ARTICLE IN PRESS

Journal of Econometrics xxx (xxxx) xxx

Contents lists available at ScienceDirect

Journal of Econometrics

journal homepage: www.elsevier.com/locate/jeconom

Time-varying parameter MIDAS models: Application to nowcasting US Real GDP☆,☆☆

Joshua C.C. Chan [a], Aubrey Poon [b], Dan Zhu [c,*]

[a] Purdue University, United States of America
[b] University of Kent, United Kingdom
[c] Monash University, Australia

ARTICLE INFO

Keywords:
MIDAS regressions
Bayesian estimation Stochastic volatility
Precision-based methods
Nowcasting

ABSTRACT

We introduce a novel time-varying parameter mixed-data sampling (TVP-MIDAS) framework. Specifically, we allow both the MIDAS weights and the coefficients representing the overall impacts of the high-frequency variables to vary over time. This is done by introducing a class of linear parameterizations, which facilitate estimation in settings with a large number of high-frequency predictors. We demonstrate the usefulness of this framework via an application of nowcasting US GDP in real-time using monthly, weekly and daily predictors. The results show that the TVP-MIDAS models produce superior nowcasts, and are particularly effective in capturing the downside risk compared to their time-invariant counterparts.

1. Introduction

Mixed-data sampling (MIDAS) regressions have garnered significant attention in empirical macroeconomics for their utility in nowcasting key macroeconomic indicators such as real GDP and inflation. A key advantage of MIDAS is that it provides a straightforward and parsimonious framework for forecasting a variable of interest using predictors sampled at different frequencies. For example, a prevalent application of MIDAS regressions involves nowcasting the real GDP, a quarterly variable, using higher frequency predictors such as monthly industrial production or employment (e.g., Marcellino and Schumacher, 2010; Kuzin et al., 2011; Foroni and Marcellino, 2014; Mogliani and Simoni, 2021). GDP figures are released with a significant delay, and the MIDAS framework allows forecasters to exploit timely, high-frequency predictors to update the GDP nowcasts. More recently, MIDAS regressions have been applied in other areas beyond empirical macroeconomics, particularly in financial applications such as forecasting stock price volatility (see Andreou, 2016; Wang et al., 2020).

Standard applications of MIDAS regressions assume constant parameters, mostly due to necessity. This is because common parameterizations of the MIDAS weighting function, such as exponential Almon lag and beta polynomials, are nonlinear in the parameters. As such, extending them to be time varying typically involves the estimation of nonlinear state space models, which is computationally intensive. But this assumption of constant parameters appears to be overly restrictive given the documented importance of the role of time-varying parameters in forecasting macroeconomic variables (see Primiceri, 2005; Koop and Korobilis, 2013; D'Agostino et al., 2013; Barnett et al., 2014; Chan and Eisenstat, 2018). Furthermore, macroeconomic variables often exhibit both time-varying conditional means and volatilities, reflecting fluctuations in government policies, global economic conditions,

☆ This article is part of a Special issue entitled: 'Mixed Frequency Data' published in Journal of Econometrics.

☆☆ We thank Todd Clark, James Mitchell and Saeed Zaman for their constructive comments and helpful suggestions on a previous version of the paper.

* Corresponding author.
E-mail address: dan.zhu@monash.edu (D. Zhu).

https://doi.org/10.1016/j.jeconom.2025.106090
Received 15 July 2024; Received in revised form 30 May 2025; Accepted 25 August 2025

technological advancements, and other socio-economic factors. Comprehensive understanding and modeling of these time-varying distributions is imperative for accurate risk assessment and informed policy formulation, especially in the aftermath of the COVID-19 pandemic.

To address the limitations inherent in existing approaches, we propose a novel MIDAS framework that accommodates time-varying parameters, stochastic volatility and COVID-19 outliers. Specifically, we first introduce a class of linear parameterizations that are both flexible and conducive to fast estimation. The proposed setup can be motivated as finite-dimensional approximations of weighting functions using suitable basis functions. This setup includes the Almon polynomial, as well as many other basis functions, such as Fourier series and B-splines. A key advantage is that all these basis functions can be represented as linear regressions. As such, extending them to time-varying parameter settings is relatively straightforward. In addition to the time-varying weighting function, we also allow other coefficients, such as the scalar parameter representing the overall impact of the lags of the high-frequency variable, to be time-varying. In order to separately identify the overall impact parameter and the parameters in the weighting function, we develop an alternative identification scheme that preserves linearity and facilitates estimation. While normalization and identification are not strictly required in MIDAS regressions, they can prove beneficial in applications where the overall impact coefficient has an interesting economic interpretation, as exemplified in Ghysels et al. (2005).

In addition to time-varying parameters in the conditional mean, the proposed TVP-MIDAS framework also includes stochastic volatility and an explicit outlier component to address the extreme movements of many macroeconomic variables at the onset of the COVID-19 pandemic. Furthermore, we also discuss how the proposed framework can handle some complex but frequently encountered data issues. These include applications with predictors in multiple high frequencies (e.g., forecasting a quarterly variable with both monthly and weekly predictors), settings in which the numbers of observations of the high-frequency variables vary across time periods (e.g., a quarter has 13 or 14 weeks) and irregularly spaced mixed-frequency data (e.g., two weekly observations are available 3 and 10 days before the release of the monthly variable).

This paper is related to a few recent MIDAS studies, each offering distinct insights into modeling the dynamics between low- and high-frequency variables. Firstly, Potjagailo and Kohns (2023) propose a Bayesian MIDAS model incorporating a time-varying trend and stochastic volatility for nowcasting UK real GDP. Their model is a restricted version of the proposed TVP-MIDAS model, as theirs solely permits time variation in the intercept of the MIDAS regression, while maintaining a time-invariant structure for the MIDAS weighting function. Secondly, Guérin and Marcellino (2013) extend the MIDAS framework to accommodate parameter changes, albeit using a Markov switching model with two regimes. Thirdly, Schumacher (2014) develops a MIDAS regression with time-varying parameters, but the weighting function is parameterized using the exponential Almon polynomial, which is nonlinear in the parameters. As such, the estimation of the nonlinear state space model requires particle filtering techniques, which are computationally burdensome. Consequently, Schumacher (2014) only allows for time variation in a single high-frequency predictor in the MIDAS regression. In contrast, the proposed approach uses linear parameterizations of the weighting function, each of which defines a linear Gaussian state space model. As such, standard estimation approaches, such as the precision-based method in Chan and Jeliazkov (2009), can be used to estimate the proposed model. The computational efficiency and flexibility of this approach thus allow the researcher to consider multiple high-frequency predictors, with different time-varying weights and impact parameters.

We demonstrate the flexibility of the proposed TVP-MIDAS framework with a nowcasting application. Specifically, we use monthly, weekly and daily predictors to nowcast the quarterly US real GDP: the monthly industrial production, the weekly National Financial Condition Index (NFCI) and a daily interest rate spread (defined as the difference between the 10-year and 3-month treasury yields) that represents the slope of the yield curve. We also consider a larger dataset with 12 additional monthly predictors. For each quarter we generate three GDP nowcasts at the end of each month within the quarter, with our evaluation period spanning from 1990Q1 to 2021Q2.

The nowcast results indicate that the TVP-MIDAS models yield superior point and density forecasts compared to a variety of machine learning-based MIDAS models with constant parameters (Mogliani and Simoni, 2021; Babii et al., 2022). For example, using the model confidence set approach of Hansen et al. (2011), 8 models are included in $\widehat{\mathcal{M}}^*_{75\%}$, the model confidence set with 75% coverage, for density nowcasts (with forecast horizon $h = 0$) and evaluation period 1990Q1-2019Q4, and they all feature time-varying parameters. For the sample that includes the COVID-19 pandemic, 16 models are included in the model confidence set $\widehat{\mathcal{M}}^*_{75\%}$, and 11 feature time-varying parameters.

Furthermore, we explore the performance of the TVP-MIDAS models in nowcasting the downside risk of real GDP. Specifically, we find that the TVP-MIDAS specifications outperform their time-invariant versions in nowcasting the left tail of the GDP distribution. In fact, the best performing specifications in the model confidence sets for nowcasting left-tail risks are all TVP-MIDAS models. These results suggest that during periods of heightened volatility, the incorporation of time variation is crucial in predicting economic slowdown or recessionary events. This conclusion is in line with the findings of Adrian et al. (2019) and Estrella and Hardouvelis (1991), who underscore the significance of financial conditions and the yield curve slope as predictors for future recessions in the economy.

The rest of the paper is organized as follows. Section 2 introduces and discusses the proposed TVP-MIDAS framework. Section 3 outlines the posterior sampler. Section 4 assesses the accuracy of the proposal linear parameterizations in finite samples through a series of Monte Carlo experiments. Section 5 presents the real-time out-of-sample nowcast results. Finally, Section 6 concludes.

2. MIDAS regressions

To illustrate the MIDAS approach, we start with a simple setting in which we are interested in forecasting the variable y_t, which is observed only at discrete times $t = 1, 2, \ldots, T$, using the history of another variable $x_t^{(m)}$, which is observed m times between

the discrete time periods. More specifically, the observations of the high-frequency variable between $t-1$ and t are denoted as $x^{(m)}_{t-k/m}, k = 0, \dots, m-1$, where $x^{(m)}_{t-(m-1)/m}$ and $x^{(m)}_t$ are, respectively, the first and last available observations between the periods. An example is the forecasting of monthly inflation y_t using daily interest rates $x^{(m)}_t$ with $m = 22$, if we assume that there are 22 daily available observations within each month. In Section 2.3, we will consider more complex settings where the numbers of observations of the high-frequency variable between discrete time periods are not constant.

2.1. MIDAS weighting functions

One challenge even in this simple setting is the proliferation of parameters when m is large. A common approach is to use the average of the high-frequency variable observations between $t-1$ and t, $\frac{1}{m}\sum_{k=0}^{m-1} x^{(m)}_{t-k/m}$, as a single predictor. More specifically, let $h \geq 1$ denote the forecast horizon, and consider the following direct forecasting approach:

$$y_{t+h} = \alpha + \beta\left(\frac{1}{m}\sum_{k=0}^{m-1} x^{(m)}_{t-k/m}\right) + \epsilon_{t+h}. \tag{1}$$

Alternatively, one could use only the last observation of the high-frequency variable between periods $t-1$ and t:

$$y_{t+h} = \alpha + \beta x^{(m)}_t + \epsilon_{t+h}. \tag{2}$$

Obviously, both approaches are ad hoc and application specific. The key feature of the MIDAS regression is the use of a parsimonious and data-driven weighting function to summarize the information of the high-frequency variable $x^{(m)}_t$ for predicting y_t. As a simple example, consider the predictive regression

$$y_{t+h} = \alpha + \beta \mathbf{w}_t' \mathbf{x}^{(m)}_t + \epsilon_{t+h},$$

where $\mathbf{w}_t$ is an $m \times 1$ vector of weights and $\mathbf{x}^{(m)}_t = [x^{(m)}_t, x^{(m)}_{t-1/m}, \dots, x^{(m)}_{t-(m-1)/m}]'$. It is easy to verify that the predictive regressions in (1) and (2) are special cases with $\mathbf{w}_t = \frac{1}{m}\mathbf{1}_m$ and $\mathbf{w}_t = [1, 0, \dots, 0]'$, respectively.

Following Ghysels et al. (2007) and Pettenuzzo et al. (2016), we consider a general MIDAS regression of the form

$$y_{t+h} = \alpha + \boldsymbol{\rho}'\mathbf{y}_t + \boldsymbol{\gamma}'\mathbf{z}_t + \beta\mathcal{B}\left(L^{1/m};\boldsymbol{\theta}\right)x^{(m)}_t + \epsilon_{t+h}, \tag{3}$$

where the scalar β captures the overall impact of the lagged values of $x^{(m)}_t$ on y_{t+h}, $\boldsymbol{\rho}$ is the vector of autoregressive coefficients on $\mathbf{y}_t = [y_t, y_{t-1}, \dots, y_{t-p_y}]'$, and $\mathbf{z}_t$ is a vector of exogenous predictors. The MIDAS weighting function $\mathcal{B}\left(L^{1/m};\boldsymbol{\theta}\right)$ is parameterized as

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

where $L^{k/m}$ is a lag operator such that $L^{k/m}x^{(m)}_t = x^{(m)}_{t-k/m}$ and each component function $B(k;\boldsymbol{\theta})$ depends on a low-dimensional vector of parameters $\boldsymbol{\theta}$.

Ghysels et al. (2007) consider two parameterizations of the component function $B(k;\boldsymbol{\theta})$: the exponential Almon lag and the beta polynomial. Both parameterizations are parsimonious, and yet flexible enough to model a wide variety of dynamic patterns. However, they are nonlinear in the parameters, which makes estimation more difficult, especially in time-varying parameter settings. A further challenge is the imposition of the conventional identification restriction: in order to separately identify β and $\boldsymbol{\theta}$, one typically normalizes the weighting function $\mathcal{B}\left(L^{1/m};\boldsymbol{\theta}\right)$, i.e., replacing the component function $B(k;\boldsymbol{\theta})$ by its normalized version

$$\widetilde{B}(k;\boldsymbol{\theta}) = \frac{B(k;\boldsymbol{\theta})}{\sum_{k=1}^{K} B(k;\boldsymbol{\theta})}. \tag{4}$$

This type of normalization further complicates the estimation procedure.[1]

To tackle these challenges, we consider a class of parameterizations that are linear in the parameters for fast estimation. They may also be motivated as finite-dimensional approximations of weighting functions with desirable properties (e.g., smooth, bounded, square-integrable). In addition, we develop an alternative identification scheme that facilitates estimation. These two features are vitally important when we generalize the MIDAS model to time-varying parameter settings in the next section.

