Time-Varying Parameter MIDAS Models: Application to Nowcasting US Real GDP

Joshua C. C. Chan
Purdue University

Aubrey Poon
University of Kent

Dan Zhu
Monash University

April 24, 2024

Abstract

We propose a novel time-varying parameter mixed-data sampling (TVP-MIDAS) framework. Specifically, we decompose the MIDAS coefficients into a scalar parameter representing the overall impact of high-frequency variables and a vector of weights, allowing both features to vary over time. Our study applies this framework to real-time forecasting for US real GDP. Our analysis demonstrates that the TVP-MIDAS model specifications produce superior point forecasts and are particularly effective in capturing left tail risk compared to their time-invariant counterparts. Additionally, our in-sample analysis reveals a significant negative trend in the influence of the National Financial Conditions Index (NFCI) on US real GDP, suggesting a progressively adverse correlation over time. Conversely, the impact of the yield curve slope on US real GDP exhibits minimal variation over time.

To do:

1. cite Bekierman and Gribisch (2021)

1 Introduction

Mixed-data sampling (MIDAS) regressions have garnered significant attention in empirical macroeconomics for their utility in nowcasting key macroeconomic indicators like real GDP and inflation. A notable advantage of MIDAS lies in its straightforward and parsimonious framework, enabling the sampling of left-hand and right-hand variables of a time series regression at differing frequencies. For instance, a prevalent application of MIDAS in the literature involves forecasting real GDP, a quarterly variable, utilizing high-frequency predictors such as monthly industrial production or employment (see Marcellino and Schumacher, 2010; Kuzin, Marcellino, and Schumacher, 2011; Foroni and Marcellino, 2014; Mogliani and Simoni, 2021). High-frequency predictors offer the advantage of timelier release compared to quarterly real GDP data. Leveraging a MIDAS framework allows forecasters to exploit this timely information, enabling the production of a nowcast of real GDP prior to the official release by the statistical agency. Moreover, MIDAS has extended its prominence beyond empirical macroeconomics into financial settings, particularly in forecasting stock price volatility (see Andreou, 2016; Wang, Ma, Liu, and Yang, 2020).

Within the realm of MIDAS regressions, the standard approach is to assume constant parameters. This is mostly due to necessity, as 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 time-varying settings typically involves the estimation of nonlinear state space models, which tends to be computationally intensive. On the other hand, assuming constant parameters in MIDAS regressions appears to be overly restrictive given the documented importance of allowing for time-variation in model parameters when forecasting macroeconomic variables (see Primiceri, 2005; Barnett, Mumtaz, and Theodoridis, 2014; Koop and Korobilis, 2013; D'Agostino, Gambetti, and Giannone, 2013; Chan and Eisenstat, 2018). Furthermore, macroeconomic variables often demonstrate time-varying conditional distributions, reflecting fluctuations in government policies, global economic conditions, technological advancements, and other socioeconomic factors. Comprehensive understanding and modeling of these time-varying distributions are imperative for accurate risk assessment, informed policy formulation, and the anticipation of future economic trends.

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. To that end, 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 imperative in MIDAS regressions, they can prove beneficial in certain applications, as exemplified in Ghysels, Santa-Clara, and Valkanov (2005), where the overall impact coefficient has an interesting economic interpretation.

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 irregularly spaced mixed-frequency data and settings with multiple high-frequency predictors.

Our research 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 variant of the proposed TVP-MIDAS model, as they solely permit 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 apply the proposed TVP-MIDAS framework to a real-time forecasting application for quarterly US real GDP. Within this framework, we incorporate two high-frequency predictors: the daily interest rate spread, representing the slope of the yield curve (defined as the difference between the 10-year and 3-month treasury yields), and a weekly national financial condition index (NFCI). Our forecasting process involves generating three forecasts or nowcasts at the end of each month within the quarter, with our evaluation period spanning from 1990Q1 to 2021Q2.

