TI 2018-026/III
Tinbergen Institute Discussion Paper

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

Paolo Gorgi[1]
Siem Jan (S.J.) Koopman[2]
Mengheng Li[3]

1: VU Amsterdam
2: VU Amsterdam; Tinbergen Institute, The Netherlands
3: VU Amsterdam

Tinbergen Institute is the graduate school and research institute in economics of Erasmus University Rotterdam, the University of Amsterdam and VU University Amsterdam.

Contact: discussionpapers@tinbergen.nl

More TI discussion papers can be downloaded at the Tinbergen Site

Tinbergen Institute has two locations:

Tinbergen Institute Amsterdam
Gustav Mahlerplein 117
1082 MS Amsterdam
The Netherlands
Tel.: +31(0)20 598 4580

Tinbergen Institute Rotterdam
Burg. Oudlaan 50
3062 PA Rotterdam
The Netherlands
Tel.: +31(0)10 408 8900

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*[a,b]

[a]Vrije Universiteit Amsterdam, The Netherlands

[b]Tinbergen Institute, The Netherlands

[c]CREATES, Aarhus University, Denmark

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 a low-frequency time series variable through the use of timely information from high-frequency variables. We aim to verify in-sample and out-of-sample performances of the model in an empirical study on the forecasting of U.S. headline inflation. In particular, we forecast monthly inflation using daily oil prices and quarterly inflation using effective federal funds rates. 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 point and density forecasts.

Keywords: Factor model; GAS model; Inflation forecasting; MIDAS; Score-driven model; Weighted maximum likelihood.

*Corresponding author: SJ Koopman, Department of Econometrics, Vrije Universiteit Amsterdam School of Business and Economics, De Boelelaan 1105, 1081 HV Amsterdam, The Netherlands. `s.j.koopman@vu.nl`

1 Introduction

In many 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 Armesto et al. (2010) and Andreou et al. (2011). In particular, in cases of economic forecasting where both economic and financial variables are relevant, the distinction between low frequency and high frequency data sampling can be substantive. 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) because it is relatively straightforward to electronically record financial transactions. On the other hand, it is more complicated and more costly to collect data on economic variables, such as inflation and gross domestic product (GDP) growth. Hence economic variables are typically available at a quarterly or monthly level. When the interest is in the forecasting of economic variables, the high-frequency financial variables may have a relevant role to play as predictors and may be capable to improve the accuracy of forecasts.

A widely used method for incorporating high frequency data to produce forecasts of low frequency variables is the Mixed Data Sampling (MIDAS) method of Ghysels et al. (2004). MIDAS is a regression-based method that transforms the high frequency variables into low frequency indicators via a parsimonious weighting scheme with possibly different weights to data sampled at high frequency (within the low frequency period). The weighting scheme can reflect the notion that more recent observations should be more informative to predict 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 such as vector autoregressive and dynamic factor models. For instance, Marcellino and Schumacher (2010) have considered a two step approach that combines principal component analysis and MIDAS regressions.

In our current study we adopt a dynamic model with score-driven time-varying location and scale parameters. Creal et al. (2013) and Harvey (2013) have developed a general framework to specify 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 and conditional on past and current time periods. The score-driven updates have an optimality property. Blasques et al. (2015) show that the score update is optimal in minimizing the Kullback-Leibler divergence with respect to an unknown true distribution. Score-driven models provide an appealing forecasting method and they have been successfully employed in empirical applications to forecast economic and financial variables, see for instance Delle Monache and Petrella (2017) on forecasting inflation and Lucas and Zhang (2016) and Blasques et al. (2016a) on forecasting exchange rates. In a more general context, the forecasting performance of GAS models is investigated in detail by Koopman et al. (2016). Finally, GAS models are appealing because they are flexible in terms of 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 computational-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 where the score innovations from the high frequency variables are transformed into the low frequency score function via a MIDAS weighting scheme. We name the resulting approach MIDAS-GAS. The MIDAS-GAS model retains all 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 heteroscedastic errors and parameter updates that are robust against outliers. Furthermore, we adopt the weighted likelihood approach of Blasques et al. (2016b) 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 MIDAS-GAS framework to produce forecasts for monthly and quarterly U.S. headline inflation. In particular, for forecasting monthly inflation we consider daily crude oil inflation as a predictor, and for quarterly inflation we take daily effective federal funds rate (EFFR) as a predictor. We present a detailed account of the forecasting study that includes comparisons with several competing models such as MIDAS regression models, autoregressive models and standard GAS models. The results show a promising performance of the MIDAS-GAS model in terms of point as well as density forecasts.

An alternative approach to MIDAS-based methods for the treatment of mixed frequency data is provided by state space time series analyses which rely on the Kalman filter. In this approach we align the data at the highest data sampling frequency and introduce missing observations for the low frequency variables. The Kalman filter is then used to handle these artificial missing observations, see Mariano and Murasawa (2003), Schumacher and Breitung (2008) and Blasques et al. (2016b) for such solutions and with interesting illustrations. A limitation of this more rigorous approach compared to our MIDAS-GAS model is that the Kalman filter requires Gaussian and homoscedastic errors. There is much empirical evidence that shows the importance of accounting for heteroscedastic errors and fat-tailed distributions to obtain more accurate forecasts for economic time series, see, for example, Creal et al. (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 with the forecasting of both quarterly and monthly U.S. headline inflation. Section 5 concludes.

2 The MIDAS-GAS model

Assume that our aim is to forecast a key economic variable 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 variable is not of interest but we assume that it can be exploited to 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 the 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 equal to the number of working days in a month. For notational convenience and simplicity of exposition, we assume that the both variables y and x are univariate. However, all results discussed below can be extended to the multivariate case straightforwardly.