More specifically, suppose we wish to approximate a function $B(s)$ using the finite-dimensional approximation

$$B(s;\boldsymbol{\theta}) = \sum_{j=0}^{p} \theta_j \phi_j(s),$$

where $\phi_0, \dots, \phi_p$ are the basis functions and $\boldsymbol{\theta} = [\theta_0, \dots, \theta_p]'$ is the associated vector of coefficients. By evaluating $B(s;\boldsymbol{\theta})$ at discrete values $s = k = 0, \dots, K$, it takes the form

$$B(k;\boldsymbol{\theta}) = \boldsymbol{\theta}'\mathbf{v}_k, \tag{5}$$

[1] While the normalization and identification of β and $\boldsymbol{\theta}$ are not necessary for our forecasting application, they are useful for other applications that focus on the economic interpretation of the impact of the high-frequency variable on the low-frequency one. See Ghysels et al. (2007) for some interesting examples.

where $\mathbf{v}_k = [\phi_0(k), \dots, \phi_p(k)]'$. As an example, this formulation recovers the widely used Almon lag polynomial by setting $\mathbf{v}_k = [1, k, k^2, \dots, k^p]'$, so that

$$B(k;\boldsymbol{\theta}) = \sum_{j=0}^{p} \theta_j k^j. \tag{6}$$

That is, the Almon lag polynomial may be viewed as using the polynomials $\phi_j(s) = s^j, j = 0, 1, \dots, p$, as basis functions.

While polynomial basis functions are simple and easy to use, they are not orthogonal and do not provide an efficient basis system. An alternative is the set of Fourier basis functions — i.e., $\phi_0(s) = 1$, $\phi_j(s) = \cos(j\omega s)$ if j is odd and $\phi_j(s) = \sin(j\omega s)$ if j is even — that forms an orthonormal basis (for square-integrable functions).[2] By setting $\omega = 2\pi/(pm)$, $B(k;\boldsymbol{\theta})$ can be represented as

$$B(k;\boldsymbol{\theta}) = \theta_0 + \sum_{j=1}^{p} \left(\theta_{j1} \cos\left(\frac{2\pi}{pm} jk\right) + \theta_{j2} \sin\left(\frac{2\pi}{pm} jk\right) \right). \tag{7}$$

This formulation opens up many possibilities, as any basis functions, such as B-splines or wavelets, can be represented using the linear parameterization in (5). Not only is the linear parameterization flexible, it also makes estimation of the unknown parameter vector $\boldsymbol{\theta}$ straightforward.

Finally, instead of following the standard normalization approach that introduces additional nonlinearities in $\boldsymbol{\theta}$, we directly impose the linear equality constraint that the component functions sum to unity:

$$\sum_{k=1}^{K} B(k;\boldsymbol{\theta}) = \sum_{k=1}^{K} \boldsymbol{\theta}'\mathbf{v}_k = 1.$$

While this identification assumption is equivalent to the standard normalization approach given in (4), estimation following the former is much easier and it generalizes well to time-varying parameter settings, as we will show in the following section.

2.2. Time-varying coefficients, stochastic volatility and outlier adjustment

The conventional MIDAS regression in (3) assumes both a time-invariant weighting function $\mathcal{B}(L^{1/m};\boldsymbol{\theta})$ and a constant overall impact of the high-frequency variable $x_t^{(m)}$ on y_t. However, when forecasting macroeconomic variables, such as GDP or inflation, these assumptions are overly restrictive. In fact, an extensive literature has highlighted the significant benefits of accommodating parameter variations over time when forecasting such macroeconomic variables (see Barnett et al., 2014; Koop and Korobilis, 2013; D'Agostino et al., 2013).

Consequently, we develop a novel TVP-MIDAS framework, wherein both the weighting function and regression coefficients are permitted to evolve over time. This facilitates the direct assessment of the evolving impact of high-frequency variable $x_t^{(m)}$ on y_t. Schumacher (2014) proposes a MIDAS regression with time-varying exponential Almon lag weights. A limitation of this setup is that the exponential Almon lag polynomial is nonlinear in the parameters, and extending it to a time-varying setting involves the estimation of a nonlinear state space model. Schumacher (2014) considers an example with only one high-frequency predictor, and estimates the model using the particle filter. The estimation entails significant computational burden, rendering real-time forecasting using multiple high-frequency predictors infeasible.[3]

In contrast, the proposed framework uses linear parameterizations for the weighting functions and can be written as a linear Gaussian state space model. Therefore, estimation can be done easily using either conventional Kalman-filter based sampling methods or the more efficient precision-based methods developed in Chan and Jeliazkov (2009). The proposed approach thus scales well to high-dimensional settings and allows the researcher to consider multiple high-frequency predictors in real-time forecasting applications.

Another crucial aspect for modeling and forecasting macroeconomic time-series is the incorporation of stochastic volatility. A large body of empirical research, such as those conducted by Clark (2011), Clark and Ravazzolo (2015), Cross and Poon (2016) and Chan and Eisenstat (2018), has underscored the significance of accommodating time-varying volatility for both in-sample and out-of-sample applications. Furthermore, Carriero et al. (2015) and Pettenuzzo et al. (2016) have emphasized the importance of incorporating stochastic volatility in the context of MIDAS regressions for forecasting key macroeconomic variables. Finally, given the extreme movements in many macroeconomic variables during the COVID-19 pandemic, the proposed framework also explicitly includes an outlier component to address any potential outliers.

Specifically, we consider the following TVP-MIDAS model with stochastic volatility

$$y_{t+h} = \alpha_t + \boldsymbol{\rho}_t'\mathbf{y}_t + \boldsymbol{\gamma}_t'\mathbf{z}_t + \beta_t \mathcal{B}\left(L^{1/m};\boldsymbol{\theta}_t\right) x_t^{(m)} + \epsilon_{t+h}, \quad \epsilon_{t+h} \sim \mathcal{N}(0, \lambda_t e^{g_t}), \tag{8}$$

where the log-volatility g_t follows a standard random walk process

$$g_t = g_{t-1} + \eta_t, \quad \eta_t \sim \mathcal{N}(0, \sigma_g^2)$$

[2] Compared to Almon lag polynomial, Fourier basis functions are less frequently used in mixed-frequency settings. A notable exception is Bekierman and Gribisch (2021), who utilize a Fourier series expansion to capture the periodic intraday patterns in their mixed-frequency stochastic volatility model for intraday returns.

[3] In addition, recent research by Cross et al. (2023) has highlighted potential shortcomings of particle filtering methods, such as poor mixing properties and path degeneracy issues.

with the initial condition $g_1 \sim \mathcal{N}(0, V_g)$. The latent variable λ_t is introduced to model potential outliers. Different distributional assumptions on λ_t imply different types of outlier-augmented specifications. An example is the mixture distribution considered in Stock and Watson (2016) and Carriero et al. (2024). In particular, let $\lambda_t = o_t^2$, where o_t follows a 2-part distribution: with probability $1-q$, $o_t = 1$; otherwise, o_t follows a uniform distribution on the interval $(2, 10)$. The point mass at 1 represents regular observations whose scale is normalized to 1; the second part captures outliers that can have 2–10 times larger standard deviations relative to regular observations. Another example is to assume a continuous distribution for λ_t, say, an inverse-gamma distribution

$$(\lambda_t \mid \delta) \sim \mathcal{IG}(\delta/2, \delta/2).$$

This choice is motivated by the fact that a t distribution with degree of freedom δ can be represented as a scale mixture of normals in which the mixing distribution is $\mathcal{IG}(\delta/2, \delta/2)$. In the empirical application, we include this t specification for comparison, as it is found to work well in forecasting applications involving post COVID-19 pandemic data (see, e.g., Bobeica and Hartwig, 2023). We emphasize that the setup in (8) can accommodate many other types of outlier-augmented specifications.

In addition to the stochastic volatility and the outlier component, another important feature of the MIDAS model in (8) is that the weighting function is time-varying: $\mathcal{B}\left(L^{1/m}; \boldsymbol{\theta}_t\right) = \sum_{k=0}^{K} B(k; \boldsymbol{\theta}_t) L^{k/m}$, where the component function takes the form $B(k; \boldsymbol{\theta}_t) = \boldsymbol{\theta}_t' \mathbf{v}_k$ for some $(p+1)$-vector $\mathbf{v}_k$ (that depends of the chosen basis functions). Since

$$\mathcal{B}\left(L^{1/m}; \boldsymbol{\theta}_t\right) x_t^{(m)} = \sum_{k=0}^{K} \boldsymbol{\theta}_t' \mathbf{v}_k L^{k/m} x_t^{(m)} = \boldsymbol{\theta}_t' \sum_{k=0}^{K} \mathbf{v}_k x_{t-k/m}^{(m)} = \boldsymbol{\theta}_t' \mathbf{V} \mathbf{x}_t^{(m)},$$

where $\mathbf{V} = [\mathbf{v}_0, \mathbf{v}_1, \dots, \mathbf{v}_K]$ is a $(p+1) \times (K+1)$ matrix and $\mathbf{x}_t^{(m)} = [x_t^{(m)}, x_{t-1/m}^{(m)}, \dots, x_{t-K/m}^{(m)}]'$ is a $(K+1)$-vector, we can rewrite (8) as

$$y_{t+h} = \alpha_t + \boldsymbol{\rho}_t' \mathbf{y}_t + \boldsymbol{\gamma}_t' \mathbf{z}_t + \beta_t \boldsymbol{\theta}_t' \mathbf{V} \mathbf{x}_t^{(m)} + \epsilon_{t+h}, \quad \epsilon_{t+h} \sim \mathcal{N}(0, \lambda_t e^{g_t}). \tag{9}$$

Let $\mathbf{b}_t$ denote the $p_{\mathbf{b}}$-vector of time-varying parameters $\mathbf{b}_t = [\alpha_t, \boldsymbol{\rho}_t', \boldsymbol{\gamma}_t', \beta_t]'$. Then, we assume that the time-varying parameters $\mathbf{b}_t$ and $\boldsymbol{\theta}_t$ evolve according to the random walks:

$$\mathbf{b}_t = \mathbf{b}_{t-1} + \mathbf{u}_{1,t}, \quad \mathbf{u}_{1,t} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Omega}), \tag{10}$$

$$\boldsymbol{\theta}_t = \boldsymbol{\theta}_{t-1} + \mathbf{u}_{2,t}, \quad \mathbf{u}_{2,t} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Xi}), \tag{11}$$

where $\boldsymbol{\Omega} = \text{diag}(\omega_1^2, \dots, \omega_{p_{\mathbf{b}}}^2)$ and $\boldsymbol{\Xi} = \text{diag}(\xi_1^2, \dots, \xi_{p+1}^2)$, with the initial conditions $\mathbf{b}_1 \sim \mathcal{N}(\mathbf{0}, \mathbf{V}_{\mathbf{b}})$ and $\theta_1 \sim \mathcal{N}(\mathbf{0}, \mathbf{V}_{\theta})$. Similar to the time-invariant case, to separately identify β_t and $\boldsymbol{\theta}_t$, for $t = 1, \dots, T$, we impose the conditions

$$\boldsymbol{\theta}_t' \mathbf{V} \mathbf{1}_{K+1} = 1,$$

where $\mathbf{1}_{K+1}$ is a $(K+1)$-column of ones.

Finally, we assume the following priors on the time-invariant parameters

$$\begin{aligned}
&\omega_i^2 \sim \mathcal{IG}(\nu_\omega, S_\omega), \ i = 1, \dots, p_{\mathbf{b}}, \\
&\xi_i^2 \sim \mathcal{IG}(\nu_\xi, S_\xi), \ i = 1, \dots, p+1, \\
&\sigma_g^2 \sim \mathcal{IG}(\nu_g, S_g),
\end{aligned}$$

with hyperparameters $\nu_\omega = 5, S_\omega = 0.004$, $\nu_\xi = 10, S_\xi = 0.001$, $\nu_g = 5$ and $S_g = 0.04$. The hyperparameters for the initial conditions are set to be $\mathbf{V}_{\mathbf{b}} = 10\mathbf{I}_{p_{\mathbf{b}}}$, $\mathbf{V}_\theta = 10\mathbf{I}_{p+1}$ and $V_g = 10$.

2.3. Irregularly spaced mixed-frequency data

In many MIDAS applications, such as those by Marcellino and Schumacher (2010), Kuzin et al. (2011), Foroni and Marcellino (2014) and Mogliani and Simoni (2021), researchers use monthly predictors to forecast quarterly variables. Since every quarter has exactly 3 months, these are examples of regularly spaced mixed-frequency applications. However, for more complex applications, such as forecasting quarterly variables using weekly or daily predictors, we face two related but distinct challenges. Firstly, the numbers of observations of the high-frequency variables can vary across time periods (e.g., there are between 61 to 64 business days within a quarter). Secondly, the observations of the high-frequency variables might be irregularly spaced relative to the low frequency one (e.g., two weekly observations are available 3 and 10 days before the release of the monthly variable). These data issues become problematic when one attempts to align the low-frequency dependent variable with the high-frequency predictors. In our application, we nowcast quarterly GDP using both weekly and daily predictors. Consequently, we need to adapt the proposed framework to allow for time-varying numbers of high-frequency observations between discrete periods and irregularly spaced high-frequency observations.