Our findings indicate that the TVP-MIDAS model specifications yield superior point forecasts compared to their corresponding static counterparts. However, when assessing density forecasts, we observe that the TVP-MIDAS model specifications with stochastic volatility demonstrate similar performance to their static counterparts. Furthermore, we explore the forecasting performance of our TVP-MIDAS model regarding the tail risk of US real GDP. Specifically, our analysis reveals that TVP-MIDAS model specifications with stochastic volatility outperform their static counterparts in forecasting the left tail of US real GDP. Conversely, for the right tail of US real GDP, the static MIDAS models exhibit superior forecasting ability compared to the TVP-MIDAS models. Therefore, our results suggest that during periods of heightened volatility, the incorporation of time-varying parameters and stochastic volatility is crucial for accurately forecasting the left tail of US real GDP or recessionary events. This conclusion resonates with the findings of Adrian, Boyarchenko, and Giannone (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. In contrast, during tranquil periods, a static time-invariant MIDAS model proves adequate for forecasting the right tail of US real GDP.

Finally, we estimate our TVP-MIDAS model with stochastic volatility using the final data vintage of 2021Q2. Our in-sample analysis reveals that the influence of the NFCI exhibits a significant negative trend, implying a progressively adverse correlation between US real GDP and NFCI over time. This finding aligns with the conclusions drawn by Adrian,

Boyarchenko, and Giannone (2019). Conversely, our analysis reveals that the impact of the yield curve slope on US real GDP dynamics exhibits minimal significance over time.

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 presents the real-time out-of-sample application and the in-sample analysis. Finally, Section 5 concludes.

2 The Econometric Framework

3 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 $x_{t-k/m}^{(m)}, k = 0, \ldots, m-1$, where $x_{t-(m-1)/m}^{(m)}$ and $x_t^{(m)}$ 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 $x_t^{(m)}$ with $m = 22$, if we assume that there are 22 daily available observations within each month. In Subsection 3.3, we will consider more complex settings where the numbers of observations of the high-frequency variable between discrete time periods are not constant.

3.1 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, $\frac{1}{m}\sum_{k=0}^{m-1} x_{t-k/m}^{(m)}$, as a single predictor. More specifically, let $h \geq 1$ denote the forecast horizon, and consider the following direct forecasting approach:

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

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

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

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 $x_t^{(m)}$ for predicting y_t. As a simple example, consider the predictive regression

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

where $\mathbf{w}_t$ is an $m \times 1$ vector of weights and $\mathbf{x}_t^{(m)} = [x_t^{(m)}, x_{t-1/m}^{(m)}, \ldots, x_{t-(m-1)/m}^{(m)}]'$. 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 $\mathbf{w}_t = [1, 0, \ldots, 0]'$, respectively.

Following Ghysels, Sinko, and Valkanov (2007) and Pettenuzzo, Timmermann, and Valkanov (2016), we consider a general MIDAS regression of the form

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

where the scalar β captures the overall impact of the lagged values of $x_t^{(m)}$ on y_{t+h}, $\boldsymbol{\rho}$ is the vector of autoregressive coefficients on $\mathbf{y}_t = [y_t, y_{t-1}, \ldots, y_{t-p_y}]'$, 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 $L^{k/m} x_t^{(m)} = x_{t-k/m}^{(m)}$ and each component function $B(k; \boldsymbol{\theta})$ depends on a low-dimensional vector of parameters $\boldsymbol{\theta}$.

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

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

where $\mathbf{v}_k = [\phi_0(k), \ldots, \phi_p(k)]'$. As an example, this formulation recovers the widely used Almon lag polynomial by setting $\mathbf{v}_k = [1, k, k^2, \ldots, k^p]'$, 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 $\phi_j(s) = s^j, j = 0, 1, \ldots, p$, 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

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

even — that forms an orthonormal basis (for square-integrable functions). 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.

3.2 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, Mumtaz, and Theodoridis, 2014; Koop and Korobilis, 2013; D'Agostino, Gambetti, and Giannone, 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.[2]

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, Clark, and Marcellino (2015) and Pettenuzzo, Timmermann, and Valkanov (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 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

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

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

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

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-

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

augmented specifications. An example is the mixture distribution considered in Stock and Watson (2016) and Carriero, Clark, Marcellino, and Mertens (2022). 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

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

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

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

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 $\boldsymbol{\Omega} = \mathrm{diag}(\omega_1^2, \ldots, \omega_{p_{\mathbf{b}}}^2)$ and $\boldsymbol{\Xi} = \mathrm{diag}(\xi_1^2, \ldots, \xi_{p+1}^2)$, with the initial conditions $\mathbf{b}_1 \sim$