2.1 The MIDAS touch

Amongst the range of forecasting methods using mixed frequency data, the MIDAS regression is regarded as a simple and direct forecasting method. Denote the h-step ahead forecast of y_T^L by $\widehat{y}_{T+h|T}^L$ where T denotes the sample size. This forecast can be constructed by the p-lag MIDAS regression

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

for $t =, \ldots, T$, where c is the intercept, $D_p(B, z) = z_0 + z_1 B + \ldots + z_p B^p$, for $z = \beta, \alpha$ is the lag polynomial function with backshift operator B, unknown parameter vectors $\beta = (\beta_0, \ldots, \beta_p)'$ and $\alpha = (\alpha_0, \ldots, \alpha_p)'$, and weighting coefficients $\omega_i(\varphi)$ for $i = 1, \ldots, n_x$ and for a parameter vector φ, and ϵ_t is an identical independently distributed (iid) error with mean zero and variance σ^2. Ghysels et al. (2004) advocate a parsimonious weighting function for $\omega_i(\varphi)$, for $i = 1, \ldots, n_x$, based on exponential Almon lag or Beta lag parameterizations. The q-th order exponential Almon lag is specified as

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

for some q-dimensional parameter vector $\varphi = (\varphi_1, \ldots, \varphi_q)'$. In practice, q is set equal to 2 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 Beta$(\cdot; \varphi_1, \varphi_2)$ is the probability density function of a Beta distribution with parameter vector $\varphi = (\varphi_1, \varphi_2)'$. Figure 1 presents illustrations of weighting functions based on second order exponential Almon lag and Beta lag; we can conclude that the shapes of these weighting functions are very flexible. The parameters of the MIDAS regression include c, β, α, φ and σ^2 which can be estimated using either nonlinear least squares or maximum likelihood (ML) and based on the MIDAS regression model (1). We obtain the h-step ahead forecast $\widehat{y}^L_{T+h|T}$ in the usual way. Andreou et al. (2011), and the references therein, present a more detailed discussion on MIDAS regressions together with empirical applications.

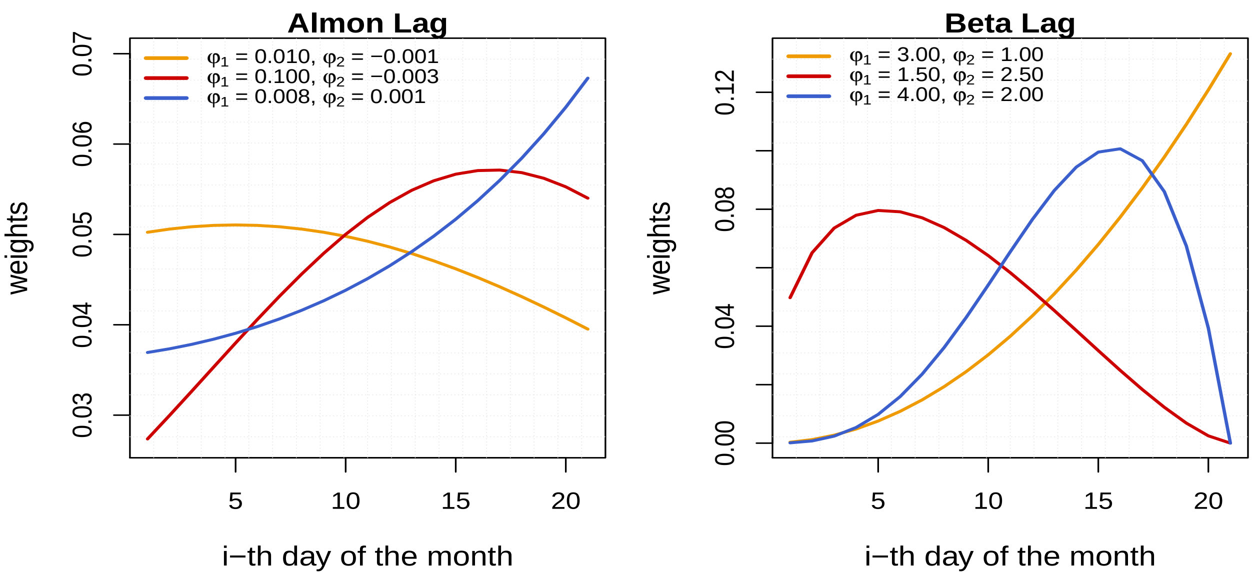

Figure 1: Weighting functions based on exponential Almon lag and Beta lag. The weights are assigned to daily observations within a month using different parameters of the exponential Almon lag (left) and the Beta lag (right) functions.

2.2 The MIDAS-GAS filter

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

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

$$y_t^L, x_t^H | f_t \sim p(y_t^L, x_t^H | f_t; \psi), \tag{2}$$

where $p(\cdot|f_t;\psi)$ is a (n_x+1)-variate conditional density, f_t is a time-varying parameter and ψ is a vector of static parameters. Further we assume that y_t^L is independent of x_t^H conditional on f_t and also that the elements of the vector x_t^H are iid conditional on f_t. These assumptions are quite standard in the literature of multivariate GAS models and also state-space models in general. Note also that the conditional independence does not imply that the variables are independent because the dependence is determined by the common time-varying parameter f_t. Under these conditions, the joint conditional density of y_t^L and x_t^H can be factorized as follows

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

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 $x_{i,t}^H$.

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

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

where $\nabla_t^y = \partial \log p_y(y_t^L|f_t;\psi)/\partial f_t$ and $\nabla_{i,t}^x = \partial \log p_x(x_{i,t}^H|f_t;\psi)/\partial f_t$. 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

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

where $s_t^y = S_t^y \nabla_t^y$ and $s_{i,t}^x = S_{i,t}^x \nabla_t^{x_i}$ for some given scaling factors S_t^y and $S_{i,t}^x$. For instance, these scaling factors can be chosen to be some transformation of the Fisher information or they can be simply set equal to one. The score innovation s_t is easy to interpret: s_t^y can be seen as the information from y_t^L to update the time-varying parameter f_t, similarly, $s_{i,t}^x$ can be seen as the information from $x_{i,t}^H$ to update f_t. In standard GAS models, s_t is directly used as innovation for the time varying parameter f_t. This means that the source of information are equally weighted over the score contributions $s_{i,t}^x$, $i = 1, \ldots, n_x$. Since our objective is forecasting the variable y_t^L, it is quite restrictive to assume that the predictive content carried by the elements $s_{i,t}^x$, $i = 1, \ldots, n_x$, is the same. For instance, it is reasonable to think that the latest score innovations withing the time period t may be more informative in predicting y_{t+1}^L, given that they are closer in time. We therefore introduce a MIDAS weighting scheme for the score innovations $s_{i,t}^x$, that is, $\sum_{i=1}^{n_x} \omega_i(\varphi) s_{i,t}^x$. In this way, the more recent score innovations are allowed to receive more weight. The resulting MIDAS-GAS filter takes the simple form

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

where δ, β, α_y, α_x and φ are static parameters to be estimated. The parameter β is the autoregressive coefficient of f_t and the parameters α_y and α_x determine the relative importance of y_t^L and x_t^H in predicting future values of y_t^L. Notice that the autoregressive dynamics of f_t in (4) is of order one, but it is straightforward to extend to higher orders. Furthermore, in general, f_t can be a vector of time-varying parameters.

