg	–	1.86	0.33	0.18	0.13	0.00	12.06	0.97	−704.68	5.21
MIDAS-GAS	–	1.81	0.32	0.21	–	–	11.75	–	−726.36	5.35

Note: The last two columns report the log-likelihood and the average AIC, respectively.

Fig. 4. The GDP growth rate and the financial risk measure. Notes: Left: The quarterly GDP growth rate. Right: The daily financial risk measure, $x_{i,t} = -\log(r_{i,t}^2)$. **Fig. 5.** Estimated MIDAS-GAS weighting functions for the four model specifications using the GDP growth rate with the measure of financial risk. Note: The horizontal axis indicates the day in the quarter (there are about 63 working days in a quarter).

forecasts are obtained for one to six quarters ahead, and the out-of-sample period is from 1989 to 2018. In the case of quarterly GDP growth, the time-varying mean of the MIDAS-GAS models is specified as an autoregressive process of order two, as this may capture some more of the richer dynamics that we associate with business cycle features. Table 5 reports the results. The MIDAS-GAS models tend to have the best performances in terms of point

and density forecasts, except for one-step-ahead forecasts, where the best performer is the multiplicative MIDAS regression, extended with Student-*t* errors and with conditional heteroscedasticity. It is evident from the results in Table 5 that using the Student-*t* distribution and including conditional heteroscedasticity are both important for improving the performances for both point and density

Table 5

Relative MSE and log score criterion values for the different model specifications using quarterly GDP growth rates with the daily measure of financial risk.

	Forecast mean squared error						Log score criterion					
	<i>h</i> = 1	<i>h</i> = 2	<i>h</i> = 3	<i>h</i> = 4	<i>h</i> = 5	<i>h</i> = 6	<i>h</i> = 1	<i>h</i> = 2	<i>h</i> = 3	<i>h</i> = 4	<i>h</i> = 5	<i>h</i> = 6
t-MIDAS-GASg	1.00	1.00	1.00	1.00	1.00	1.00	−2.23	−2.22	−2.25	−2.26	−2.29	−2.30
t-MIDAS-GAS	0.99	1.00	1.01	1.01	1.02	1.02	−2.32	−2.31*	−2.33	−2.34	−2.35	−2.35
MIDAS-GASg	1.00	1.01	1.01	1.02	1.02	1.02	−2.23	−2.26	−2.33*	−2.44*	−2.49*	−2.48*
MIDAS-GAS	0.99	1.01	1.01	1.02	1.02	1.02	−2.35*	−2.36**	−2.37*	−2.38	−2.38	−2.39
t-MIDASg(2)	1.01	1.05*	1.04	1.03	1.06	1.08	−2.22	−2.26	−2.33	−2.26	−2.40	−2.43
t-MIDAS(2)	1.00	1.06*	1.06	1.04	1.05	1.07	−2.32	−2.36**	−2.38*	−2.38*	−2.39	−2.41
MIDASg(2)	1.01	1.05*	1.04	1.03	1.06	1.08	−2.26	−2.32*	−2.45	−2.56	−2.66	−2.62
MIDAS(2)	1.01	1.06*	1.05	1.04	1.04	1.07	−2.37*	−2.41***	−2.43**	−2.42**	−2.44*	−2.46**
t-MIDASg(4)	0.98	1.05	1.06	1.07	1.11	1.11	−2.20	−2.29*	−2.37*	−2.38	−2.46*	−2.48*
t-MIDAS(4)	0.99	1.05	1.07*	1.08	1.11	1.11	−2.31	−2.34*	−2.38*	−2.40*	−2.42*	−2.42*
MIDASg(4)	0.99	1.05	1.06	1.08	1.11	1.11	−2.26	−2.37	−2.44	−2.58	−2.69	−2.71
MIDAS(4)	0.99	1.05	1.06	1.08	1.11	1.10	−2.35*	−2.39**	−2.42**	−2.44**	−2.46**	−2.46**
t-ARg(2)	1.00	1.04*	1.05	1.05	1.04	1.04	−2.21	−2.23	−2.28	−2.30	−2.33	−2.35
t-AR(2)	1.00	1.04*	1.04	1.04	1.03	1.03	−2.29	−2.31	−2.34	−2.35	−2.34	−2.35
ARg(2)	1.00	1.05*	1.06	1.06	1.05	1.05	−2.26	−2.31	−2.36	−2.49	−2.53	−2.52
AR(2)	1.00	1.05*	1.05	1.05	1.03	1.03	−2.37*	−2.41***	−2.44***	−2.45**	−2.45*	−2.45*
t-ARg(4)	1.00	1.04*	1.05	1.04	1.04	1.03	−2.21	−2.23	−2.28	−2.30	−2.33	−2.34
t-AR(4)	1.00	1.05	1.05	1.05	1.03	1.03	−2.28	−2.30	−2.34	−2.34	−2.34	−2.34
ARg(4)	1.00	1.05*	1.06	1.05	1.05	1.05	−2.26	−2.31	−2.36	−2.47	−2.53	−2.52
AR(4)	0.99	1.04	1.05	1.04	1.03	1.02	−2.37*	−2.39**	−2.41**	−2.41**	−2.41*	−2.41*
t-GASg	1.02	1.06**	1.05*	1.04	1.02	1.01	−2.23	−2.24	−2.28	−2.30	−2.32	−2.33
t-GAS	1.02	1.05**	1.05*	1.03	1.02	1.01	−2.32	−2.33*	−2.35*	−2.34	−2.34	−2.34
GASg	1.02	1.05*	1.05*	1.04	1.02	1.01	−2.29	−2.30*	−2.33*	−2.37	−2.37	−2.38
GAS	1.00	1.05**	1.05*	1.04	1.02	1.01	−2.37*	−2.39**	−2.41**	−2.41*	−2.41*	−2.40
fMIDAS	1.02	1.07**	1.06*	1.06	1.04	1.04	−2.37*	−2.41***	−2.44***	−2.45**	−2.45*	−2.45*