$\mathcal{N}(\mathbf{0}, \mathbf{V_b})$ and $\boldsymbol{\theta}_1 \sim \mathcal{N}(\mathbf{0}, \mathbf{V}_{\boldsymbol{\theta}})$. Similar to the time-invariant case, to separately identify β_t and $\boldsymbol{\theta}_t$, for $t = 1, \ldots, T$, 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

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

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

3.3 Irregularly Spaced Mixed-Frequency Data

In many MIDAS applications, such as those by Marcellino and Schumacher (2010), Kuzin, Marcellino, and Schumacher (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 $x_{t-k/m_t}^{(m)}, k = 0, \ldots, m_t - 1$. 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 $\mathbf{v}_k = [\phi_0(k), \ldots, \phi_p(k)]'$ is the vector of functional values of the basis functions $\phi_0(s), \ldots, \phi_p(s)$ 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 $x_{t-k/m_t}^{(m)}, k = 0, \ldots, m_t - 1$ to denote the m_t observations, but they are available at times $0 \leq s_{t,0} < s_{t,1} \cdots < s_{t,m_t-1} < 1$ 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 $s_{t,k} = k/m_t, k = 0, \ldots, m_t - 1$. 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 $\mathbf{v}_k = [\phi_0(k), \ldots, \phi_p(k)]'$ by $\mathbf{v}_{t,k} = [\phi_0(s_{t,k}), \ldots, \phi_p(s_{t,k})]'$.

3.4 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 $x_t^{(m_1)}, \ldots, x_t^{(m_n)}$, 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 $\mathbf{V}_j = [\mathbf{v}_{j,0}, \mathbf{v}_{j,1}, \ldots, \mathbf{v}_{j,K_j}]$ and $\mathbf{x}_t^{(m_j)} = [x_t^{(m_j)}, x_{t-1/m_j}^{(m_j)}, \ldots, x_{t-K_j/m_j}^{(m_j)}]'$,

the TVP-MIDAS model in (9) can be extended to include multiple high-frequency predictors:

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

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.

4 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 $\mathbf{b} = (\mathbf{b}_1', \ldots, \mathbf{b}_T')'$ and $\boldsymbol{\theta} = (\boldsymbol{\theta}_1', \ldots, \boldsymbol{\theta}_T')'$ and discuss efficient sampling from these posterior distributions.

We start with the conditional posterior distribution of $\mathbf{b}$. To that end, stack $\mathbf{y} = (y_{1+h}, \ldots, y_{T+h})'$ and $\boldsymbol{\epsilon} = (\epsilon_{1+h}, \ldots, \epsilon_{T+h})'$, 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 $\boldsymbol{\Sigma} = \mathrm{diag}(\lambda_1 \mathrm{e}^{g_1}, \ldots, \lambda_T \mathrm{e}^{g_T})$ and $\mathbf{X}_1 = \mathrm{diag}(\mathbf{x}_{\mathbf{b},1}', \ldots, \mathbf{x}_{\mathbf{b},T}')$ is a $T \times p_{\mathbf{b}}$ matrix whose t-th 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 equation (10) over $t = 1, \ldots, T$ yields

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

where $\mathbf{u}_1 = (\mathbf{u}_{1,1}', \ldots, \mathbf{u}_{1,T}')'$, $\mathbf{S}_1 = \mathrm{diag}(\mathbf{V}_{\mathbf{b}}, \boldsymbol{\Omega}, \ldots, \boldsymbol{\Omega})$, 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

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

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 $(\mathbf{b} \,|\, \mathbf{y}, \boldsymbol{\theta}, \boldsymbol{\Sigma}, \boldsymbol{\Omega})$ 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 $t = 1, \ldots, T$ to obtain $\tilde{\mathbf{y}} = (\tilde{y}_t, \ldots, \tilde{y}_T)'$. 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 $\mathbf{X}_2 = \operatorname{diag}(\beta_1 \mathbf{x}_1^{(m)\prime} \mathbf{V}', \ldots, \beta_T \mathbf{x}_T^{(m)\prime} \mathbf{V}')$. Furthermore, stacking the state equation (11) over $t = 1, \ldots, T$, 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 $\mathbf{u}_2 = (\mathbf{u}_{2,1}', \ldots, \mathbf{u}_{2,T}')'$ and $\mathbf{S}_2 = \operatorname{diag}(\mathbf{V}_{\boldsymbol{\theta}}, \boldsymbol{\Xi}, \ldots, \boldsymbol{\Xi})$ 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 $t = 1, \ldots, T$. 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 $\mathcal{S}$:

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

where

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

There are efficient algorithms that can be used to sample from $\mathcal{N}_{\mathcal{S}}(\boldsymbol{\mu}_{\boldsymbol{\theta}}, \mathbf{K}_{\boldsymbol{\theta}}^{-1})$, such as Algorithm 2.6 in Rue and Held (2005) and Algorithm 2 in Cong, Chen, and Zhou (2017). In particular, we can first sample $\tilde{\boldsymbol{\theta}} \sim \mathcal{N}(\boldsymbol{\mu}_{\boldsymbol{\theta}}, \mathbf{K}_{\boldsymbol{\theta}}^{-1})$ 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

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

Other steps of the posterior sampler are standard. For example, the log-volatility can be sampled using the auxiliary mixture sampler of Kim, Shephard, and Chib (1998), with the adjustment (for the latent variables $\lambda_1, \ldots, \lambda_T$) 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).

5 Empirical Application

5.1 Design of the Real-time Forecasting Exercise

To assess the performance of the proposed TVP-MIDAS model, we undertake a real-time forecasting (nowcasting) exercise focusing on quarterly US real GDP. This evaluation employs the real-time quarterly datasets of US real GDP sourced from the Philadelphia Federal Reserve Real-Time Datasets for Macroeconomists, spanning from 1990Q1 to 2021Q2. In our TVP-MIDAS model, we incorporate two high-frequency predictors for forecasting US real GDP: the daily interest rate spread, representing the slope of the yield curve (defined as the difference between the 10-year and 3-month treasury yields), and a weekly NFCI. Building upon the seminal work of Estrella and Hardouvelis (1991), empirical studies have consistently demonstrated the significant predictive capability of the yield curve slope in forecasting future US real GDP (Estrella, Rodrigues, and Schich

(2003); Rudebusch and Williams (2009)). Moreover, recent work by Adrian, Boyarchenko, and Giannone (2019) suggests that tightening financial conditions are associated with a notable increase in downside risk for US real GDP. Additionally, a recent study by Poon and Zhu (2024) underscores the importance of financial conditions as crucial predictors for forecasting recessions across various countries.

We acquire the daily interest rate spread data from the St. Louis FRED database. However, for the NFCI, weekly data vintages are exclusively accessible from the Archival Federal Reserve Economic Data (ALFRED) database starting from 2011. To compensate for this limitation, we utilize the weekly NFCI data compiled by Amburgey and McCracken (2023), which includes weekly data vintages from 1988 onwards. Our forecasting design aligns with that of Guérin and Marcellino (2013), who conduct a real-time forecasting exercise for quarterly US real GDP using a Markov-Switching MIDAS framework. Our approach involves generating forecasts for US real GDP one quarter ahead at the conclusion of each month of a quarter. Table 1 exemplifies our real-time forecast for US real GDP in 2000Q1. Notably, US real GDP data exhibits a one-quarter release delay. Consequently, when formulating a forecast at the conclusion of January 2000, our information set only encompasses daily and weekly predictors up to the conclusion of the first month of the quarter. Progressing to the end of February 2000, our information set expands to include information on daily and weekly predictors up to the conclusion of the second month of the quarter. By the conclusion of March 2000, our information set encompasses daily and weekly predictors for the entire quarter. Formally, we denote the forecasts 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 estimation sample for US real GDP spans from 1982Q1 to 1989Q4, with recursive expansion continuing until the end of 2021Q1. This temporal progression is mirrored in the timeframe for our daily and weekly predictors. The evaluation period for forecasting US Real GDP ranges from 1990Q1 to 2021Q2. Finally, we exclusively apply a data transformation to US real GDP, which involves multiplying the quarterly change in the natural logarithm of US real GDP by a factor of 400.