The MIDAS-GAS model specified by equations (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 and densities such as the ordered logit are considered. We refer the reader to Creal et al. (2014) for possible applications in this context. We present some MIDAS-GAS specifications with dynamic mean and variance in Section 3. These specifications are well suited to obtain point and density forecasts of economic variables.

2.3 Weighted likelihood estimation

One of the appealing features of GAS models is that the likelihood function is available in closed form through a prediction error decomposition. Therefore ML estimation is easy to implement and computationally fast. This is also the case for our MIDAS-GAS model in (2)-(4). For our MIDAS-GAS model, we need to estimate the parameter vector $\theta = (\psi', \delta, \beta, \alpha_y, \alpha_x, \varphi')'$ that contains all the static parameters of the model. In the following, instead of ML, we consider a weighted maximum likelihood (WML) approach as proposed by Blasques et al. (2016b). Blasques et al. (2016b) show that in a Gaussian dynamic factor model with mixed frequency data, forecasting can be improved using parameters estimated via a weighted maximum likelihood (WML) procedure. Its novelty is to introduce variable-specific weights instead of observation-specific weights in the likelihood function. In our setting, this method is particularly appealing because we are only interested in forecasting the low frequency variable y_t^L and the high frequency variable x_t^H is just used as a predictor. In our framework, following Blasques et al. (2016b), the weighted likelihood of the MIDAS-GAS model can be written 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 $W \in [0, 1]$. If the weigh W is equal to one, then the weighted likelihood function is equal to the usual likelihood function, instead, if W is equal to zero, then only the likelihood of the variable y_t^L is considered. We highlight that setting $W = 0$ may lead to lack of identifiability of same parameters and this needs to be accounted when dealing with specific models. Some parameter restrictions to ensure identifiability are discussed in Section 4. Maximization of the weighted likelihood function can be done via standard numerical routines. The weight W cannot be estimated together with the other parameters. A cross validation approach can be used to select the weight. In the empirical application of Section 4, we set the weight to zero and therefore consider only the likelihood contribution of the univariate time series y_t^L. As shown in Blasques et al. (2016b), in principle, a different choice of the weight may provide better out-of-sample results. However, we consider $W = 0$ to obtain a simple form of the likelihood function and to have a more realistic forecasting comparison with other models. In case we set W equal to zero, the MIDAS-GAS model can be regarded as a univariate model for y_t^L of the form

$$y_t^L \sim p_y(y_t^L | f_t; \psi), \qquad f_{t+1} = \delta + \beta f_t + \alpha_y s_t^y + \alpha_x \sum_{i=1}^{n_x} \omega_i(\varphi) s_{i,t}^x.$$

Although the weight W is set to zero, the high frequency variable x_t^H still enters into the MIDAS-GAS time-varying parameter f_t through the score innovations $s_{i,t}^x$. This model can be viewed a univariate GAS model with an exogenous predictor x_t^H. Therefore, standard asymptotic theory for ML estimation of GAS models in Blasques et al. (2014) applies in this case.

3 A MIDAS-GAS factor model

In this section, we develop some factor models for mixed frequency data based on the general MIDAS-GAS specification in (2)-(4). We can formulate a factor forecasting model, which is essentially a multivariate time-varying location and scale model, with mixed frequency data as follows

$$\begin{bmatrix} y_t^L \\ x_t^H \end{bmatrix} = \begin{bmatrix} 1 \\ \lambda_\mu \mathbf{1}_{n_x} \end{bmatrix} \mu_t + \sigma_t \begin{bmatrix} \epsilon_t^y \\ \epsilon_t^x \end{bmatrix}, \tag{5}$$

where ϵ_t^y and $\epsilon_t^x = (\epsilon_{1,t}^x, \ldots, \epsilon_{n_x,t}^x)'$ are independent disturbance terms that follow some parametric distributions. The error ϵ_t^y has mean zero and unit variance, instead, $\epsilon_{i,t}^x$ has mean zero and variance λ_σ. The vector $\mathbf{1}_{n_x}$ is a vector of length n_x that contains ones, and λ_μ is a loading on the factor μ_t for x_t^H. The specification in (5) entails a factor structure for both μ_t and σ_t^2 where λ_σ and λ_μ are the corresponding loadings for the high frequency variable.

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

$$\begin{aligned} \mu_{t+1} &= \delta_\mu + \beta_\mu \mu_t + \alpha_\mu^y s_t^y + \alpha_\mu^x \sum_{i=1}^{n_x} \omega_i(\varphi) s_{i,t}^x, \\ \sigma_{t+1}^2 &= \delta_\sigma + \beta_\sigma \sigma_t^2 + \alpha_\sigma^y v_t^y + \alpha_\sigma^x \sum_{i=1}^{n_x} \omega_i(\varphi) v_{i,t}^x, \end{aligned} \tag{6}$$

where s_t^y and $s_{i,t}^x$ are the score contributions for the time varying mean μ_t and 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 and 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 the following we consider two different specifications for the error terms: the

Normal and the Student-t distribution. The use of a Student-t distribution leads to robust updating of the time varying mean and variance.

In the first specification, we consider the errors to be normally distributed, i.e. $\epsilon_t^y \sim N(0,1)$ and $\epsilon_{i,t}^x \sim N(0,\lambda_\sigma)$. Under this specification and considering a scaling factor for the scores based on the Fisher information, we obtain that, up to a constant scale, the score innovations are given by

$$\begin{aligned} s_t^y &= y_t^L - \mu_t, & s_{i,t}^x &= x_{i,t}^H - \lambda_\mu \mu_t, \\ v_t^y &= (y_t^L - \mu_t)^2 - \sigma_t^2, & v_{i,t}^x &= (x_{i,t}^H - \lambda_\mu \mu_t)^2 - \lambda_\sigma \sigma_t^2. \end{aligned} \tag{7}$$

In the second specification, we model ϵ_t^y considering a Student-t distribution with zero mean and unit variance, i.e. $\epsilon_t^y \sim t_v(0,1)$. In this way, 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 error 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 shall see, the Student-t delivers a robust update for our MIDAS-GAS model and this can also lead to an improvement in point forecasts. Note that a Student-t distribution for the error term ϵ_t^x can also be considered. In the following, for simplicity, we consider a normal distribution for ϵ_t^x. However, we highlight that the choice of the distribution of ϵ_t^x is less relevant than the one of ϵ_t^y because we are not interested in forecasting x_t^H and also the estimation of the model is based on WML with $W = 0$. Under this second specification, we obtain that, up to some scaling constant, the score innovations have the following form