To tackle the first challenge, let m_t denote the number of observations of the high-frequency variable $x_t^{(m)}$ between periods $t-1$ and t. Suppose for now that these observations are regularly spaced. That is, between the two periods, we observe $x_{t-k/m_t}^{(m)}, k = 0, \dots, m_t - 1$. The weighting function then becomes

$$\mathcal{B}\left(L^{1/m_t}; \boldsymbol{\theta}_t\right) x_t^{(m)} = \sum_{k=0}^{K} \boldsymbol{\theta}_t' \mathbf{v}_k L^{k/m_t} x_t^{(m)} = \boldsymbol{\theta}_t' \sum_{k=0}^{K} \mathbf{v}_k x_{t-k/m_t}^{(m)}, \tag{12}$$

where $\mathbf{v}_k = [\phi_0(k), \dots, \phi_p(k)]'$ is the vector of functional values of the basis functions $\phi_0(s), \dots, \phi_p(s)$ evaluated at $s = k$. Note that as long as we fix the number of basis functions that determines the dimension of $\mathbf{v}_k$, the number of coefficients that need to be estimated remains constant, even though the number of observations of $x_t^{(m)}$ may vary across t.

Now, suppose the number of observations between periods $t-1$ and t remains to be m_t, but these observations are irregularly spaced. Even so, we maintain the notation $x_{t-k/m_t}^{(m)}, k = 0, \dots, m_t - 1$ to denote the m_t observations, but they are available at times $0 \le s_{t,0} < s_{t,1} \cdots < s_{t,m_t-1} < 1$ from period t. That is, $x_{t-k/m_t}^{(m)}$ is available at time $t - s_{t,k}$. This formulation provides a very flexible framework to handle irregularly spaced observations. Naturally, we can recover the regularly spaced case by setting $s_{t,k} = k/m_t, k = 0, \dots, m_t - 1$. Finally, the weighting function has exactly the same form as in (12); one only needs to evaluate the basis functions at different points. Specifically, we replace $\mathbf{v}_k = [\phi_0(k), \dots, \phi_p(k)]'$ by $\mathbf{v}_{t,k} = [\phi_0(s_{t,k}), \dots, \phi_p(s_{t,k})]'$.

2.4. Data in multiple high frequencies

The proposed framework can be generalized to the case of multiple high-frequency variables with different numbers of observations between discrete periods. More specifically, suppose we have n high-frequency variables $x_t^{(m_1)}, \dots, x_t^{(m_n)}$, where $x_t^{(m_j)}$ is observed m_j times between time periods $t-1$ and t. Let $\mathcal{B}_j\left(L^{1/m_j}; \boldsymbol{\theta}_{j,t}\right)$ denote the weighting function for $x_t^{(m_j)}$, which takes the form

$$\mathcal{B}_j\left(L^{1/m_j}; \boldsymbol{\theta}_{j,t}\right) = \sum_{k=0}^{K_j} \boldsymbol{\theta}_{j,t}' \mathbf{v}_{j,k} L^{k/m_j},$$

where $\boldsymbol{\theta}_{j,t}$ is a (p_j+1)-vector of parameters and $\mathbf{v}_{j,k}$ is the corresponding vector of basis function values. If we define $\mathbf{V}_j = [\mathbf{v}_{j,0}, \mathbf{v}_{j,1}, \dots, \mathbf{v}_{j,K_j}]$ and $\mathbf{x}_t^{(m_j)} = [x_t^{(m_j)}, x_{t-1/m_j}^{(m_j)}, \dots, x_{t-K_j/m_j}^{(m_j)}]'$, the TVP-MIDAS model in (9) can be extended to include multiple high-frequency predictors:

$$y_{t+h} = \alpha_t + \boldsymbol{\rho}_t' \mathbf{y}_t + \boldsymbol{\gamma}_t' \mathbf{z}_t + \sum_{j=1}^{n} \beta_{j,t} \boldsymbol{\theta}_{j,t}' \mathbf{V}_j \mathbf{x}_t^{(m_j)} + \epsilon_{t+h}, \quad \epsilon_{t+h} \sim \mathcal{N}(0, \lambda_t e^{g_t}),$$

where $\beta_{j,t}$ captures the overall impact of $x_t^{(m_j)}$ on y_{t+h} at time t. This formulation again defines a linear Gaussian state space model in the time-varying parameters, and it can be efficiently estimated.

3. Posterior simulation

In this section, we outline the posterior sampler for estimating the proposed TVP-MIDAS model. In particular, we derive the conditional posterior distributions of the time-varying parameters $\mathbf{b} = (\mathbf{b}_1', \dots, \mathbf{b}_T')'$ and $\boldsymbol{\theta} = (\boldsymbol{\theta}_1', \dots, \boldsymbol{\theta}_T')'$ and discuss efficient sampling from these posterior distributions.

We start with the conditional posterior distribution of $\mathbf{b}$. To that end, stack $\mathbf{y} = (y_{1+h}, \dots, y_{T+h})'$ and $\boldsymbol{\epsilon} = (\epsilon_{1+h}, \dots, \epsilon_{T+h})'$, and rewrite (9) as

$$\mathbf{y} = \mathbf{X}_1 \mathbf{b} + \boldsymbol{\epsilon}, \quad \boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}), \tag{13}$$

where $\boldsymbol{\Sigma} = \text{diag}(\lambda_1 e^{g_1}, \dots, \lambda_T e^{g_T})$ and $\mathbf{X}_1 = \text{diag}(\mathbf{x}_{\mathbf{b},1}', \dots, \mathbf{x}_{\mathbf{b},T}')$ is a $T \times p_{\mathbf{b}}$ matrix whose tth row is $\mathbf{x}_{\mathbf{b},t} = [1, \mathbf{y}_t', \mathbf{z}_t', \boldsymbol{\theta}_t' \mathbf{V} \mathbf{x}_t^{(m)}]'$.

Next, stacking the state Eq. (10) over $t = 1, \dots, T$ yields

$$\mathbf{H}_1 \mathbf{b} = \mathbf{u}_1, \quad \mathbf{u}_1 \sim \mathcal{N}(\mathbf{0}, \mathbf{S}_1), \tag{14}$$

where $\mathbf{u}_1 = (\mathbf{u}_{1,1}', \dots, \mathbf{u}_{1,T}')'$, $\mathbf{S}_1 = \text{diag}(\mathbf{V}_{\mathbf{b}}, \boldsymbol{\Omega}, \dots, \boldsymbol{\Omega})$, and $\mathbf{H}_1$ is a first-difference matrix

$$\mathbf{H}_1 = \begin{bmatrix} \mathbb{I}_{p_{\mathbf{b}}} & \mathbf{O}_{p_{\mathbf{b}}} & \cdots & \cdots & \mathbf{O}_{p_{\mathbf{b}}} \\ -\mathbb{I}_{p_{\mathbf{b}}} & \mathbb{I}_{p_{\mathbf{b}}} & & & \vdots \\ \mathbf{O}_{p_{\mathbf{b}}} & \ddots & \ddots & & \vdots \\ & & \ddots & \mathbb{I}_{p_{\mathbf{b}}} & \mathbf{O}_{p_{\mathbf{b}}} \\ \mathbf{O}_{p_{\mathbf{b}}} & \cdots & \cdots & -\mathbb{I}_{p_{\mathbf{b}}} & \mathbb{I}_{p_{\mathbf{b}}} \end{bmatrix}.$$

Since the determinant of $\mathbf{H}_1$ is one, it is invertible. By a change of variable, we have $\mathbf{b} \sim \mathcal{N}(\mathbf{0}, (\mathbf{H}_1' \mathbf{S}_1^{-1} \mathbf{H}_1)^{-1})$. Combining (13) and (14) and using standard linear regression results, the conditional posterior for $\mathbf{b}$ is then obtained as

$$(\mathbf{b} \mid \mathbf{y}, \boldsymbol{\theta}, \boldsymbol{\Sigma}, \boldsymbol{\Omega}) \sim \mathcal{N}(\boldsymbol{\mu}_{\mathbf{b}}, \mathbf{K}_{\mathbf{b}}^{-1}),$$

where

$$\mathbf{K}_{\mathbf{b}} = \mathbf{H}_1' \mathbf{S}_1^{-1} \mathbf{H}_1 + \mathbf{X}_1' \boldsymbol{\Sigma}^{-1} \mathbf{X}_1, \quad \boldsymbol{\mu}_{\mathbf{b}} = \mathbf{K}_{\mathbf{b}}^{-1} (\mathbf{X}_1' \boldsymbol{\Sigma}^{-1} \mathbf{y}).$$

Since the precision matrix $\mathbf{K}_{\mathbf{b}}$ is a band matrix, sampling from $(\mathbf{b} \mid \mathbf{y}, \boldsymbol{\theta}, \boldsymbol{\Sigma}, \boldsymbol{\Omega})$ can be efficiently accomplished using the algorithm in Chan and Jeliazkov (2009).

The conditional posterior distribution of $\boldsymbol{\theta}$ can be derived similarly. More specifically, let $\tilde{y}_t = y_{t+h} - \alpha_t - \boldsymbol{\rho}_t' \mathbf{y}_t - \boldsymbol{\gamma}_t' \mathbf{z}_t$ and stack $\tilde{y}_t$ over $t = 1, \dots, T$ to obtain $\tilde{\mathbf{y}} = (\tilde{y}_t, \dots, \tilde{y}_T)'$. Then, (9) can be rewritten as

$$\tilde{\mathbf{y}} = \mathbf{X}_2 \boldsymbol{\theta} + \boldsymbol{\epsilon}, \quad \boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}), \tag{15}$$

where $\mathbf{X}_2 = \text{diag}(\beta_1 \mathbf{x}_1^{(m)\prime} \mathbf{V}', \dots, \beta_T \mathbf{x}_T^{(m)\prime} \mathbf{V}')$. Furthermore, stacking the state Eq. (11) over $t = 1, \dots, T$, we have

$$\mathbf{H}_2 \boldsymbol{\theta} = \mathbf{u}_2, \quad \mathbf{u}_2 \sim \mathcal{N}(\mathbf{0}, \mathbf{S}_2), \tag{16}$$

where $\mathbf{u}_2 = (\mathbf{u}_{2,1}', \dots, \mathbf{u}_{2,T}')'$ and $\mathbf{S}_2 = \text{diag}(\mathbf{V}_\theta, \boldsymbol{\Xi}, \dots, \boldsymbol{\Xi})$ and

$$\mathbf{H}_2 = \begin{bmatrix} \mathbb{I}_{p+1} & \mathbf{O}_{p+1} & \cdots & \cdots & \mathbf{O}_{p+1} \\ -\mathbb{I}_{p+1} & \mathbb{I}_{p+1} & & & \vdots \\ \mathbf{O}_{p+1} & \ddots & \ddots & & \vdots \\ & & \ddots & \mathbb{I}_{p+1} & \mathbf{O}_{p+1} \\ \mathbf{O}_{p+1} & \cdots & \cdots & -\mathbb{I}_{p+1} & \mathbb{I}_{p+1} \end{bmatrix}.$$

Here $\mathbf{H}_2$ is a first-difference matrix with unit determinant. It follows that $\boldsymbol{\theta} \sim \mathcal{N}(\mathbf{0}, (\mathbf{H}_2' \mathbf{S}_2^{-1} \mathbf{H}_2)^{-1})$. Without imposing any restrictions on $\boldsymbol{\theta}$, its conditional posterior distribution is again Gaussian. A slight complication is the imposition of the identification restrictions $\boldsymbol{\theta}_t' \mathbf{V} \mathbf{1}_{K+1} = 1$ for $t = 1, \dots, T$. Specifically, let $\mathcal{S}$ denote the hyperplane defined by the T linear equality restrictions

$$\mathcal{S} \stackrel{\text{def}}{=} \left\{ \boldsymbol{\theta} \in \mathbb{R}^{T(p+1)} : (\mathbb{I}_T \otimes (\mathbf{1}_{K+1}' \mathbf{V}')) \boldsymbol{\theta} = \mathbf{1}_T \right\}.$$

Then, the conditional posterior of $\boldsymbol{\theta}$ is a Gaussian distribution truncated to the hyperplane $\mathcal{S}$:

$$(\boldsymbol{\theta} \mid \mathbf{y}, \mathbf{b}, \boldsymbol{\Sigma}, \boldsymbol{\Xi}) \sim \mathcal{N}_{\mathcal{S}}(\boldsymbol{\mu}_\theta, \mathbf{K}_\theta^{-1}),$$

where

$$\mathbf{K}_\theta = \mathbf{H}_2' \mathbf{S}_2^{-1} \mathbf{H}_2 + \mathbf{X}_2' \boldsymbol{\Sigma}^{-1} \mathbf{X}_2, \quad \boldsymbol{\mu}_\theta = \mathbf{K}_\theta^{-1} (\mathbf{X}_2' \boldsymbol{\Sigma}^{-1} \tilde{\mathbf{y}}).$$

There are efficient algorithms that can be used to sample from $\mathcal{N}_{\mathcal{S}}(\boldsymbol{\mu}_\theta, \mathbf{K}_\theta^{-1})$, such as Algorithm 2.6 in Rue and Held (2005) and Algorithm 2 in Cong et al. (2017). In particular, we can first sample $\tilde{\boldsymbol{\theta}} \sim \mathcal{N}(\boldsymbol{\mu}_\theta, \mathbf{K}_\theta^{-1})$ using the algorithm in Chan and Jeliazkov (2009). Then, we impose the identification restrictions $\mathbf{M}\boldsymbol{\theta} = \mathbf{1}_T$, where $\mathbf{M} = \mathbb{I}_T \otimes (\mathbf{1}_{K+1}' \mathbf{V}')$ by computing

$$\boldsymbol{\theta} = \tilde{\boldsymbol{\theta}} + \mathbf{K}_\theta^{-1} \mathbf{M}' (\mathbf{M} \mathbf{K}_\theta^{-1} \mathbf{M}')^{-1} (\mathbf{1}_T - \mathbf{M}\tilde{\boldsymbol{\theta}}).$$

Other steps of the posterior sampler are standard. For example, the log-volatility can be sampled using the auxiliary mixture sampler of Kim et al. (1998), with the adjustment (for the latent variables $\lambda_1, \dots, \lambda_T$) outlined in Chan and Hsiao (2014). The degree of freedom parameter δ can be sampled using a Metropolis–Hastings step described in Chan and Hsiao (2014).