Notes: The first six columns of the table report the ratio of the MSE of each model to that of the benchmark model (t-MIDAS-GASg) for several forecasting horizons (from 1 to 6 steps ahead). A value greater than one indicates that a model is underperforming the benchmark model, while the opposite is true when the MSE ratio is smaller than one. The last six columns report the log score criterion for several forecasting horizons (from one to six steps ahead). Bold values indicate the best performing model for a given horizon.

***Indicates the significance level of the DM test at 0.1%.

**Indicates the significance level of the DM test at 1%.

*Indicates the significance level of the DM test at 5%.

forecasting. The DM test shows that several models produce forecasts that are significantly less accurate than the t-MIDAS-GASg model, especially for the case of density forecasts. Overall, then, we conclude that MIDAS-GAS models are able to deliver accurate and competitive forecasts.

4.2.3. Nowcasting

A strong motivation for the use of MIDAS can also be found in the nowcasting literature; see for example [Marcellino and Schumacher \(2010\)](#). We therefore also illustrate our MIDAS-GAS approach for the nowcasting of U.S. GDP. In our framework, nowcasting can be implemented in a straightforward way by shifting time index *t* one period forward for the high-frequency variable. Assume that the aim is to predict y_{T+1} and that the last observations available are y_T and $x_{s,T+1}$. For instance, in the case of GDP, in quarter $T + 1$ we have the financial risk measure being observed up to some day *s* in quarter $T + 1$. In this nowcasting setting, we define the *s*-period shifted variable $x_t^s = (x_{1+s,t}, \dots, x_{n_x,t}, x_{1,t+1}, \dots, x_{s,t+1})'$. The specification of the MIDAS-GAS filter then becomes

$$f_{t+1} = \delta + \beta f_t + \alpha_y y_t^y + \alpha_x \sum_{i=1}^{n_x} \omega_i(\varphi) s_{i,t}^{x,s},$$

where $s_{i,t}^{x,s}$ denotes the score innovation from the shifted variable $x_{i,t}^s$, which is the *i*th element of the vector x_t^s . The

estimation of the parameters in the model is carried out by WML; see the discussions above.