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

The score innovations s_t^y and v_t^y in (8) are robust against otliers. This is one of the peculiarities of the GAS approach when using fat-tailed distributions, see Harvey and Luati (2014) and Harvey (2013) for a discussion on 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. Secondly, the dynamic factor process is non-linear due to the score updating mechanism, which locally minimizes the KL divergence between the true measure of the data and the one implied by the factor model as discussed in Blasques et al. (2015). Third, the MIDAS touch incorporated in the dynamic factor process (6) further balances the predictive information from y_t^L and x_t^H. This differs from the MIDAS factor model of Marcellino and Schumacher (2010) which extract factors from x_t^H and treat the high frequency factor as regressors in the standard MIDAS regression. It is also different from the MIDAS dynamic factor model of Frale 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 require Gaussian errors and a linear structure so that principal component analysis (PCA) and Kalman filter can be used to extract factors and produce forecasts. The GAS filter in our model has the appealing feature that it does not limit us in a Gaussian and linear world. In time series models, density forecasts are usually improved by incorporating stochastic volatility. This has been extensively documented by the rich literature studying the forecasting performance of models with stochastic volatility, for example Kim et al. (1998), Tse and Tsui (2002), and Chib et al. (2002) among others. However, estimation of such models is computationally demanding and Bayesian methods are often used. Instead, the estimation of the MIDAS-GAS is straightforward.

4 Forecasting headline inflation

In this section, we employ the MIDAS-GAS factor model for forecasting U.S. headline inflation. We consider two applications. In the first application, we forecast monthly headline inflation using daily crude oil inflation. Instead, in the second application, we forecast quarterly headline inflation using daily EFFR. In this way, we can access the performance of the MIDAS-GAS model with different frequencies and different variables as predictors.

4.1 Forecasting monthly inflation with daily oil prices

4.1.1 The dataset and in-sample results

There is evidence in the literature that oil prices have relevant predictive content for U.S. inflation, see Clark and Terry (2010) and Stock and Watson (2003) for instance. We consider time series of monthly U.S. headline inflation and crude oil inflation from January 1986 to November 2017. Figure 2 shows the plot of the time series.

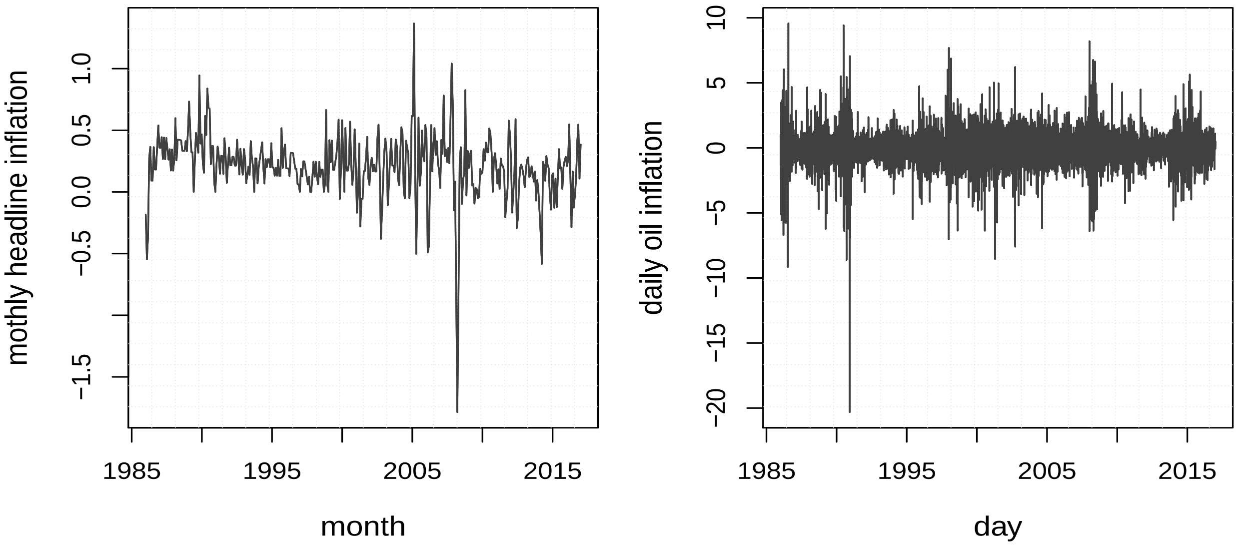

Figure 2: The headline CPI inflation and oil price inflation. Left: The monthly headline inflation computed as the first difference of the logarithm of monthly CPI indices. Right: The daily oil price inflation computed as the first difference of the logarithm of daily WTI crude oil prices.

We consider the MIDAS-GAS factor model presented in Section 3 based on WML with $W = 0$. We impose some parameter restrictions. In particular, we consider a random walk process for the time varying mean by imposing $\beta_\mu = 1$ and $\delta_\mu = 0$. Furthermore, we set $\lambda_\sigma = 1$ to ensure identifiability. Note that λ_σ could be set equal to any positive value and the resulting model would 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 above restrictions can be written as univariate model for y_t^L of the form

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

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

We estimate 4 different specifications of the MIDAS-GAS model in (9). The t-MIDAS-GASg has a Student-t error for y_t^L and the score innovations given in (8), the t-MIDAS-GAS is the same as the t-MIDAS-GASg but without conditional heteroscedasticity, i.e. $\beta_\sigma = \alpha_\sigma^x = \alpha_\sigma^y = 0$, the MIDAS-GASg has a normal error for y_t^L and the score innovations given in (7), and the MIDAS-GAS is the same as the MIDAS-GASg but without conditional heteroscedasticity, i.e. $\beta_\sigma = \alpha_\sigma^x = \alpha_\sigma^y = 0$. Table 1 reports the estimation results. We can see that the

	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	4.92	0.18	0.14	0.08	0.34	0.00	0.02	0.78	294.86	**-569.72**
t-MIDAS-GAS	4.11	1.24	0.35	0.18	-	-	0.05	-	265.20	-516.40
MIDAS-GASg	-	0.31	0.13	0.10	0.26	0.00	0.04	0.91	162.04	-306.08
MIDAS-GAS	-	1.76	0.33	0.15	-	-	0.05	-	122.81	-233.62

Table 1: Full-Sample parameter estimates of MIDAS-GAS factor models using monthly inflation with daily oil prices. The last two columns report the log-likelihood AIC criterion, respectively.