4. Assessing the accuracy of the linear parameterizations

The proposed linear parameterization may be motivated as a finite-dimensional approximation of certain classes of weighting functions. As a specific example, consider the Hilbert space of L^2 or square-integrable functions on $(0, \infty)$, i.e., $B \in L^2$ if $\|B\|_2 = \left(\int_0^\infty |B(x)|^2 \text{d}x \right)^{1/2} < \infty$, and fix a basis $\{\phi_j\}_{j=0}^\infty$. This family of functions is flexible and contains many commonly-used weighting functions, such as those that are bounded with finite support. For a weighting function $B \in L^2$, the proposed linear parameterization with basis functions $\phi_0, \dots, \phi_p$, $B(s; \boldsymbol{\theta}) = \sum_{j=0}^p \theta_j \phi_j(s)$, can therefore be viewed as a finite-dimensional approximation of $B(s)$, where the approximation error vanishes in $\|\cdot\|_2$ norm as $p \to \infty$.

Of course, for any finite p, $B(s; \boldsymbol{\theta}) = \sum_{j=0}^p \theta_j \phi_j(s)$ is an approximation. In addition to the approximation error, in practice one needs to estimate the coefficients $\theta_0, \dots, \theta_p$, which entails estimation errors. Below we assess the accuracy of the proposal linear parameterizations in finite samples through a series of Monte Carlo experiments.

4.1. The time-invariant case

We first consider MIDAS models where the coefficients of the weighting functions are constant. Specifically, we generate data from an autoregressive distributed lag MIDAS model with a nonlinear, constant-coefficient weighting function. Then, we estimate a MIDAS model with the proposed linear parameterizations using the simulated data. To assess its accuracy, we compare its forecast performance relative to the parametric MIDAS model from which the data are generated.

We follow the simulation design in Babii et al. (2022), particularly the data generating process (DGP):

$$y_t = \rho_1 y_{t-1} + \rho_2 y_{t-2} + \sum_{j=1}^n \frac{1}{m} \sum_{i=1}^m \widetilde{B}\left(\frac{i-1}{m}; \boldsymbol{\theta}_j \right) x_{t-(i-1)/m, j} + u_t, \quad u_t \sim \mathcal{N}(0, \sigma_u^2),$$

where y_t is the low-frequency variable of interest and $x_{t,j}, j = 1, \dots, n$, are the predictors with the same high-frequency. Following Babii et al. (2022), we set $\sigma_u^2 = 1$, $\rho_1 = 0.3$, $\rho_2 = 0.01$, $n = 3$ and $m = 12$. For the weighting functions $\widetilde{B}\left(s; \boldsymbol{\theta}_j\right)$, we consider two types. The first is based on the beta density $f(x; a, b) = \frac{\Gamma(a+b)}{\Gamma(a)\Gamma(b)} x^{a-1}(1-x)^{b-1}$:

$$B\left(s; \boldsymbol{\theta}_1\right) = f(s; 1, 3), \quad B\left(s; \boldsymbol{\theta}_2\right) = f(s; 2, 3), \quad B\left(s; \boldsymbol{\theta}_3\right) = f(s; 2, 2),$$

Table 1
Average MSFE and CRPS across $R = 20$ replications. The numerical standard errors are in the parentheses.

	$T = 50$			$T = 100$			$T = 200$			$T = 500$		
	Beta	Fourier	Almon	Beta	Fourier	Almon	Beta	Fourier	Almon	Beta	Fourier	Almon
MSFE	1.17	1.32	1.26	1.04	1.14	1.10	1.11	1.14	1.15	1.04	1.05	1.05
	(0.10)	(0.09)	(0.09)	(0.05)	(0.07)	(0.07)	(0.06)	(0.07)	(0.07)	(0.03)	(0.03)	(0.03)
CRPS	0.61	0.66	0.64	0.58	0.60	0.59	0.59	0.60	0.60	0.58	0.58	0.58
	(0.03)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.01)	(0.01)	(0.01)
	Exp. Almon	Fourier	Almon	Exp. Almon	Fourier	Almon	Exp. Almon	Fourier	Almon	Exp. Almon	Fourier	Almon
MSFE	1.08	1.25	1.19	0.93	1.00	0.98	1.08	1.10	1.10	1.04	1.05	1.05
	(0.08)	(0.10)	(0.09)	(0.07)	(0.07)	(0.07)	(0.04)	(0.04)	(0.04)	(0.03)	(0.03)	(0.03)
CRPS	0.58	0.63	0.62	0.54	0.56	0.55	0.59	0.59	0.59	0.58	0.58	0.58
	(0.02)	(0.03)	(0.03)	(0.02)	(0.02)	(0.02)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)

and we normalize the weighting functions via $\widetilde{B}(s;\boldsymbol{\theta}_j) = B(s;\boldsymbol{\theta}_j)/\sum_{i=1}^{m} B((i-1)/m;\boldsymbol{\theta}_j)$. The second is based on the exponential Almon polynomials:

$$B\left(s;\boldsymbol{\theta}_1\right) = \exp\left(7s \times 10^{-4} - s^2 \times 10^{-4}\right),$$
$$B\left(s;\boldsymbol{\theta}_2\right) = \exp\left(6s \times 10^{-3} - 5s^2 \times 10^{-4}\right),$$
$$B\left(s;\boldsymbol{\theta}_3\right) = \exp\left(3s \times 10^{-2} - 7s^2 \times 10^{-4}\right),$$

which are similarly normalized to obtain $\widetilde{B}(s;\boldsymbol{\theta}_j), j = 1, 2, 3$. We generate the high-frequency predictors $x_{t,j}$ according to the AR(1) process:

$$x_h = 0.7x_{h-1} + \epsilon_h, \quad \epsilon_h \sim \mathcal{N}(0, 1). \tag{17}$$

Finally, the initial conditions are set as $x_0 \sim \mathcal{N}(0, 1/(1-0.7^2))$ and $y_0 \sim \mathcal{N}(0, \sigma_u^2(1-\rho_2)/((1+\rho_2)((1-\rho_2)^2 - \rho_1^2))$.

We estimate the proposed MIDAS models (restricting the coefficients to be constant in this case) with two types of basis functions: the Fourier series and the Almon lag polynomials. The accuracy of the out-of-sample forecasts is assessed against the true parametric MIDAS model from which the data are generated (here the functional form of the weighting function is assumed to be known but the coefficients are estimated from the data). We consider a range of sample sizes from $T = 50$ to $T = 500$. In each case, the final 25% of the sample is designated as the evaluation period, and we recursively compute the one-step-ahead forecasts from the models using an expanding window. We repeat this forecasting process $R = 20$ times, i.e., 20 time-series are generated for each simulation design. For each dataset, the models are estimated using Markov chain Monte Carlo methods with 10,000 posterior draws after a burn-in period of 5000.

Table 1 presents the average mean squared forecast error (MSFE) and average continuous ranked probability score (CRPS) across the replications for the true parametric MIDAS models (beta or exponential Almon) and the proposed MIDAS models with two types of basis functions (Fourier and Almon). As expected, in all cases the true parametric models forecast the best, due to the efficiency gain in assuming that the DGP is known. Interestingly, the forecast performance of the proposed MIDAS models with Fourier series and Almon lag polynomial basis functions is similar to the true model, especially when the sample size is large. These results show that the proposed linear parameterizations can effectively approximate nonlinear MIDAS weight functions.

In terms of computational costs, the proposed approach based on the linear parameterizations is much faster than fitting the parametric MIDAS models. This is because the weighting functions in the former are linear in the parameters whereas they are highly nonlinear in the latter. Consequently, the former can be estimated using a standard Gibbs sampler, whereas the latter requires the Metropolis–Hastings algorithm. For example, for the parametric MIDAS model based on the beta density with sample size $T = 500$, obtaining 10,000 posterior draws using a tailored Metropolis–Hastings algorithm takes about 4 min on a standard desktop with an Intel Xeon W-2223 @ 3.6 GHz processor and 16 GB of RAM; for the proposed approach, obtaining 10,000 posterior draws takes about 10 s instead.

4.2. The time-varying case

We now examine whether the proposed TVP-MIDAS models can accurately recover the coefficients on the weighting functions when these nonlinear weighting functions are approximated using the linear parameterizations. Specifically, we consider the following DGP with nonlinear weighting functions:

$$y_t = \beta_{1,t} + \beta_{2,t}\widetilde{B}_1(k;\boldsymbol{\theta}_{1,t})x_{1,t}^{(m_1)} + \beta_{3,t}\widetilde{B}_2(k;\boldsymbol{\theta}_{2,t})x_{2,t}^{(m_2)} + u_t, \quad u_t \sim \mathcal{N}(0, \sigma_u^2),$$

where y_t is the low-frequency variable (e.g., quarterly), the high-frequency predictors $x_{1,t}^{(m_1)}$ and $x_{2,t}^{(m_2)}$ are generated according to the AR process in (17), and the time-varying coefficients follow the independent random walks

$$\beta_{i,t} = \beta_{i,t-1} + u_{i,t}^{\beta}, \quad u_{i,t}^{\beta} \sim \mathcal{N}(0, \sigma_\beta^2)$$

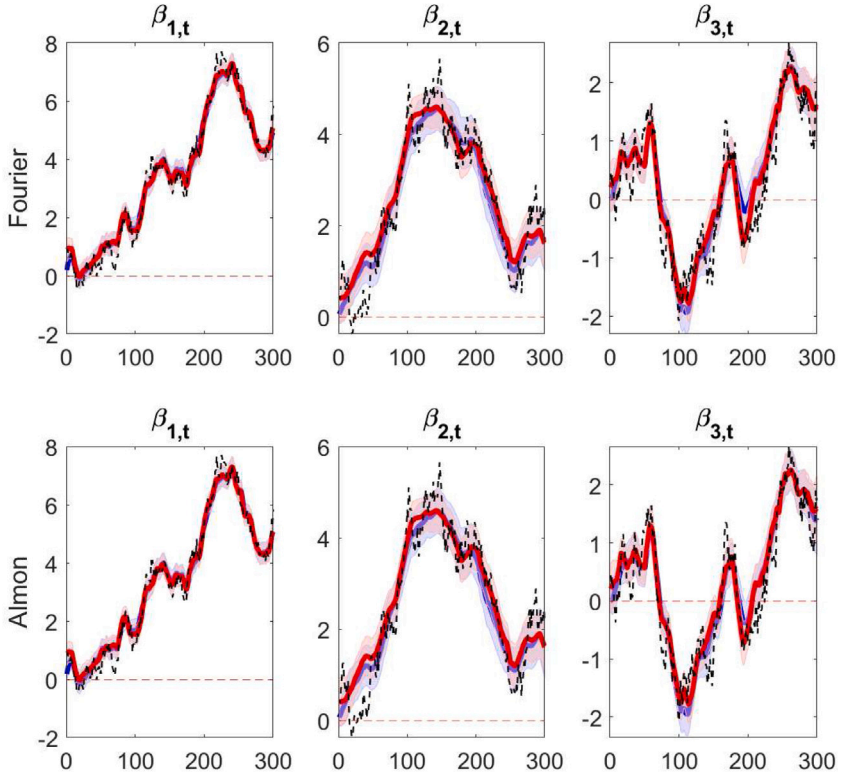

Fig. 1. Posterior estimates of $\beta_{i,t}, i = 1, 2, 3$, from the proposed TVP-MIDAS models with the Fourier series and Almon lag polynomial basis functions (solid blue line) against the true values (dashed black line) and estimates from the TVP-MIDAS model with known beta weighting functions (solid red line). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

for $i = 1, 2, 3$. We set $\sigma_u^2 = 1$, $\sigma_\beta^2 = 0.05$, $T = 300$, $m_1 = 60$ (60 days per quarter) and $m_2 = 12$ (12 weeks per quarter). Similar to the constant-coefficient case, we consider two types of weighting functions: the first type is based on the beta density $f(x; a, b)$

$$\widetilde{B}_1(k; \boldsymbol{\theta}_{1,t}) = \frac{f(\frac{k}{m_1}, \theta_{1,t}, \theta_{2,t})}{\sum_{i=1}^{m_1} f(\frac{i}{m_1}, \theta_{1,t}, \theta_{2,t})}, \quad \widetilde{B}_2(k; \boldsymbol{\theta}_{2,t}) = \frac{f(\frac{k}{m_2}, \theta_{3,t}, \theta_{4,t})}{\sum_{i=1}^{m_2} f(\frac{i}{m_2}, \theta_{3,t}, \theta_{4,t})},$$

and the second on the exponential Almon polynomials

$$\widetilde{B}_1(k; \boldsymbol{\theta}_{1,t}) = \frac{\exp\left(\theta_{1,t}\frac{k}{m_1} + \theta_{2,t}(\frac{k}{m_1})^2\right)}{\sum_{i=1}^{m_1} \exp\left(\theta_{1,t}\frac{i}{m_1} + \theta_{2,t}(\frac{i}{m_1})^2\right)}, \quad \widetilde{B}_2(k; \boldsymbol{\theta}_{2,t}) = \frac{\exp\left(\theta_{3,t}\frac{k}{m_2} + \theta_{4,t}(\frac{k}{m_2})^2\right)}{\sum_{i=1}^{m_2} \exp\left(\theta_{3,t}\frac{i}{m_2} + \theta_{4,t}(\frac{i}{m_2})^2\right)}.$$

The parameters in the weighting function are assumed to follow the following autoregressive process:

$$\theta_{i,t} = \theta_{i,t-1} + u_{i,t}^\theta, \quad u_{i,t}^\theta \sim \mathcal{TN}_{(c_1,c_2)}(0, \sigma_\theta^2)$$

for $i = 1, \ldots, 4$, where $\mathcal{TN}_{(c_1,c_2)}(0, \sigma_\theta^2)$ denotes the normal distribution with mean 0 and variance σ_θ^2 truncated to the interval (c_1, c_2). We set $\sigma_\theta^2 = 0.001$. For the first type based on the beta density, we specify $c_1 = 1$ and $c_2 = 20$; for the exponential Almon lag polynomials, we set $c_1 = 0$ and $c_2 = 0.001$.