We consider the nowcasting of quarterly GDP growth by utilizing the daily risk measure of the financial condition that becomes available for the days in the first and second months of the corresponding quarter. For example, for one-month nowcasting, the vector of daily measures that enters the model includes the daily observations in the last two months of the previous quarter plus those in the first month of the most recent quarter. [Table 6](#) reports the results of our nowcasting exercise. The results show similar performances (our benchmark model here is MIDAS-GASg) across the models. There is no clear winner for the point forecasting; however, in case of density predictions, we find that models without conditional heteroscedasticity have significantly lower nowcast accuracies than MIDAS-GASg. Overall, we can conclude that the performances of the MIDAS-GAS models and the multiplicative MIDAS regressions are comparable in this nowcasting exercise for U.S. GDP.

5. Conclusion

We have introduced the MIDAS-GAS model as a novel modelling approach to forecasting and nowcasting with

Table 6

Relative MSE and log score criterion values for the different model specifications using quarterly GDP growth with daily financial risk measure.

	One month		Two months	
	FMSE	Log score	FMSE	Log score
t-MIDAS-GASg	1.02	−2.20	1.03	−2.20
t-MIDAS-GAS	1.01	−2.30*	1.00	−2.30*
MIDAS-GASg	1.00	−2.20	1.00	−2.20
MIDAS-GAS	1.00	−2.35**	0.99	−2.34**
t-MIDASg(2)	1.01	−2.19	1.02*	−2.19
t-MIDAS(2)	1.02*	−2.30*	1.02	−2.30*
MIDASg(2)	1.02*	−2.23	1.02*	−2.23
MIDAS(2)	1.01	−2.35**	1.01	−2.35**
t-MIDASg(4)	0.99	−2.18	1.00	−2.19
t-MIDAS(4)	1.00	−2.30	1.00	−2.29
MIDASg(4)	1.00	−2.23	1.00	−2.24
MIDAS(4)	0.99	−2.34*	1.02	−2.35**
fMIDAS	1.06*	−2.42***	1.10	−2.42***

Notes: The first two columns of the table report the relative MSE relative to the benchmark model (MIDAS-GASg) and the log score criterion for one-month-ahead nowcasting. The last two columns report the results for two-month-ahead nowcasting. Bold values indicate the best performing model for a given horizon.

***Indicates the significance level of the DM test at 0.1%.

**Indicates the significance level of the DM test at 1%.

*Indicates the significance level of the DM test at 5%.

mixed-frequency data on economic variables. The MIDAS-GAS model transforms the predictive log-likelihood score contributions of the high-frequency variables through a MIDAS weighting scheme. The proposed approach has several advantages, as it retains all of the appealing features of the GAS models while at the same time accounting for mixed frequencies using MIDAS. Based on the general MIDAS-GAS framework, we have developed a novel forecasting model with dynamic factor structures for the location and scale. The method has shown promising forecasting performances in two key economic applications. In the first case, the forecasting of monthly U.S. headline inflation using daily crude oil prices is considered. In the second case, the forecasting and nowcasting of quarterly GDP growth using a daily measure of financial condition are considered. For both cases, we report favorable results for MIDAS-GAS in both forecasting and nowcasting.