Student-t distribution improves the in-sample fit. This can be noted from the better performance of the models with Student-t error in terms of Akaike information criterion (AIC). Additionally, the estimated degrees of freedom parameters ν are very small (around 4), indicating the presence of fat tails. The better fit of the Student-t is not surprising since the headline inflation series exhibits several extreme observations, see Figure 2. Furthermore, we also note that the conditional heteroscedasticity of the error term delivers a clear in-sample improvement in terms of AIC. Also this finding is coherent with the volatility clustering of the inflation series that we can see in Figure 2.

Finally, Figure 3 reports the estimated MIDAS weighting functions for the four different model specifications. We can see that estimated functions give more weight to the more recent observations of the high frequency variable. This result is coherent with the idea that observations that are farther away in time should be less relevant for forecasting the future values of the variable of interest. Furthermore, we also note that the results seem consistent across the different models since the shape of the four weighting functions is very similar.

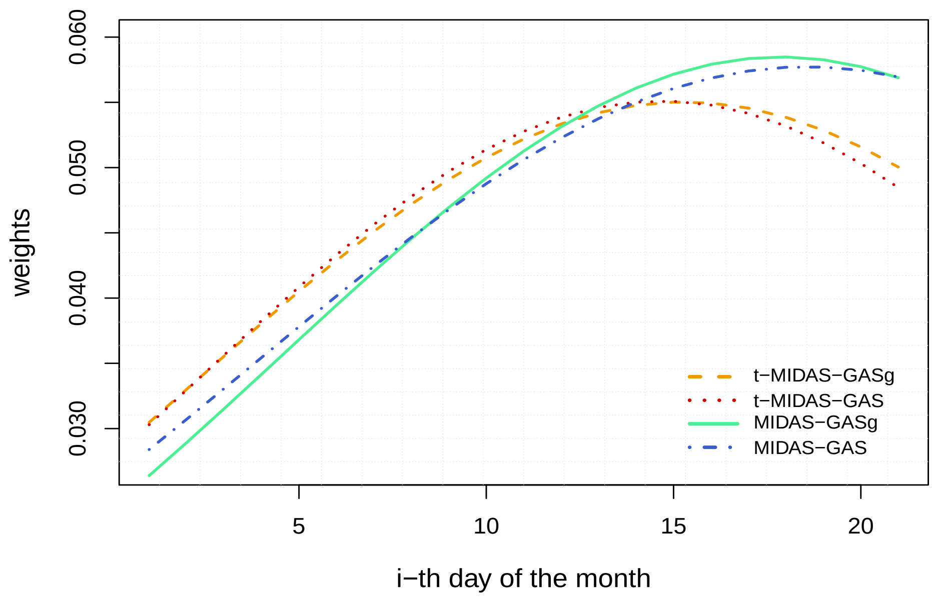

Figure 3: Estimated MIDAS-GAS weighting functions for the four model specifications using monthly inflation with daily oil prices. The horizontal axis indicates the day in the month (there are about 21 working days in a month). More weight on the last days of the month indicates that more recent observations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

4.1.2 Out-of-sample exercise

We consider two forecasting tasks: point forecast and density forecast. Point forecast is usually the central task carried out in central banks. For example, inflation forecast facilitates the use of forward-looking monetary policy which helps calculate the ex ante real interest rate to determine the aggregate demand or IS curve for an economy. Additionally, density forecast is important because it provides a risk metric measuring how accurate the point forecast is. We evaluate the performance of point forecasts using the Forecast Mean Squared Error (FMSE) and 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).

We split the full data sample, which consists of 383 months, into two subsamples: the first 200 months are used as in-sample training period and the remaining months are used for the out-of-sample evaluation period. We consider a rolling window forecasting exercise. Therefore the length of the in-sample estimation period is equal to 200 for all forecasts. We consider multi-step forecasts: from 1 step ahead to 6 steps ahead. Besides our 4 MIDAS-GAS models, we include several competing models in the comparison. We include 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 illustrate the specification of the competing models included in the comparison. We note that the MIDAS regression models and the autoregressive models are estimated as a direct forecasting method for each forecasting horizon.

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

Table 2: Specification of the competing models used in the out-of-sample exercise.

Table 3 reports the results of the forecasting study. We can see that MIDAS-GAS models tend to have the best performance in terms of point forecasts, except for 1 step ahead forecasts. Furthermore, we note that the inclusion of conditional heteroscedasticity and Student-t errors also plays a major. This can be noted from the fact that the t-MIDAS-GASg model tends to have the best performance among the 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 deliver accurate forecasts compared to a wide pool of competing models.