Fig. 1 presents the posterior estimates of $\beta_{1,t}, \beta_{2,t}$ and $\beta_{3,t}$ from the proposed TVP-MIDAS models using two types of basis functions (Fourier series and Almon lag polynomials), where the DGP is based on the time-varying beta weighting functions specified above. To provide a benchmark of these estimates, we also fit a TVP-MIDAS model in which the beta weighting functions are assumed to be known and only the coefficients $\beta_{1,t}, \beta_{2,t}$ and $\beta_{3,t}$ are estimated. All the posterior estimates are based on 10,000 MCMC draws after a burn-in period of 5000. It is clear from the figure that the posterior estimates from the proposed TVP-MIDAS models closely track the true values and are very similar to those from the infeasible TVP-MIDAS model with known weighting functions.

To further evaluate the in-sample fit of the proposed models, we compute the Bayesian R^2 measure proposed by Gelman et al. (2019). Similar to the conventional R^2, this measure is always between 0 and 1 by construction. The posterior distributions of this measure from the proposed TVP-MIDAS models are reported in Fig. 2. For comparison we also report the Bayesian R^2 from the TVP-MIDAS model with known weighting functions. The results show the two TVP-MIDAS models using the linear

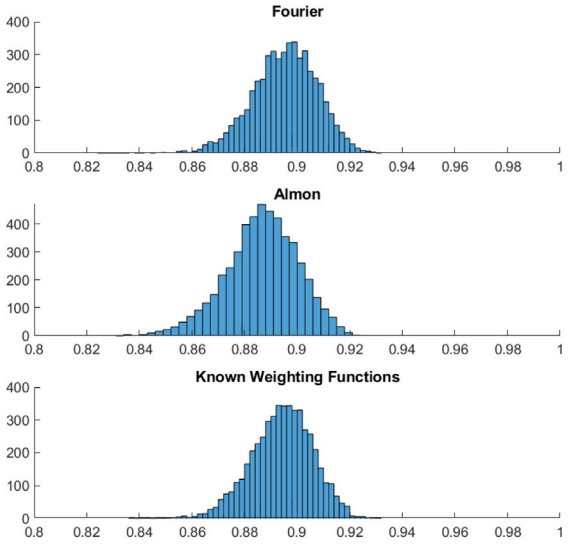

Fig. 2. Posterior distributions of the Bayesian R^2 from the proposed TVP-MIDAS models with Fourier series and Almon lag polynomial basis functions and TVP-MIDAS model with known beta weighting functions.

parameterizations achieve very similar Bayesian R^2 values compared to the oracle, suggesting that the loss in in-sample fit using the linear parameterizations is not substantial.

A similar set of results are observed for the DGP using weighting functions based on the exponential Almon lag polynomial; see Online Appendix B for details. All in all, these findings suggest that the proposed TVP-MIDAS models using the linear parameterizations can effectively approximate commonly-used non-linear weighting functions in a dynamic TVP-MIDAS setting.

5. Empirical application: Nowcasting US GDP

5.1. Design of the real-time nowcasting application

We assess the performance of the proposed TVP-MIDAS framework via a real-time nowcasting application: we use monthly, weekly and daily variables to nowcast the quarterly US real GDP (see Cascaldi-Garcia et al., 2024, for a recent survey on nowcasting GDP). In particular, we consider two sets of predictors: a small-scale case and a large-scale case. For the small-scale case, the predictors include the monthly industrial production, the weekly National Financial Conditions Index (NFCI), and a daily interest rate spread (defined as the difference between the 10-year and 3-month treasury yields) that captures the slope of the yield curve. The industrial production is a standard monthly measure of real economic activity.[4] The NFCI is widely used in nowcasting GDP, following the influential work of Adrian et al. (2019), which shows that tightening financial conditions are associated with a notable increase in downside risk for US real GDP. Finally, the yield curve slope has been consistently shown to improve forecasts of US real GDP (Estrella and Hardouvelis, 1991; Estrella et al., 2003; Rudebusch and Williams, 2009). Moreover, a recent study by Poon and Zhu (2024) underscores the importance of financial conditions as crucial predictors for forecasting recessions across various countries. For the large-scale case, we include all the predictors in the small-scale case, as well as 12 additional monthly predictors that are often used in nowcasting GDP. More details on these predictors are provided in Online Appendix C.

We construct our real-time datasets from a few data sources. The quarterly vintages of US real GDP and monthly vintages of industrial production are sourced from the Philadelphia Federal Reserve Real-Time Datasets for Macroeconomists, spanning from 1990Q1 to 2021Q2. The weekly NFCI vintages, starting from 2011, are obtained from the Archival Federal Reserve Economic Data (ALFRED) database. For the earlier sample, we utilize the weekly NFCI dataset compiled by Amburgey and McCracken (2023), which includes weekly data vintages from 1988 onwards. Finally, we acquire the daily interest rate spread data from the St. Louis FRED database. We transform the GDP and industrial production to annualized growth rates—i.e., we multiply the quarterly and monthly changes in the natural logarithms of GDP and industrial production by factors of 400 and 100, respectively. The additional 12 monthly predictors are also sourced from the ALFRED database and are transformed similarly; details of the transformation are provided in Online Appendix C.

[4] In an earlier version of this paper, we also implemented models with the monthly employment growth instead of the industrial production. On average, we find that models utilizing the industrial production provide slightly better nowcasts than those using the employment growth. These results are available upon request.

Table 2
Nowcasting scheme for 2000Q1.

	Forecast origin		
	January 2000	February 2000	March 2000
Forecast horizon	$h = 2/3$	$h = 1/3$	$h = 0$
GDP data up to quarter	1999Q4	1999Q4	1999Q4
Industrial production data up to month	December 1999	January 2000	February 2000
NFCI data up to month	January 2000	February 2000	March 2000
Interest rate spread data up to month	January 2000	February 2000	March 2000

Table 3
Competing MIDAS models nested within the proposed framework.

Model	Description
MIDAS	MIDAS with constant parameters and constant volatility
MIDAS-SV	MIDAS with constant parameters and stochastic volatility (SV)
MIDAS-SVt	MIDAS with constant parameters, SV and t errors
TVP-MIDAS	MIDAS with time-varying parameters (TVP) and constant volatility
TVP-MIDAS-SV	MIDAS with TVP and SV
TVP-MIDAS-SVt	MIDAS with TVP, SV and t errors
TVP-MIDAS-R1	MIDAS with time-varying $\boldsymbol{\beta}_t$, $\boldsymbol{\theta}_t = \boldsymbol{\theta}$ and constant volatility
TVP-MIDAS-SV-R1	MIDAS with time-varying $\boldsymbol{\beta}_t$, $\boldsymbol{\theta}_t = \boldsymbol{\theta}$ and SV
TVP-MIDAS-SVt-R1	MIDAS with time-varying $\boldsymbol{\beta}_t$, $\boldsymbol{\theta}_t = \boldsymbol{\theta}$, SV and t errors
TVP-MIDAS-R2	MIDAS with time-varying $\boldsymbol{\theta}_t$, $\boldsymbol{\beta}_t = \boldsymbol{\beta}$ and constant volatility
TVP-MIDAS-SV-R2	MIDAS with time-varying $\boldsymbol{\theta}_t$, $\boldsymbol{\beta}_t = \boldsymbol{\beta}$ and SV
TVP-MIDAS-SVt-R2	MIDAS with time-varying $\boldsymbol{\theta}_t$, $\boldsymbol{\beta}_t = \boldsymbol{\beta}$, SV and t errors

Our nowcasting design aligns with that of Guérin and Marcellino (2013), who nowcast quarterly US real GDP using a Markov-switching MIDAS framework. Our approach involves generating nowcasts of US real GDP at the conclusion of each month of the quarter. Table 2 gives an example of the data vintages used to construct nowcasts of US real GDP in 2000Q1. US GDP data are released about one month after the end of the quarter. Consequently, when constructing a nowcast of 2000Q1 GDP at the conclusion of January 2000, our information set encompasses the GDP data up to 1999Q4 and the daily, weekly and monthly predictors released up to the end of January 2000. Progressing to the end of February 2000, our information set expands to include information on daily, weekly and monthly predictors up to the conclusion of the second month of the quarter. By the conclusion of March 2000, our information set encompasses daily, weekly and monthly predictors for the entire quarter. Formally, we denote the nowcasts at the conclusion of the first, second, and third months of the quarter as $h = 2/3$, $h = 1/3$ and $h = 0$, respectively.

Our initial sample spans from 1982Q1 to 1989Q4, with recursive expansion continuing until the end of sample. This temporal progression is mirrored in the timeframe for our daily, weekly, and monthly predictors. We focus on the evaluation period that starts from 1990Q1 and ends in 2019Q4, given the extreme, unexpected movements of US GDP at the onset of the COVID-19 pandemic. But we also assess the performance of the TVP-MIDAS models using a sample that ends in 2021Q1.

5.2. Out-of-sample nowcast performance

We evaluate the nowcast performance of the proposed TVP-MIDAS framework against a wide variety of MIDAS specifications with two goals in mind. First, since the proposed TVP-MIDAS framework can accommodate many different types of nonlinearities and time-variations, it is useful to see what type of model flexibility matters most in nowcasting US GDP. To that end, we evaluate a range of MIDAS configurations nested within the proposed framework by switching on and off different features. In particular, we include MIDAS specifications in which we fix either the coefficients on the high-frequency variables or the parameters in the weighting functions to be constant, i.e., $\boldsymbol{\beta}_t = \boldsymbol{\beta}$ and $\boldsymbol{\theta}_t = \boldsymbol{\theta}$ for $t = 1, \dots, T$, respectively, to assess which type of time-variation is the most useful. We also consider different combinations of volatility assumptions (constant variance or stochastic volatility) and error distributions (normal or t distributions). All in all, we consider 11 variants that are nested within the proposed framework, as well as an AR(2) model as a benchmark. These competing models are summarized in Table 3. This evaluation is done using the first set of monthly, weekly and daily predictors to nowcast GDP.

The second goal of this subsection is to demonstrate that the proposed TVP-MIDAS models are competitive against state-of-the-art machine learning-based MIDAS specifications in a data-rich environment. To that end, we nowcast US GDP using a larger dataset that includes the monthly, weekly and daily variables in the small-scale case, as well as 12 additional monthly variables. Naturally, incorporating this extensive set of variables introduces the risk of overfitting given the very flexible TVP-MIDAS framework. To address this challenge, we adopt the noncentered parameterization developed in Frühwirth-Schnatter and Wagner (2010) and Bitto and Frühwirth-Schnatter (2019), which facilitates the implementation of global-local shrinkage priors. Following the methodology in Huber et al. (2021), we employ these priors to induce shrinkage on the time-varying parameters. Specifically, we consider three widely-used global-local shrinkage priors from the macroeconometrics literature: the Dirichlet–Laplace, normal-gamma, and horseshoe priors (see Online Appendix D for technical details).

Table 4
Competing large-scale MIDAS models.

Model	Description
BMIDAS-AGL	Bayesian MIDAS with adaptive group Lasso in Mogliani and Simoni (2021)
BMIDAS-AGL-SS	Bayesian MIDAS with adaptive group Lasso and spike and slab prior in Mogliani and Simoni (2021)
MIDAS-SG-LASSO	Sparse group Lasso MIDAS in Babii et al. (2022)
U-MIDAS-SG-LASSO	Sparse group Lasso unrestricted MIDAS in Babii et al. (2022)
U-MIDAS-AGL	Unrestricted MIDAS with adaptive group Lasso in Mogliani and Simoni (2021)
U-MIDAS-AGL-SS	Unrestricted MIDAS with adaptive group Lasso and spike and slab prior in Mogliani and Simoni (2021)
TVP-MIDAS-DL	TVP-MIDAS with Dirichlet–Laplace prior and constant volatility
TVP-MIDAS-SV-DL	TVP-MIDAS with Dirichlet–Laplace prior and SV
TVP-MIDAS-SVt-DL	TVP-MIDAS with Dirichlet–Laplace prior, SV and t errors
TVP-MIDAS-NG	TVP-MIDAS with normal-gamma prior and constant volatility
TVP-MIDAS-SV-NG	TVP-MIDAS with normal-gamma prior and SV
TVP-MIDAS-SVt-NG	TVP-MIDAS with normal-gamma prior, SV and t errors
TVP-MIDAS-HS	TVP-MIDAS with horseshoe prior and constant volatility
TVP-MIDAS-SV-HS	TVP-MIDAS with horseshoe prior and SV
TVP-MIDAS-SVt-HS	TVP-MIDAS with horseshoe prior, stochastic volatility and t errors

Table 5a
The RMSFEs of the variants nested within the proposed TVP-MIDAS framework benchmarked against the AR(2).