References

- Adrian, T., Boyarchenko, N., & Giannone, D. (2018). Vulnerable growth. *American Economic Review*, (forthcoming).
- Andreou, E., Ghysels, E., & Kourtellis, A. (2011). Forecasting with mixed-frequency data. In M. P. Clements, & D. F. Hendry (Eds.), *The Oxford handbook of economic forecasting*. Oxford: Oxford University Press.
- Armesto, M. T., Engemann, K. M., & Owyang, M. T. (2010). Forecasting with mixed frequencies. *Federal Reserve Bank of St. Louis Review*, 92(6), 521–536.
- Bai, J., Ghysels, E., & Wright, J. H. (2013). State space models and MIDAS regressions. *Econometric Reviews*, 32(7), 779–813.
- Blasques, F., Ji, J., & Lucas, A. (2016). Semiparametric score driven volatility models. *Computational Statistics & Data Analysis*, 100, 58–69.
- Blasques, F., Koopman, S. J., & Lucas, A. (2014). *Maximum likelihood estimation for generalized autoregressive score models: Technical report*, Tinbergen Institute Discussion Paper.
- Blasques, F., Koopman, S. J., & Lucas, A. (2015). Information-theoretic optimality of observation-driven time series models for continuous responses. *Biometrika*, 102(2), 325–343.

- Blasques, F., Koopman, S., Mallee, M., & Zhang, Z. (2016). Weighted maximum likelihood for dynamic factor analysis and forecasting with mixed frequency data. *Journal of Econometrics*, 193(2), 405–417.
- Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 31(3), 307–327.
- Brunnermeier, M. K., & Sannikov, Y. (2014). A macroeconomic model with a financial sector. *American Economic Review*, 104(2), 379–421.
- Carriero, A., Clark, T. E., & Marcellino, M. (2018). Measuring uncertainty and its impact on the economy. *Review of Economics and Statistics*, 100(5), 799–815.
- Chen, X., & Ghysels, E. (2010). News good or bad and its impact on volatility predictions over multiple horizons. *The Review of Financial Studies*, 24(1), 46–81.
- Chib, S., Nardari, F., & Shephard, N. (2002). Markov chain Monte Carlo methods for stochastic volatility models. *Journal of Econometrics*, 108(2), 281–316.
- Clark, T. E., & Terry, S. J. (2010). Time variation in the inflation passthrough of energy prices. *Journal of Money, Credit and Banking*, 42(7), 1419–1433.
- Creal, D., Koopman, S. J., & Lucas, A. (2013). Generalized autoregressive score models with applications. *Journal of Applied Econometrics*, 28(5), 777–795.
- Creal, D., Schwaab, B., Koopman, S. J., & Lucas, A. (2014). Observation-driven mixed-measurement dynamic factor models with an application to credit risk. *The Review of Economics and Statistics*, 96(5), 898–915.
- D'Agostino, A., Gambetti, L., & Giannone, D. (2013). Macroeconomic forecasting and structural change. *Journal of Applied Econometrics*, 28(1), 82–101.
- Delle Monache, D., & Petrella, I. (2017). Adaptive models and heavy tails with an application to inflation forecasting. *International Journal of Forecasting*, 33(2), 482–501.
- Diebold, F. X., & Mariano, R. S. (1995). Comparing predictive accuracy. *Journal of Business & Economic Statistics*, 13, 253–265.
- Doz, C., Giannone, D., & Reichlin, L. (2011). A two-step estimator for large approximate dynamic factor models based on Kalman filtering. *Journal of Econometrics*, 164(1), 188–205.
- Engle, R. F. (1982). Autoregressive conditional heteroskedasticity with estimates of the variance of the United Kingdom inflation. *Econometrica*, 50, 987–1007.
- Frale, C., & Monteforte, L. (2011). *FaMIDAS: A mixed frequency factor model with MIDAS structure: Technical report*, Bank of Italy.
- Gali, J., & Gambetti, L. (2015). The effects of monetary policy on stock market bubbles: Some evidence. *American Economic Journal: Macroeconomics*, 7(1), 233–257.
- Geweke, J., & Amisano, G. (2011). Optimal prediction pools. *Journal of Econometrics*, 164(1), 130–141.
- Ghysels, E., & Marcellino, M. (2016). The econometric analysis of mixed frequency data sampling. *Journal of Econometrics*, 193(2), 291–293.
- Ghysels, E., Santa-Clara, P., & Valkanov, R. (2004). *The MIDAS touch: Mixed data sampling regression models: Technical report*.
- Giacomini, R., & White, H. (2006). Tests of conditional predictive ability. *Econometrica*, 74(6), 1545–1578.
- Gourio, F. (2012). Disaster risk and business cycles. *American Economic Review*, 102(6), 2734–2766.
- Harvey, A. (2013). *Dynamic models for volatility and heavy tails: with applications to financial and economic time series*. New York: Cambridge University Press.
- Harvey, A., & Luati, A. (2014). Filtering with heavy tails. *Journal of the American Statistical Association*, 109(507), 1112–1122.
- Johannes, M., Lochstoer, L. A., & Mou, Y. (2016). Learning about consumption dynamics. *The Journal of Finance*, 71(2), 551–600.
- Kim, S., Shephard, N., & Chib, S. (1998). Stochastic volatility: likelihood inference and comparison with ARCH models. *Review of Economic Studies*, 65(3), 361–393.
- Koopman, S. J., Lucas, A., & Scharth, M. (2016). Predicting time-varying parameters with parameter-driven and observation-driven models. *The Review of Economics and Statistics*, 98(1), 97–110.
- Lucas, A., & Zhang, X. (2016). Score-driven exponentially weighted moving averages and value-at-risk forecasting. *International Journal of Forecasting*, 32(2), 293–302.
- Marcellino, M., & Schumacher, C. (2010). Factor MIDAS for nowcasting and forecasting with ragged-edge data: A model comparison for German GDP. *Oxford Bulletin of Economics and Statistics*, 72(4), 518–550.