	Forecast Mean Squared Error						Log score criterion					
	$h=1$	$h=2$	$h=3$	$h=4$	$h=5$	$h=6$	$h=1$	$h=2$	$h=3$	$h=4$	$h=5$	$h=6$
t-MIDAS-GASg	1.00	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	0.43	**0.27**	**0.15**	**0.14**	**0.15**	0.11
t-MIDAS-GAS	0.93	1.04	1.06	1.06	1.08	1.06	0.38	0.17	0.05	0.05	0.11	0.05
MIDAS-GASg	1.00	1.10	**1.00**	1.06	1.09	1.02	-0.42	-0.94	-1.10	-1.13	-0.92	-1.02
MIDAS-GAS	0.93	1.17	1.18	1.16	1.18	1.11	-0.55	-1.00	-1.12	-1.18	-1.14	-1.26
t-MIDASg(2)	0.91	**1.00**	1.04	1.04	1.03	1.03	**0.49**	0.21	0.11	0.01	0.21	-0.06
t-MIDAS(2)	0.93	1.08	1.05	1.09	1.06	1.05	0.30	0.20	-0.11	-0.00	0.05	-0.09
MIDASg(2)	**0.87**	1.03	1.08	1.09	1.10	1.11	-0.40	-0.91	-1.19	-1.25	-1.19	-1.18
MIDAS(2)	0.93	1.04	1.21	1.19	1.20	1.18	-0.73	-1.22	-1.28	-1.43	-1.32	-1.34
t-MIDASg(4)	0.99	1.08	1.04	1.01	1.02	1.04	0.47	0.24	0.12	0.12	0.14	**0.12**
t-MIDAS(4)	0.99	1.12	1.07	1.02	1.02	1.08	0.31	0.11	0.06	-0.03	0.05	-0.00
MIDASg(4)	0.94	1.11	1.12	1.08	1.17	1.06	-0.38	-0.87	-1.12	-1.21	-1.09	-1.12
MIDAS(4)	0.96	1.14	1.13	1.13	1.19	1.12	-0.67	-1.26	-1.19	-1.47	-1.31	-1.27
t-ARg(2)	0.97	1.12	1.13	1.13	1.15	1.15	0.43	0.19	0.11	0.10	0.09	-0.07
t-AR(2)	0.97	1.17	1.17	1.19	1.19	1.17	0.26	0.18	-0.14	-0.08	0.05	-0.14
ARg(2)	1.01	1.22	1.19	1.19	1.25	1.16	-0.47	-0.83	-1.25	-1.36	-1.23	-1.21
AR(2)	1.03	1.24	1.24	1.26	1.25	1.26	-0.72	-1.20	-1.33	-1.44	-1.38	-1.43
t-ARg(4)	0.98	1.16	1.13	1.15	1.19	1.16	0.42	0.15	0.10	0.03	0.04	-0.09
t-AR(4)	0.98	1.22	1.20	1.19	1.23	1.21	0.30	0.11	0.04	-0.11	-0.06	-0.13
ARg(4)	0.93	1.10	1.10	1.10	1.13	1.10	-0.35	-0.91	-1.16	-1.24	-1.03	-1.11
AR(4)	0.92	1.16	1.15	1.18	1.16	1.18	-0.69	-1.20	-1.27	-1.44	-1.41	-1.37
t-GASg	1.01	1.11	1.12	1.11	1.12	1.14	0.20	0.11	0.04	0.09	0.08	0.05
t-GAS	0.95	1.11	1.20	1.24	1.21	1.20	0.34	0.15	0.03	0.02	0.04	-0.02
GASg	0.91	1.15	1.14	1.18	1.14	1.18	-0.39	-0.91	-1.33	-1.35	-1.12	-1.21
GAS	0.93	1.21	1.22	1.26	1.22	1.23	-0.46	-1.03	-1.34	-1.51	-1.23	-1.31
fMIDAS	0.97	1.13	1.12	1.14	1.14	1.09	-0.48	-0.80	-1.21	-1.19	-1.04	-1.27

Table 3: Relative MSE and log score criterion for the different model specifications using monthly inflation with daily oil prices. The first 6 columns of the table report the ratio between the MSE of each model with respect to the MSE of benchmark model (t-MIDAS-GASg) for several forecasting horizons (from 1 to 6 steps ahead). A value greater than 1 indicates that a model is underperforming the benchmark model instead the opposite is true when the MSE ratio is smaller than 1. The last 6 columns report the log score criterion for several forecasting horizons (from 1 to 6 steps ahead).

4.2 Forecasting quarterly inflation with EFFR

4.2.1 The dataset and in-sample results

In the following, we report the results of a second empirical study where we employ the MIDAS-GAS model in (9) for forecasting quarterly U.S. headline inflation using daily EFFR. In this case, we consider quarterly instead of monthly inflation to see how the MIDAS-GAS model performs at different frequencies. We also use a different high frequency variable: the EFFR, which is available daily since 1955. The use of EFFR to forecast inflation is quite natural since the EFFR is one of the main tools used by the Federal Reserve to control inflation, see Armesto et al. (2010) for an application on predicting inflation with EFFR. The time series we consider are from the first quarter of 1955 to the first quarter of 2017.

Figure 4 shows the quarterly headline inflation series and the daily EFFR series.

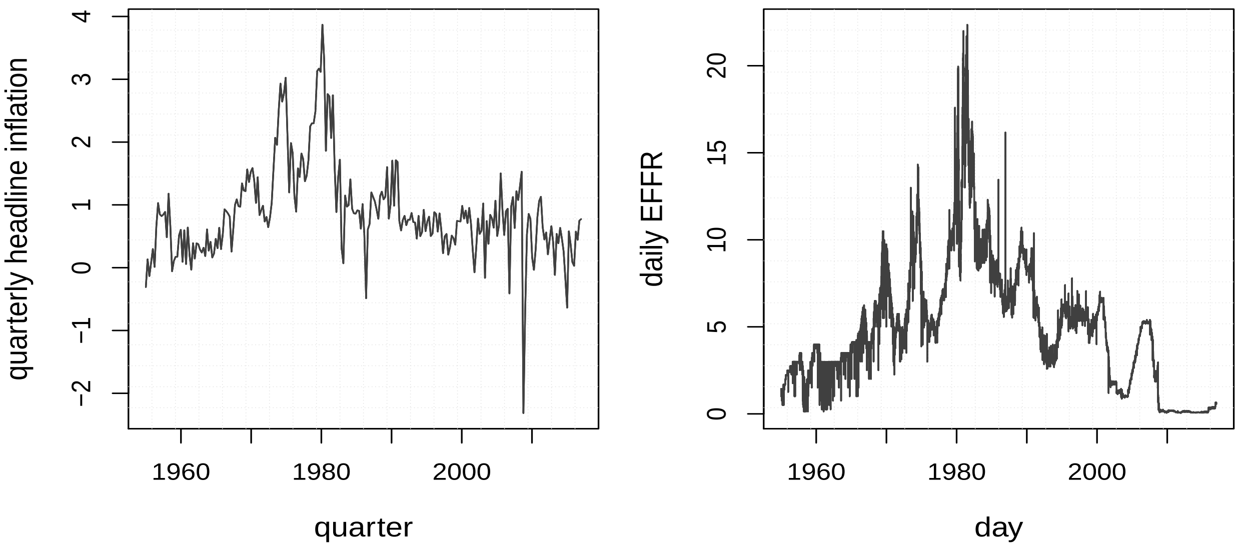

Figure 4: The headline CPI inflation and EFFR series. Left: The quarterly headline inflation. Right: The daily effective federal founds rate series.