Forecast horizon	$h = 2/3$		$h = 1/3$		$h = 0$	
Model	RMSFE	p_{MCS}	RMSFE	p_{MCS}	RMSFE	p_{MCS}
Fourier						
MIDAS	0.95	0.00	0.95	0.03	0.96	0.05
MIDAS-SV	0.93	0.11^{+}	0.94	0.03	0.94	0.05
MIDAS-SVt	0.93	0.11^{+}	0.94	0.03	0.94	0.05
TVP-MIDAS	0.93^{*}	0.11^{+}	0.93^{*}	0.11^{+}	0.92^{*}	0.05
TVP-MIDAS-SV	0.93^{*}	0.11^{+}	0.92^{*}	0.11^{+}	0.92^{*}	0.05
TVP-MIDAS-SVt	0.93^{*}	0.11^{+}	0.92^{*}	0.11^{+}	0.92^{*}	0.05
TVP-MIDAS-R1	0.90^{**}	0.84^{++}	0.89^{**}	0.14^{+}	0.88^{**}	0.26^{++}
TVP-MIDAS-SV-R1	0.90^{**}	0.99^{++}	0.89^{**}	0.14^{+}	0.89^{**}	0.26^{++}
TVP-MIDAS-SVt-R1	0.90^{**}	0.99^{++}	0.89^{**}	0.14^{+}	0.89^{**}	0.26^{++}
TVP-MIDAS-R2	0.92^{*}	0.84^{++}	0.93	0.11^{+}	0.92^{*}	0.05
TVP-MIDAS-SV-R2	0.94	0.11^{+}	0.93	0.11^{+}	0.92^{*}	0.05
TVP-MIDAS-SVt-R2	0.94	0.11^{+}	0.92^{*}	0.11^{+}	0.93^{*}	0.05
Almon lag						
MIDAS	0.97	0.08	0.95	0.03	0.95	0.05
MIDAS-SV	0.95	0.08	0.93	0.11^{+}	0.94	0.05
MIDAS-SVt	0.95	0.08	0.93	0.11^{+}	0.94	0.05
TVP-MIDAS	0.93^{*}	0.11^{+}	0.93^{*}	0.11^{+}	0.92^{*}	0.05
TVP-MIDAS-SV	0.92^{*}	0.11^{+}	0.92^{*}	0.11^{+}	0.92^{*}	0.05
TVP-MIDAS-SVt	0.92^{*}	0.11^{+}	0.92^{*}	0.11^{+}	0.92^{*}	0.05
TVP-MIDAS-R1	0.92^{*}	0.30^{++}	0.90^{**}	0.14^{+}	0.90^{**}	0.26^{++}
TVP-MIDAS-SV-R1	0.92^{*}	0.14^{+}	0.91^{*}	0.11^{+}	0.92^{*}	0.05
TVP-MIDAS-SVt-R1	0.92^{*}	0.15^{+}	0.91^{*}	0.13^{+}	0.91^{*}	0.05
TVP-MIDAS-R2	0.93^{*}	0.10^{+}	0.92^{*}	0.11^{+}	0.91^{*}	0.05
TVP-MIDAS-SV-R2	0.93	0.10^{+}	0.93^{*}	0.11^{+}	0.92^{*}	0.05
TVP-MIDAS-SVt-R2	0.93	0.10^{+}	0.93	0.11^{+}	0.92	0.05

Notes: *, **, *** denote the 10, 5, and 1 percent significant level of the Diebold–Mariano predictability test. $^{+}$ and $^{++}$ denote the forecasts in model confidence sets $\widehat{\mathcal{M}}^{*}_{90\%}$ and $\widehat{\mathcal{M}}^{*}_{75\%}$ of Hansen et al. (2011). The evaluation period starts from 1990Q1 and ends in 2019Q4.

To benchmark these TVP-MIDAS models, we compare their nowcast performance to the sparse group LASSO-based MIDAS approach developed by Babii et al. (2022).[5] Furthermore, we also include Bayesian versions of these penalized MIDAS models introduced by Mogliani and Simoni (2021), particularly the adaptive group LASSO and the adaptive group LASSO with a spike-and-slab prior. For consistency, the MIDAS models of Babii et al. (2022) and Mogliani and Simoni (2021) are implemented in an unrestricted (U)-MIDAS setting, where the temporal aggregation is fixed at 12 weeks and 60 days per quarter.[6] Table 4 provides a list of all MIDAS specifications considered in the large-scale case.

Thanks to the linear parameterizations, estimating the proposed TVP-MIDAS models is about as fast as fitting univariate time-varying coefficients regressions. Take the TVP-MIDAS model with the horseshoe prior and stochastic volatility as an example. With

[5] They focus mainly on point prediction; therefore, in what follows we compute only the point nowcasts from this framework.

[6] In our nowcasting application, we adopt the two-group structure of Bayesian MIDAS penalized models proposed by Mogliani and Simoni (2021) for simplicity. Using our real-time datasets, the unrestricted and restricted (weighting functions) models incorporate over 50 and 100 predictors, respectively.

Table 5b
The RMSFEs of large-scale MIDAS models benchmarked against the AR(2) model.

Forecast horizon	$h = 2/3$		$h = 1/3$		$h = 0$	
Models	RMSFE	p_{MCS}	RMSFE	p_{MCS}	RMSFE	p_{MCS}
U-MIDAS-AGL	1.09	0.08	1.14	0.03	1.13	0.05
U-MIDAS-AGL-SS	1.13	0.08	1.15	0.03	1.12	0.05
U-MIDAS-SG-LASSO	0.93	0.74^{++}	0.93	0.14^{+}	0.93	0.26^{++}
Fourier						
BMIDAS-AGL	1.00	0.08	0.99	0.03	0.99	0.05
BMIDAS-AGL-SS	1.01	0.08	0.99	0.03	0.99	0.05
MIDAS-SG-LASSO	0.95	0.11^{+}	0.95	0.11^{+}	0.95	0.05
TVP-MIDAS-DL	0.94	0.09	0.93	0.03	0.92	0.05
TVP-MIDAS-SV-DL	0.92	0.11^{+}	0.91	0.11^{+}	0.90^{*}	0.05
TVP-MIDAS-SVt-DL	0.92	0.11^{+}	0.92	0.11^{+}	0.90^{*}	0.05
TVP-MIDAS-NG	1.09	0.01	1.06	0.03	1.03	0.05
TVP-MIDAS-SV-NG	0.94	0.10^{+}	0.93	0.11^{+}	0.92	0.05
TVP-MIDAS-SVt-NG	0.94	0.10^{+}	0.94	0.11^{+}	0.93	0.05
TVP-MIDAS-HS	1.00	0.08	0.96	0.03	0.99	0.02
TVP-MIDAS-SV-HS	0.88^{**}	0.74^{++}	0.87^{**}	0.14^{+}	0.86^{**}	0.26^{++}
TVP-MIDAS-SVt-HS	0.87^{**}	1.00^{++}	0.86^{**}	0.14^{+}	0.85^{**}	0.28^{++}
Almon						
BMIDAS-AGL	1.00	0.08	0.99	0.03	0.98	0.05
BMIDAS-AGL-SS	1.00	0.08	0.99	0.03	0.98	0.05
MIDAS-SG-LASSO	0.97	0.10	0.94	0.11^{+}	0.94	0.05
TVP-MIDAS-DL	0.97	0.08	0.96	0.03	0.95	0.05
TVP-MIDAS-SV-DL	0.93	0.10^{+}	0.92	0.11^{+}	0.91	0.05
TVP-MIDAS-SVt-DL	0.94	0.08	0.92	0.11^{+}	0.91	0.05
TVP-MIDAS-NG	1.17	0.01	1.20	0.00	1.32	0.00
TVP-MIDAS-SV-NG	0.95	0.09	0.95	0.11^{+}	0.94	0.05
TVP-MIDAS-SVt-NG	0.96	0.08	0.97	0.03	0.96	0.05
TVP-MIDAS-HS	1.07	0.01	1.05	0.00	1.14	0.01
TVP-MIDAS-SV-HS	0.88^{**}	0.74^{++}	0.87^{**}	0.14^{+}	0.85^{**}	0.26^{++}
TVP-MIDAS-SVt-HS	0.87^{**}	0.74^{++}	0.86^{**}	1.00^{++}	0.85^{**}	1.00^{++}

Notes: *, **, *** denote the 10, 5, and 1 percent significant level of the Diebold–Mariano predictability test. $^{+}$ and $^{++}$ denote the forecasts in model confidence sets $\widehat{\mathcal{M}}^{*}_{90\%}$ and $\widehat{\mathcal{M}}^{*}_{75\%}$ of Hansen et al. (2011). The evaluation period starts from 1990Q1 and ends in 2019Q4.

10 monthly predictors, it takes about half a minute to sample 1000 posterior draws on a standard desktop with an Intel Xeon W-2223 @ 3.6 GHz processor and 16 GB of RAM; with 40 monthly predictors, the computation time is about 2 min.

5.2.1. Out-of-sample nowcast performance before the COVID-19 period

We first assess the point nowcast performance of the MIDAS specifications using the root mean squared forecast error (RMSFE) for the subsample that ends in 2019Q4. Tables 5a and 5b present the RMSFEs of MIDAS models relative to a simple AR(2) benchmark. Values less than one indicate better nowcast performance relative to the benchmark. Additionally, we report the associated model confidence set (MCS) p-values proposed in Hansen et al. (2011), calculated based on the full set of all MIDAS models listed in Tables 3 and 4.

First, Table 5a reports the relative RMSFEs of the MIDAS variants nested within the proposed framework. All values are less than one, underscoring the superior point nowcast performance of the MIDAS models compared to the AR(2) benchmark. Among the TVP-MIDAS specifications, the versions that allow only time-varying coefficients on the high-frequency predictors (i.e., $\boldsymbol{\beta}_t$ is time-varying) have the best point nowcast performance, although other TVP-MIDAS specifications perform similarly. For MIDAS models with constant coefficients, adding stochastic volatility or t errors slightly improves point nowcasts. But for TVP-MIDAS models, allowing more flexible errors does not seem to substantially improve (nor degrade) their nowcast accuracy. In fact, the only specification that is consistently excluded from $\widehat{\mathcal{M}}^{*}_{90\%}$, the model confidence set with 90% coverage, is the conventional MIDAS model with constant coefficients and homoskedastic Gaussian errors. Overall, these findings demonstrate the importance of allowing some form of time-variation in MIDAS models. In particular, incorporating time-varying parameters into MIDAS models can often enhance nowcast accuracy.

Next, Table 5b presents nowcast results for a range of large-scale MIDAS models, including various machine learning-based MIDAS models developed by Mogliani and Simoni (2021) and Babii et al. (2022). Since these machine learning-based MIDAS models all assume constant parameters, comparing the TVP-MIDAS models with these benchmarks helps illustrate the value of allowing time-valuation in the presence of a large number of predictors. The results show that the proposed TVP-MIDAS models with stochastic volatility generally perform well compared to these state-of-the-art benchmarks. In particular, the best performing specifications are the TVP-MIDAS models with the horseshoe prior and stochastic volatility, although TVP-MIDAS models with other shrinkage priors also perform similarly well. The results based on the model confidence set confirm these findings: among the 9

Table 6a
The average CRPSs of the variants nested within the proposed TVP-MIDAS framework benchmarked against the AR(2).

Forecast horizon	$h = 2/3$		$h = 1/3$		$h = 0$	
Models	CRPS	p_{MCS}	CRPS	p_{MCS}	CRPS	p_{MCS}
Fourier						
MIDAS	0.95	0.00	0.96	0.00	0.97	0.00
MIDAS-SV	0.93*	0.11^{+}	0.94	0.03	0.95	0.00
MIDAS-SVt	0.93*	0.11^{+}	0.94	0.03	0.95	0.00
TVP-MIDAS	0.92**	0.11^{+}	0.92**	0.03	0.92**	0.01
TVP-MIDAS-SV	0.93*	0.11^{+}	0.93**	0.03	0.93**	0.01
TVP-MIDAS-SVt	0.93*	0.11^{+}	0.93*	0.03	0.93**	0.01
TVP-MIDAS-R1	0.91**	0.84^{++}	0.90***	0.42^{++}	0.90***	0.28^{++}
TVP-MIDAS-SV-R1	0.90***	0.99^{++}	0.89***	0.46^{++}	0.89***	0.28^{++}
TVP-MIDAS-SVt-R1	0.90***	0.99^{++}	0.89***	0.88^{++}	0.89***	0.57^{++}
TVP-MIDAS-R2	0.92**	0.84^{++}	0.93*	0.03	0.92*	0.28^{++}
TVP-MIDAS-SV-R2	0.94*	0.11^{+}	0.93*	0.03	0.93*	0.01
TVP-MIDAS-SVt-R2	0.94*	0.11^{+}	0.93*	0.03	0.93*	0.01
Almon						
MIDAS	0.96	0.00	0.95	0.00	0.96	0.00
MIDAS-SV	0.95	0.11^{+}	0.94	0.03	0.94	0.01
MIDAS-SVt	0.94	0.11^{+}	0.94*	0.03	0.94	0.01
TVP-MIDAS	0.92**	0.11^{+}	0.92**	0.03	0.92**	0.01
TVP-MIDAS-SV	0.93*	0.11^{+}	0.93**	0.03	0.92**	0.01
TVP-MIDAS-SVt	0.93*	0.11^{+}	0.93*	0.03	0.93**	0.01
TVP-MIDAS-R1	0.92**	0.11^{+}	0.92**	0.03	0.91**	0.01
TVP-MIDAS-SV-R1	0.92**	0.11^{+}	0.92**	0.03	0.91**	0.01
TVP-MIDAS-SVt-R1	0.92**	0.11^{+}	0.91**	0.04	0.91**	0.01
TVP-MIDAS-R2	0.92**	0.11^{+}	0.92**	0.03	0.91**	0.01
TVP-MIDAS-SV-R2	0.94*	0.11^{+}	0.93*	0.03	0.93*	0.01
TVP-MIDAS-SVt-R2	0.94*	0.11^{+}	0.93*	0.03	0.93*	0.01

Notes: *, **, *** denote the 10, 5, and 1 percent significant level of the Diebold–Mariano predictability test. $^{+}$ and $^{++}$ denote the forecasts in model confidence sets $\widehat{\mathcal{M}}^{*}_{90\%}$ and $\widehat{\mathcal{M}}^{*}_{75\%}$ of Hansen et al. (2011). The evaluation period starts from 1990Q1 and ends in 2019Q4.

models included in $\widehat{\mathcal{M}}^{*}_{75\%}$ for forecast horizon $h = 0$, 8 feature time-varying parameters. These findings again highlight the empirical relevance of accommodating time-variation in MIDAS models.