- Mariano, R. S., & Murasawa, Y. (2003). A new coincident index of business cycles based on monthly and quarterly series. *Journal of Applied Econometrics*, 18(4), 427–443.
- Orlik, A., & Veldkamp, L. (2014). *Understanding uncertainty shocks and the role of black swans: Technical report*, National Bureau of Economic Research.
- Schumacher, C., & Breitung, J. (2008). Real-time forecasting of German GDP based on a large factor model with monthly and quarterly data. *International Journal of Forecasting*, 24(3), 386–398.
- Stock, J., & Watson, M. (2003). Forecasting output and inflation: The role of asset prices. *Journal of Economic Literature*, 41(3), 788–829.
- Tse, Y. K., & Tsui, A. K. C. (2002). A multivariate generalized autoregressive conditional heteroscedasticity model with time-varying correlations. *Journal of Business & Economic Statistics*, 20(3), 351–362.

Paolo Gorgi is Assistant Professor of Econometrics at Vrije Universiteit Amsterdam and junior research fellow at Tinbergen Institute since 2017. He obtained his Ph.D. in 2017 from University of Padua. His research

interests include various topics in time series econometrics, observation-driven models, asymptotic theory and forecasting.

Siem Jan Koopman is Professor of Econometrics at Vrije Universiteit Amsterdam and research fellow at Tinbergen Institute since 1999. He further is a regular Visiting Professor at CREATES, Aarhus University. His Ph.D. is from the London School of Economics (LSE) and dates back to 1992. He had earlier positions at the LSE and CentER, Tilburg University. His research interests include various topics in time series econometrics, financial econometrics and economic forecasting. In particular, his research is focussed on state space models and generalized autoregressive score models.

Mengheng Li is Lecturer at the Economics Group of the University of Technology Sydney Business School. He obtained his Ph.D. in 2018 from the Vrije Universiteit Amsterdam School of Business and Economics, and the Tinbergen Institute, in the Netherlands. His research interests include time series econometrics, Bayesian methods, financial econometrics and economic forecasting.