	ν	λ^x_μ	α^y_μ	α^x_μ	α^y_σ	α^x_σ	δ_σ	β_σ	llik	AIC
t-MIDAS-GASg	7.36	1.49	0.46	0.14	0.23	0.07	2.21	0.87	-319.83	**659.66**
t-MIDAS-GAS	5.16	1.38	0.45	0.13	-	-	2.87	-	-333.89	681.78
MIDAS-GASg	-	1.62	0.54	0.11	0.23	0.21	2.31	0.81	-476.68	971.36
MIDAS-GAS	-	1.61	0.51	0.09	-	-	3.44	-	-511.36	1034.72

Table 4: Full-Sample parameter estimates of MIDAS-GAS factor models using quarterly inflation with daily EFFR. The last column reports the log-likelihood of the estimated models. The last two columns report the log-likelihood AIC criterion, respectively.

Table 4 reports the estimates of the MIDAS-GAS models. The in-sample e results are in line with those obtained for the monthly inflation. In particular, the Student-t distribution gives a better fit since outliers are present in the series, see Figure 4. Furthermore, also including conditional heteroscedasticity improves the in-sample results. Finally, Figure 5 reports the estimated MIDAS weighting functions. We can see that estimated functions give much more weight to the more recent observations of the high frequency variable. This finding shows how different variables can lead to different weighting schemes. A possible interpretation of this high weight on the very last observations of the quarter is that the EFFR is adjusted by the Federal Reserve to account for some economic and financial conditions and therefore most predictive content of the EFFR for inflation is lagged to the last 20 days of the quarter.

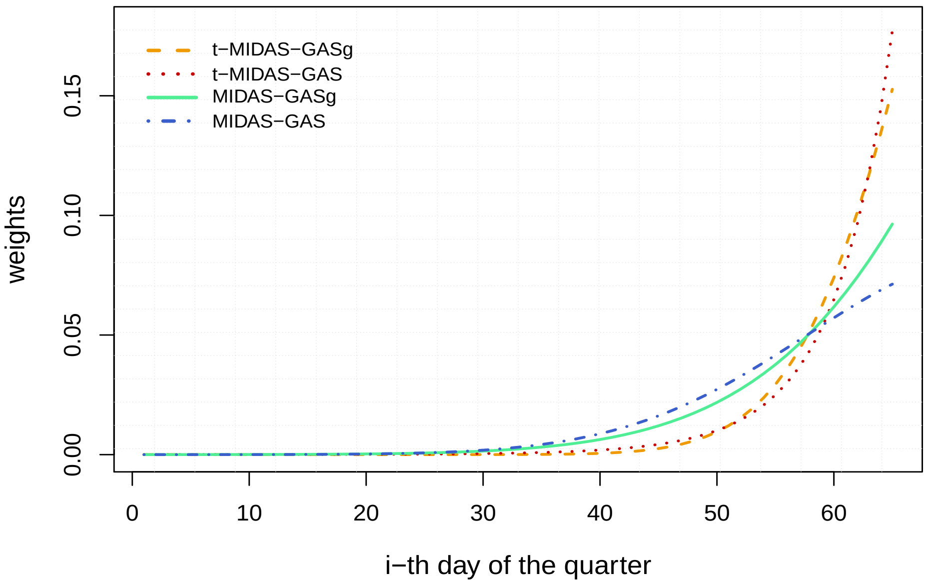

Figure 5: Estimated MIDAS-GAS weighting functions for the four model specifications using quarterly inflation with daily EFFR. The horizontal axis indicates the day in the quarter (there are about 63 working days in a quarter). More weight on the last days of the quarter indicates that more recent observations are more relevant.

4.2.2 Out-of-sample exercise

We perform a rolling window forecasting exercise as considered in the previous empirical application. The out-of-sample period is from 1993 to 2017. Table 5 reports the results of the forecasting study for the quarterly inflation series using daily EFFR as predictor. We can see that the MIDAS-GAS models tend to have the best performance in terms of point forecasts, except for 1 step ahead forecasts. In this case however the MIGAS-GASg seems to have the best performance for most horizons. As concerns the density forecasts, we see that the MIDAS-GAS models have the best performance for all forecasting horizons and in this case the Student-t distribution plays a major role. As before, we conclude that MIDAS-GAS models give accurate forecasts.

	Forecast Mean Squared Error						Log score criterion					
	$h=1$	$h=2$	$h=3$	$h=4$	$h=5$	$h=6$	$h=1$	$h=2$	$h=3$	$h=4$	$h=5$	$h=6$
t-MIDAS-GASg	1.00	1.00	**1.00**	1.00	1.00	1.00	-1.42	**-1.55**	**-1.48**	**-1.61**	**-1.59**	**-1.58**
t-MIDAS-GAS	1.03	1.00	1.02	1.04	1.05	1.02	**-1.41**	-1.56	-1.55	-1.64	-1.62	-1.67
MIDAS-GASg	1.02	**0.99**	1.01	**0.95**	**0.96**	**0.95**	-2.23	-2.19	-2.18	-2.20	-2.20	-2.12
MIDAS-GAS	1.00	1.00	1.07	1.03	1.01	1.01	-2.22	-2.26	-2.21	-2.24	-2.22	-2.21
t-MIDASg(2)	**0.96**	1.05	1.10	1.02	1.04	1.09	-1.98	-2.12	-1.81	-1.94	-2.10	-1.84
t-MIDAS(2)	0.99	1.13	1.14	1.07	1.11	1.15	-1.76	-1.86	-2.02	-2.04	-1.98	-1.86
MIDASg(2)	0.98	1.12	1.13	1.04	1.06	1.09	-2.22	-2.25	-2.25	-2.25	-2.24	-2.26
MIDAS(2)	0.99	1.13	1.13	1.06	1.10	1.13	-2.30	-2.28	-2.34	-2.32	-2.31	-2.32
t-MIDASg(4)	1.04	1.12	1.09	1.06	1.13	1.13	-1.82	-1.69	-2.00	-1.82	-1.66	-1.85
t-MIDAS(4)	1.02	1.13	1.14	1.07	1.16	1.15	-1.81	-1.73	-2.02	-1.84	-1.79	-2.18
MIDASg(4)	1.02	1.11	1.10	1.02	1.10	1.12	-2.25	-2.25	-2.24	-2.25	-2.27	-2.27
MIDAS(4)	1.01	1.13	1.13	1.08	1.15	1.14	-2.33	-2.27	-2.36	-2.29	-2.30	-2.33
t-ARg(2)	1.06	1.09	1.10	1.03	1.14	1.14	-1.78	-1.95	-2.12	-1.88	-1.83	-1.90
t-AR(2)	1.03	1.12	1.13	1.03	1.16	1.14	-1.77	-2.13	-2.20	-1.94	-1.88	-2.32
ARg(2)	0.98	1.09	1.08	1.00	1.12	1.13	-2.17	-2.18	-2.26	-2.41	-2.16	-2.18
AR(2)	1.00	1.10	1.09	1.02	1.12	1.14	-2.28	-2.24	-2.36	-2.43	-2.26	-2.22
t-ARg(4)	1.07	1.12	1.11	1.04	1.16	1.13	-1.75	-1.99	-2.11	-1.94	-1.81	-1.82
t-AR(4)	1.04	1.13	1.15	1.04	1.16	1.15	-1.76	-2.12	-2.15	-1.95	-1.84	-2.38
ARg(4)	0.99	1.10	1.11	1.06	1.09	1.12	-2.25	-2.23	-2.29	-2.44	-2.27	-2.27
AR(4)	0.98	1.10	1.14	1.06	1.16	1.13	-2.27	-2.26	-2.35	-2.43	-2.27	-2.27
t-GASg	1.04	1.06	1.01	1.02	1.02	1.00	-1.59	-1.74	-1.72	-1.86	-1.86	-1.86
t-GAS	1.06	1.09	1.03	1.04	1.05	1.03	-1.71	-1.86	-1.83	-1.98	-1.97	-1.95
GASg	0.97	1.08	1.04	1.03	1.01	0.96	-2.13	-2.20	-2.21	-2.26	-2.22	-2.17
GAS	1.01	1.12	1.07	1.03	1.02	1.01	-2.22	-2.29	-2.26	-2.28	-2.25	-2.25
fMIDAS	1.00	1.11	1.07	1.03	1.01	1.01	-2.22	-2.26	-2.31	-2.30	-2.28	-2.30