We evaluate the density nowcasts of all MIDAS specifications using the Continuous Ranked Probability Score (CRPS). Tables 6a and 6b present the average CRPS values relative to the AR(2) benchmark for each specification. These results are consistent with the findings from the point nowcasts. For example, the TVP-MIDAS models with only time-varying β_t deliver the most accurate density forecasts among all the variants nested within the proposed TVP-MIDAS framework; the TVP-MIDAS models with the horseshoe prior and stochastic volatility consistently outperform other global-local priors and machine learning-based MIDAS models; the model confidence sets include mostly MIDAS models with time-varying parameters. Overall, these results show that extending constant-coefficients MIDAS models to incorporate time-variation can significantly enhance both point and density nowcast accuracy.

In nowcasting GDP, the focus is often in quantifying the left tail risk—the risk of a downturn or recession. We therefore delve deeper into the performance of the MIDAS specifications concerning the left tail of the density nowcasts. Drawing from the methodology outlined by Gneiting and Ranjan (2011), we calculate the predictive quantile score at time t for a given quantile τ, expressed as:

$$QS_{\tau,t} = (y_t - Q_{\tau,t}) - (\tau - 1\{y_t \leq Q_{\tau,t}\}),$$

where $1\{y_t \leq Q_{\tau,t}\}$ takes the value 1 if the realized value of the GDP is at or below the predictive quantile and 0 otherwise. We assess the performance of the quantile score in the left tail by setting $\tau = 0.1$.

We present the average quantile scores for each model relative to the AR(2) benchmark in Tables 7a and 7b. Among the MIDAS variants nested within the proposed framework, results in Table 7a show that TVP-MIDAS models consistently outperform their time-invariant counterparts in nowcasting the left tail of the distribution. In particular, in line with the results based on the whole nowcast density, the versions that allow only time-varying β_t perform best in terms of quantifying the left tail risk.

Among the large-scale MIDAS models, results reported in Table 7b show that the TVP-MIDAS models with global-local shrinkage priors, particularly the Horseshoe prior, achieve comparable performance in nowcasting the left tail of the GDP distribution relative to machine learning-based MIDAS models. But these large-scale models tend to be dominated by small-scale MIDAS models that use only the industrial production, NFCI and an interest rate spread as predictors. In fact, the only model included in the model confidence sets $\widehat{\mathcal{M}}^{*}_{90\%}$ and $\widehat{\mathcal{M}}^{*}_{75\%}$ of Hansen et al. (2011) is the small-scale TVP-MIDAS model with stochastic volatility, time-varying β_t but constant-coefficients weighting functions based on the Almon lag polynomials. These findings are visually represented in Fig. 3, which illustrates the rolling average of CRPS and quantile scores for $h = 0$ over time for selected TVP-MIDAS specifications.

Table 6b
The average CRPSs of large-scale MIDAS models benchmarked against the AR(2) model.

Forecast horizon	$h = 2/3$		$h = 1/3$		$h = 0$	
Models	CRPS	p_{MCS}	CRPS	p_{MCS}	CRPS	p_{MCS}
U-MIDAS-AGL	1.08	0.00	1.12	0.00	1.11	0.00
U-MIDAS-AGL-SS	1.12	0.00	1.13	0.00	1.10	0.00
Fourier						
BMIDAS-AGL	1.00	0.00	0.99	0.01	0.99	0.00
BMIDAS-AGL-SS	1.00	0.00	0.99	0.00	0.99	0.00
TVP-MIDAS-DL	0.96	0.00	0.96	0.00	0.96	0.00
TVP-MIDAS-SV-DL	0.96	0.00	0.95	0.00	0.95	0.00
TVP-MIDAS-SVt-DL	0.95	0.00	0.96	0.00	0.95	0.00
TVP-MIDAS-NG	1.09	0.00	1.07	0.00	1.04	0.00
TVP-MIDAS-SV-NG	0.96	0.00	0.96	0.01	0.96	0.00
TVP-MIDAS-SVt-NG	0.96	0.00	0.97	0.00	0.96	0.00
TVP-MIDAS-HS	0.99	0.00	0.98	0.00	1.01	0.00
TVP-MIDAS-SV-HS	0.90^{**}	0.84^{++}	0.89^{**}	0.42^{++}	0.89^{**}	0.28^{++}
TVP-MIDAS-SVt-HS	0.90^{**}	1.00^{++}	0.89^{**}	0.46^{++}	0.88^{**}	0.57^{++}
Almon						
BMIDAS-AGL	1.00	0.11	0.99	0.03	0.98	0.01
BMIDAS-AGL-SS	1.00	0.11	0.99	0.01	0.98	0.01
TVP-MIDAS-DL	0.98	0.00	0.98	0.00	0.96	0.00
TVP-MIDAS-SV-DL	0.96	0.00	0.95	0.00	0.95	0.00
TVP-MIDAS-SVt-DL	0.96	0.00	0.95	0.00	0.94	0.00
TVP-MIDAS-NG	1.17	0.00	1.22	0.00	1.30	0.00
TVP-MIDAS-SV-NG	0.97	0.00	0.97	0.00	0.97	0.00
TVP-MIDAS-SVt-NG	0.98	0.00	0.99	0.00	0.97	0.00
TVP-MIDAS-HS	1.08	0.00	1.06	0.00	1.13	0.00
TVP-MIDAS-SV-HS	0.90^{**}	0.96^{++}	0.89^{**}	0.42^{++}	0.88^{**}	0.57^{++}
TVP-MIDAS-SVt-HS	0.90^{**}	0.99^{++}	0.89^{**}	1.00^{++}	0.88^{**}	1.00^{++}

Notes: *, **, *** denote the 10, 5, and 1 percent significant level of the Diebold–Mariano predictability test. $^{+}$ and $^{++}$ denote the forecasts in model confidence sets $\widehat{\mathcal{M}}^{*}_{90\%}$ and $\widehat{\mathcal{M}}^{*}_{75\%}$ of Hansen et al. (2011). The evaluation period starts from 1990Q1 and ends in 2019Q4.

Table 7a
The average quantile scores at $\tau = 0.1$ for the MIDAS models nested within the proposed TVP-MIDAS framework relative to the AR(2) benchmark.

Forecast horizon	$h = 2/3$		$h = 1/3$		$h = 0$	
Models	QS	p_{MCS}	QS	p_{MCS}	QS	p_{MCS}
Fourier						
MIDAS	0.97	0.00	0.98	0.00	0.98	0.00
MIDAS-SV	0.95	0.00	0.97	0.00	0.97	0.00
MIDAS-SVt	0.95	0.00	0.97	0.00	0.97	0.00
TVP-MIDAS	0.89^{***}	0.00	0.90^{***}	0.00	0.90^{***}	0.00
TVP-MIDAS-SV	0.79^{***}	0.05	0.79^{***}	0.04	0.79^{***}	0.04
TVP-MIDAS-SVt	0.80^{***}	0.00	0.80^{***}	0.00	0.80^{***}	0.00
TVP-MIDAS-R1	0.78^{***}	0.02	0.79^{***}	0.01	0.79^{***}	0.01
TVP-MIDAS-SV-R1	0.74^{***}	0.05	0.74^{***}	0.04	0.74^{***}	0.04
TVP-MIDAS-SVt-R1	0.74^{***}	0.05	0.74^{***}	0.04	0.74^{***}	0.04
TVP-MIDAS-R2	0.85^{***}	0.00	0.86^{***}	0.00	0.86^{***}	0.00
TVP-MIDAS-SV-R2	0.80^{***}	0.02	0.80^{***}	0.01	0.80^{***}	0.01
TVP-MIDAS-SVt-R2	0.81^{***}	0.00	0.80^{***}	0.00	0.81^{***}	0.00
Almon						
MIDAS	0.95	0.00	0.95	0.00	0.95	0.00
MIDAS-SV	0.95^{*}	0.00	0.95^{*}	0.00	0.95	0.00
MIDAS-SVt	0.95^{*}	0.00	0.95^{*}	0.00	0.95	0.00
TVP-MIDAS	0.89^{***}	0.00	0.90^{***}	0.00	0.90^{***}	0.00
TVP-MIDAS-SV	0.79^{***}	0.02	0.79^{***}	0.01	0.79^{***}	0.01
TVP-MIDAS-SVt	0.80^{***}	0.00	0.80^{***}	0.00	0.80^{***}	0.00
TVP-MIDAS-R1	0.77^{***}	0.05	0.77^{***}	0.04	0.77^{***}	0.04
TVP-MIDAS-SV-R1	0.72^{***}	1.00^{++}	0.73^{***}	1.00^{++}	0.73^{***}	1.00^{++}
TVP-MIDAS-SVt-R1	0.73^{***}	0.05	0.73^{***}	0.04	0.73^{***}	0.04

(*continued on next page*)

Table 7 (*continued*).

Forecast horizon	$h = 2/3$		$h = 1/3$		$h = 0$	
Models	QS	p_{MCS}	QS	p_{MCS}	QS	p_{MCS}
TVP-MIDAS-R2	0.87***	0.00	0.88***	0.00	0.87***	0.00
TVP-MIDAS-SV-R2	0.79***	0.05	0.79***	0.04	0.79***	0.04
TVP-MIDAS-SVt-R2	0.80***	0.00	0.80***	0.01	0.79***	0.01

Notes: *, **, *** denote the 10, 5, and 1 percent significant level of the Diebold–Mariano predictability test. $^{+}$ and $^{++}$ denote the forecasts in model confidence sets $\widehat{\mathcal{M}}^{*}_{90\%}$ and $\widehat{\mathcal{M}}^{*}_{75\%}$ of Hansen et al. (2011). The evaluation period starts from 1990Q1 and ends in 2019Q4.

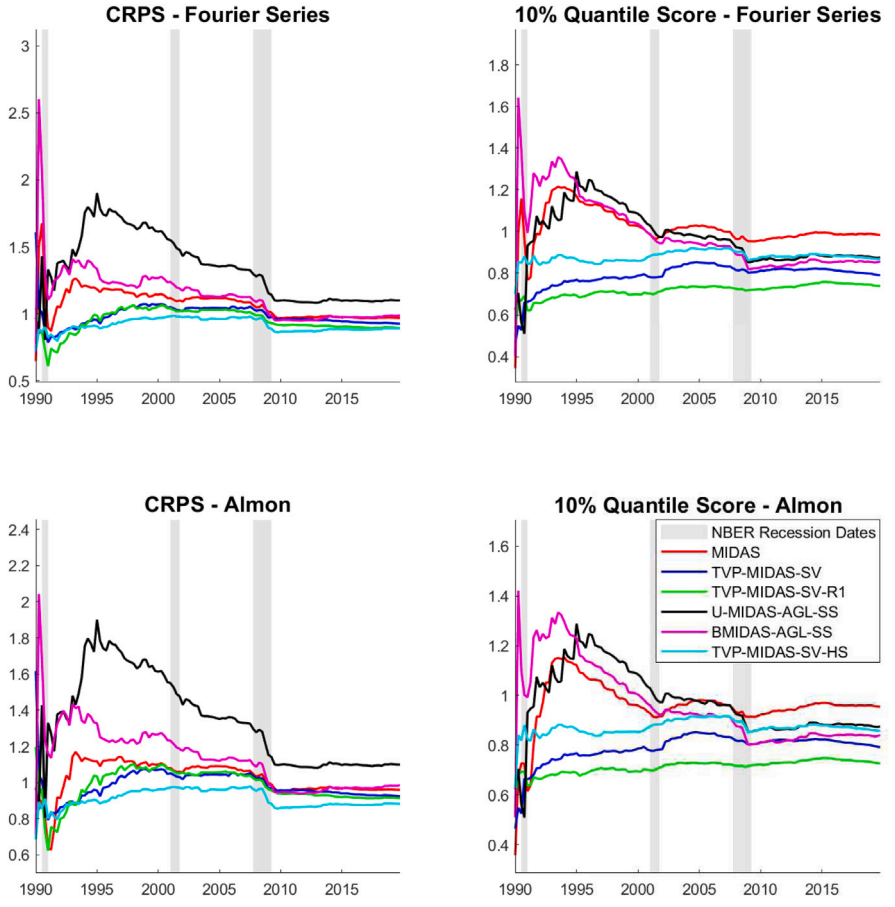

Fig. 3. Plot of the rolling average CRPS and quantile scores for $h = 0$ across the selected six MIDAS models relative to the AR(2) model, over the pre-COVID-19 period.

Our results indicate that during periods of heightened volatility, using a parsimonious set of predictors, alongside the inclusion of time-varying parameters and stochastic volatility, is essential for accurately nowcasting the left tail of GDP or economic downturns. This conclusion is consistent with the findings of Adrian et al. (2019) and Estrella and Hardouvelis (1991), who highlight the predictive importance of financial conditions and the slope of the yield curve for future recessions.

5.2.2. Out-of-sample nowcast performance through the COVID-19 period

Next, we assess the performance of the MIDAS models using a sample that includes the COVID-19 pandemic, with the evaluation period ending in 2021Q2. Specifically, we produce the same sets of nowcast results — RMSFEs, CRPSs and quantile scores — for all the MIDAS models described in Tables 3 and 4. Due to space constraints, the detailed results are reported in Online Appendix A. Given the extreme movements of US GDP and other macroeconomic variables during the COVID-19 lockdown and the subsequent reopening, all models incur significant nowcast errors during this period, and this naturally makes inference more difficult. Despite this, the main conclusion remains consistent: TVP-MIDAS models with stochastic volatility generally outperform their time-invariant counterparts, especially for density nowcasting. For instance, among the 16 models included in the model confidence set $\widehat{\mathcal{M}}^{*}_{75\%}$ for density nowcasts (CRPSs) with horizon $h = 0$, 11 feature time-varying parameters.

Table 7b
The average quantile scores at $\tau = 0.1$ for large-scale MIDAS models relative to the AR(2) benchmark.

Forecast horizon	$h = 2/3$		$h = 1/3$		$h = 0$	
Models	QS	p_{MCS}	QS	p_{MCS}	QS	p_{MCS}
U-MIDAS-AGL	0.85**	0.00	0.90	0.00	0.89*	0.00
U-MIDAS-AGL-SS	0.88**	0.00	0.90*	0.00	0.87**	0.00
Fourier						
BMIDAS-AGL	0.84	0.00	0.85***	0.00	0.86**	0.00
BMIDAS-AGL-SS	0.84	0.00	0.85***	0.00	0.86**	0.00
TVP-MIDAS-DL	0.96	0.00	0.97	0.00	0.98	0.00
TVP-MIDAS-SV-DL	0.97	0.00	0.98	0.00	0.99	0.00
TVP-MIDAS-SVt-DL	0.97*	0.00	0.99*	0.00	0.99	0.00
TVP-MIDAS-NG	0.97**	0.00	1.00	0.00	0.95	0.00
TVP-MIDAS-SV-NG	0.94***	0.00	0.94***	0.00	0.94*	0.00
TVP-MIDAS-SVt-NG	0.94***	0.00	0.95***	0.00	0.94*	0.00
TVP-MIDAS-HS	0.87***	0.00	0.86***	0.00	0.91**	0.00
TVP-MIDAS-SV-HS	0.86**	0.00	0.86	0.00	0.87***	0.00
TVP-MIDAS-SVt-HS	0.86**	0.00	0.87*	0.00	0.87***	0.00
Almon						
BMIDAS-AGL	0.83***	0.05	0.84***	0.04	0.84***	0.01
BMIDAS-AGL-SS	0.83***	0.02	0.84***	0.01	0.84***	0.00
TVP-MIDAS-DL	0.92**	0.00	0.92**	0.00	0.91**	0.00
TVP-MIDAS-SV-DL	0.96	0.00	0.97	0.00	0.97	0.00
TVP-MIDAS-SVt-DL	0.96	0.00	0.97	0.00	0.97	0.00
TVP-MIDAS-NG	1.00	0.00	1.00	0.00	1.08	0.00
TVP-MIDAS-SV-NG	0.93**	0.00	0.95**	0.00	0.95	0.00
TVP-MIDAS-SVt-NG	0.94*	0.00	0.96***	0.00	0.95	0.00
TVP-MIDAS-HS	0.93	0.00	0.91***	0.00	0.91*	0.00
TVP-MIDAS-SV-HS	0.86***	0.00	0.86	0.00	0.86***	0.00
TVP-MIDAS-SVt-HS	0.86***	0.00	0.86*	0.00	0.86***	0.00

Notes: *, **, *** denote the 10, 5, and 1 percent significant level of the Diebold–Mariano predictability test. $^{+}$ and $^{++}$ denote the forecasts in model confidence sets $\widehat{\mathcal{M}}^{*}_{90\%}$ and $\widehat{\mathcal{M}}^{*}_{75\%}$ of Hansen et al. (2011). The evaluation period starts from 1990Q1 and ends in 2019Q4.

In addition, for nowcasting left-tail risks, small-scale MIDAS models dominate. For example, the only models included in the model confidence set $\widehat{\mathcal{M}}^{*}_{75\%}$ are small-scale TVP-MIDAS models. Overall, our findings underscore the importance of incorporating time-varying parameters and stochastic volatility for nowcasting US GDP. The proposed framework is robust, flexible, and well-suited to capture unpredictable volatility arising from extreme events, such as the COVID-19 pandemic.

6. Concluding remarks and future research

This paper introduces a novel TVP-MIDAS framework that is flexible and easy to estimate. We evaluate the effectiveness of the proposed framework using a real-time application of nowcasting US real GDP. Leveraging on three high-frequency predictors — the monthly industrial production, the weekly NFCI and a daily interest rate spread — our results demonstrate that TVP-MIDAS specifications incorporating stochastic volatility consistently outperform their time-invariant counterparts. Specifically, our findings reveal that the proposed TVP-MIDAS framework yields superior nowcasts, particularly in capturing the left tail risk of the GDP. Among all TVP-MIDAS specifications, the versions that allow only time-varying coefficients on the high-frequency predictors tend to nowcast better than those with only time-varying weights, suggesting that time-varying weights are less empirically important than other forms of time-variations.

For future work, it would be interesting to extend the proposed TVP-MIDAS framework to the VAR setting, building upon the MIDAS-VAR approach suggested in Ghysels (2016). This multivariate extension is especially useful for jointly nowcasting multiple variables using higher frequency predictors or computing conditional forecasts based on the future paths of certain variables.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.jeconom.2025.106090.

References

Adrian, T., Boyarchenko, N., Giannone, D., 2019. Vulnerable growth. Am. Econ. Rev. 109 (4), 1263–1289.
Amburgey, A.J., McCracken, M.W., 2023. On the real-time predictive content of financial condition indices for growth. J. Appl. Econometrics 38 (2), 137–163.
Andreou, E., 2016. On the use of high frequency measures of volatility in MIDAS regressions. J. Econometrics 193 (2), 367–389.
Babii, A., Ghysels, E., Striaukas, J., 2022. Machine learning time series regressions with an application to nowcasting. J. Bus. Econom. Statist. 40 (3), 1094–1106.
Barnett, A., Mumtaz, H., Theodoridis, K., 2014. Forecasting UK GDP growth and inflation under structural change. A comparison of models with time-varying parameters. Int. J. Forecast. 30 (1), 129–143.
Bekierman, J., Gribisch, B., 2021. A mixed frequency stochastic volatility model for intraday stock market returns. J. Financ. Econ. 19 (3), 496–530.
Bitto, A., Frühwirth-Schnatter, S., 2019. Achieving shrinkage in a time-varying parameter model framework. J. Econometrics 210 (1), 75–97.
Bobeica, E., Hartwig, B., 2023. The COVID-19 shock and challenges for inflation modelling. Int. J. Forecast. 39 (1), 519–539.
Carriero, A., Clark, T.E., Marcellino, M., 2015. Realtime nowcasting with a Bayesian mixed frequency model with stochastic volatility. J. R. Stat. Soc. Ser. A: Stat. in Soc. 178 (4), 837–862.
Carriero, A., Clark, T.E., Marcellino, M., Mertens, E., 2024. Addressing COVID-19 outliers in BVARs with stochastic volatility. Rev. Econ. Stat. 106 (5), 1403–1417, MIT Press 255 Main Street, 9th Floor, Cambridge, Massachusetts 02142, USA ....
Cascaldi-Garcia, D., Luciani, M., Modugno, M., 2024. Lessons from nowcasting GDP across the world. In: Handbook of Research Methods and Applications in Macroeconomic Forecasting. Edward Elgar Publishing, pp. 187–217.
Chan, J.C., Eisenstat, E., 2018. Bayesian model comparison for time-varying parameter VARs with stochastic volatility. J. Appl. Econometrics 33 (4), 509–532.
Chan, J.C., Hsiao, C.Y., 2014. Estimation of stochastic volatility models with heavy tails and serial dependence. In: Bayesian Inference in the Social Sciences. Wiley Online Library, pp. 155–176.
Chan, J.C., Jeliazkov, I., 2009. Efficient simulation and integrated likelihood estimation in state space models. Int. J. Math. Model. Numer. Optim. 1 (1–2), 101–120.
Clark, T.E., 2011. Real-time density forecasts from Bayesian vector autoregressions with stochastic volatility. J. Bus. Econom. Statist. 29 (3), 327–341.
Clark, T.E., Ravazzolo, F., 2015. Macroeconomic forecasting performance under alternative specifications of time-varying volatility. J. Appl. Econometrics 30 (4), 551–575.
Cong, Y., Chen, B., Zhou, M., 2017. Fast simulation of hyperplane-truncated multivariate normal distributions. Bayesian Anal. 12 (4), 1017–1037.
Cross, J., Hou, C., Koop, G., Poon, A., 2023. Large stochastic volatility in mean VARs. J. Econometrics 236 (1), 105469.
Cross, J., Poon, A., 2016. Forecasting structural change and fat-tailed events in Australian macroeconomic variables. Econ. Model. 58, 34–51.
D'Agostino, A., Gambetti, L., Giannone, D., 2013. Macroeconomic forecasting and structural change. J. Appl. Econometrics 28 (1), 82–101.
Estrella, A., Hardouvelis, G.A., 1991. The term structure as a predictor of real economic activity. J. Financ. 46 (2), 555–576.
Estrella, A., Rodrigues, A.P., Schich, S., 2003. How stable is the predictive power of the yield curve? Evidence from Germany and the United States. Rev. Econ. Stat. 85 (3), 629–644.
Foroni, C., Marcellino, M., 2014. A comparison of mixed frequency approaches for nowcasting Euro area macroeconomic aggregates. Int. J. Forecast. 30 (3), 554–568.
Frühwirth-Schnatter, S., Wagner, H., 2010. Stochastic model specification search for Gaussian and partial non-Gaussian state space models. J. Econometrics 154, 85–100.
Gelman, A., Goodrich, B., Gabry, J., Vehtari, A., 2019. R-squared for Bayesian regression models. Amer. Statist..
Ghysels, E., 2016. Macroeconomics and the reality of mixed frequency data. J. Econometrics 193 (2), 294–314.
Ghysels, E., Santa-Clara, P., Valkanov, R., 2005. There is a risk-return trade-off after all. J. Financ. Econ. 76 (3), 509–548.
Ghysels, E., Sinko, A., Valkanov, R., 2007. MIDAS regressions: Further results and new directions. Econometric Rev. 26 (1), 53–90.
Gneiting, T., Ranjan, R., 2011. Comparing density forecasts using threshold-and quantile-weighted scoring rules. J. Bus. Econom. Statist. 29 (3), 411–422.
Guérin, P., Marcellino, M., 2013. Markov-switching MIDAS models. J. Bus. Econom. Statist. 31 (1), 45–56.
Hansen, P.R., Lunde, A., Nason, J.M., 2011. The model confidence set. Econometrica 79 (2), 453–497.
Huber, F., Koop, G., Onorante, L., 2021. Inducing sparsity and shrinkage in time-varying parameter models. J. Bus. Econom. Statist. 39 (3), 669–683.
Kim, S., Shephard, N., Chib, S., 1998. Stochastic volatility: Likelihood inference and comparison with ARCH models. Rev. Econ. Stud. 65 (3), 361–393.
Koop, G., Korobilis, D., 2013. Large time-varying parameter VARs. J. Econometrics 177 (2), 185–198.
Kuzin, V., Marcellino, M., Schumacher, C., 2011. MIDAS vs. mixed-frequency VAR: Nowcasting GDP in the Euro area. Int. J. Forecast. 27 (2), 529–542.
Marcellino, M., Schumacher, C., 2010. Factor MIDAS for nowcasting and forecasting with ragged-edge data: A model comparison for German GDP. Oxf. Bull. Econ. Stat. 72 (4), 518–550.
Mogliani, M., Simoni, A., 2021. Bayesian MIDAS penalized regressions: estimation, selection, and prediction. J. Econometrics 222 (1), 833–860.
Pettenuzzo, D., Timmermann, A., Valkanov, R., 2016. A MIDAS approach to modeling first and second moment dynamics. J. Econometrics 193 (2), 315–334.
Poon, A., Zhu, D., 2024. Do recessions and bear markets occur concurrently across countries? A multinomial logistic approach. J. Financ. Econ..
Potjagailo, G., Kohns, D., 2023. Flexible Bayesian MIDAS: time-variation, group-shrinkage and sparsity. In: Bank of England Working Paper.
Primiceri, G.E., 2005. Time varying structural vector autoregressions and monetary policy. Rev. Econ. Stud. 72 (3), 821–852.
Rudebusch, G.D., Williams, J.C., 2009. Forecasting recessions: The puzzle of the enduring power of the yield curve. J. Bus. Econom. Statist. 27 (4), 492–503.
Rue, H., Held, L., 2005. Gaussian Markov Random Fields: Theory and Applications, CRC Press.
Schumacher, C., 2014. MIDAS Regressions with Time-Varying Parameters: An Application to Corporate Bond Spreads and GDP in the Euro Area. Working Paper.
Stock, J.H., Watson, M.W., 2016. Core inflation and trend inflation. Rev. Econ. Stat. 98 (4), 770–784.
Wang, L., Ma, F., Liu, J., Yang, L., 2020. Forecasting stock price volatility: New evidence from the GARCH-MIDAS model. Int. J. Forecast. 36 (2), 684–694.