Table 5: Relative MSE and log score criterion for the different model specifications using quarterly inflation with daily EFFR. The first 6 columns of the table report the ratio between the MSE of each model with respect to the MSE of benchmark model (t-MIDAS-GASg) for several forecasting horizons (from 1 to 6 steps ahead). A value greater than 1 indicates that a model is underperforming the benchmark model instead the opposite is true when the MSE ratio is smaller than 1. The last 6 columns report the log score criterion for several forecasting horizons (from 1 to 6 steps ahead).

5 Conclusion

In this paper, we have introduced a novel approach for forecasting with mixed frequency data: the MIDAS-GAS model. The MIDAS-GAS model transforms the score contributions of the high frequency variables through a MIDAS weighting scheme. The proposed approach has several advantages as it retains all the appealing features of GAS models while accounting for mixed frequencies. Based on the general MIDAS-GAS framework, we have developed a novel forecasting model with dynamic factor structures for mean and variance. The method has shown a promising forecasting performance in two empirical applications on forecasting U.S. headline inflation using crude oil prices and EFFR. Finally, the MIDAS-GAS modeling framework can also be used for nowcasting. We leave this for future research.

References

Andreou, E., Ghysels, E., and Kourtellos, A. (2011). Forecasting with mixed-frequency data. In Clements, M. P. and Hendry, D. F., editors, *The Oxford Handbook of Economic Forecasting*. Oxford University Press, Oxford.

Armesto, M. T., Engemann, K. M., and Owyang, M. T. (2010). Forecasting with mixed frequencies. *Federal Reserve Bank of St. Louis Review*, 92(6):521–36.

Blasques, F., Ji, J., and Lucas, A. (2016a). Semiparametric score driven volatility models. *Computational Statistics & Data Analysis*, 100:58–69.

Blasques, F., Koopman, S., Mallee, M., and Zhang, Z. (2016b). Weighted maximum likelihood for dynamic factor analysis and forecasting with mixed frequency data. *Journal of Econometrics*, 193(2):405–417.

Blasques, F., Koopman, S. J., and Lucas, A. (2014). Maximum likelihood estimation for generalized autoregressive score models. Technical report, Tinbergen Institute Discussion Paper.

Blasques, F., Koopman, S. J., and Lucas, A. (2015). Information-theoretic optimality of observation-driven time series models for continuous responses. *Biometrika*, 102(2):325–343.

Chib, S., Nardari, F., and Shephard, N. (2002). Markov chain Monte Carlo methods for stochastic volatility models. *Journal of Econometrics*, 108(2):281–316.

Clark, T. E. and 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., and Lucas, A. (2013). Generalized autoregressive score models with applications. *Journal of Applied Econometrics*, 28(5):777–795.

Creal, D., Schwaab, B., Koopman, S. J., and 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.

Delle Monache, D. and Petrella, I. (2017). Adaptive models and heavy tails with an application to inflation forecasting. *International Journal of Forecasting*, 33(2):482–501.

Frale, C. and Monteforte, L. (2011). FaMIDAS: A mixed frequency factor model with MIDAS structure. Technical report, Bank of Italy.

Geweke, J. and Amisano, G. (2011). Optimal prediction pools. *Journal of Econometrics*, 164(1):130–141.

Ghysels, E. and Marcellino, M. (2016). The econometric analysis of mixed frequency data sampling. *Journal of Econometrics*, 193(2):291–293.

Ghysels, E., Santa-Clara, P., and Valkanov, R. (2004). The MIDAS touch: Mixed data sampling regression models. *Technical Report*.

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. and Luati, A. (2014). Filtering with heavy tails. *Journal of the American Statistical Association*, 109(507):1112–1122.

Kim, S., Shephard, N., and Chib, S. (1998). Stochastic volatility: likelihood inference and comparison with ARCH models. *The Review of Economic Studies*, 65(3):361–393.

Koopman, S. J., Lucas, A., and Scharth, M. (2016). Predicting time-varying parameters with parameter-driven and observation-driven models. *The Review of Economics & Statistics*, 98(1):97–110.

Lucas, A. and Zhang, X. (2016). Score-driven exponentially weighted moving averages and value-at-risk forecasting. *International Journal of Forecasting*, 32(2):293–302.

Marcellino, M. and Schumacher, C. (2010). Factor MIDAS for nowcasting and forecasting with ragged-edge data: A model comparison for German GDP. *Oxford Bulletin of Economics & Statistics*, 72(4):518–550.

Mariano, R. S. and Murasawa, Y. (2003). A new coincident index of business cycles based on monthly and quarterly series. *Journal of Applied Econometrics*, 18(4):427–443.

Schumacher, C. and 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. and Watson, M. (2003). Forecasting output and inflation: The role of asset prices. *Journal of Economic Literature*, 41(3):788–829.

Tse, Y. K. and 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.