Table 1: Forecasting Scheme for 2000Q1

	Actual Observed Dates		
	January 2000	February 2000	March 2000
Real GDP data up to month	1999Q4	1999Q4	1999Q4
Daily Interest Rate data up to month	January 2000	February 2000	March 2000
Weekly NFCI data up to month	January 2000	February 2000	March 2000
Forecast Horizon	$h = 2/3$	$h = 1/3$	$h = 0$

5.2 Out-of-Sample Results

In this real-time forecasting exercise, we evaluate the predictive accuracy of six distinct MIDAS specifications in comparison to a simple AR(1) model. Our analysis encompasses various MIDAS configurations, including those with and without time-varying parameters, as well as different error specifications. Each of these six MIDAS specifications are outline upon in Table 2. Additionally, we explore the application of both Fourier and Almon lag polynomials to define the functional form of the MIDAS weights $\mathcal{B}$.

Table 2: Competing Models

Models	Information
AR(1)	Autoregressive model with one lag
MIDAS	MIDAS with time-invariant parameters and constant volatility
MIDAS-SV	MIDAS with time-invariant parameters and stochastic volatility
MIDAS-SVt	MIDAS with time-invariant parameters and fat-tail stochastic volatility errors
TVP-MIDAS	MIDAS with time-varying parameters and constant volatility
TVP-MIDAS-SV	MIDAS with time-varying parameters and stochastic volatility
TVP-MIDAS-SVt	MIDAS with time-varying parameters and fat-tail stochastic volatility errors

Table 3 presents the root mean squared forecast error (RMSFE) for six model specifications relative to a simple AR(1) model. Across all model specifications, the RMSFE values remain consistently similar across three forecast horizons, indicating superior point forecasting performance compared to the simple AR(1) model. Notably, there is no discernible enhancement in point forecasts when information is accumulated monthly within

each quarter, suggesting that the high-frequency nature of the predictors already contains sufficient informational content for accurately predicting US real GDP regardless of the timing within the quarter. Upon closer examination of specifications incorporating constant volatility, the TVP-MIDAS model demonstrates superior point forecast accuracy compared to its time-invariant counterpart in both Fourier and Almon Lag cases. This observation suggests that the inclusion of time-varying parameters within a static MIDAS framework indeed leads to improved point forecasting accuracy. Furthermore, the addition of SV within the model specification also contributes to enhanced point forecasts. However, the point forecasts for both MIDAS-SV and TVP-MIDAS-SV models exhibit minimal divergence, indicating that the inclusion of time-varying parameters offers marginal improvement when SV is incorporated into the specification. Finally, it is noteworthy that there is no observable improvement in forecast accuracy from incorporating fat-tails within the SV specification.

Table 3: The RMSFE for the six models benchmarked against AR(1) model

Forecast Horizon	MIDAS	MIDAS-SV	MIDAS-SV-t	TVP-MIDAS	TVP-MIDAS-SV	TVP-MIDAS-SV-t
			Fourier Series			
$h = 2/3$	0.90	0.80	0.80	0.83	0.83	0.81
$h = 1/3$	0.90	0.81	0.80	0.86	0.84	0.82
$h = 0$	0.90	0.81	0.80	0.84	0.83	0.80
Average	0.90	0.80	0.80	0.85	0.83	0.81
			Almon Lag			
$h = 2/3$	0.94	0.82	0.82	0.87	0.87	0.84
$h = 1/3$	0.94	0.82	0.82	0.87	0.87	0.84
$h = 0$	0.94	0.82	0.82	0.87	0.86	0.90
Average	0.94	0.82	0.82	0.87	0.87	0.86

Notes: *,**,*** denotes the 10, 5, and 1 percent significant level of the Diebold-Mariano predictability test.

We evaluate the performance of density forecasts using the Continuous Ranked Probability Score (CRPS) across six model specifications. Table 4 presents the average CRPS values relative to the AR(1) model for each specification, revealing consistent CRPS values across the three forecast horizons. Notably, among the six specifications, the TVP-MIDAS model emerges as the only one yielding density forecasts inferior to the simple

AR(1) model. Conversely, incorporating SV within both the time-invariant and TVP-MIDAS frameworks enhances density forecast precision. However, overall, the MIDAS specifications with SV yield comparable CRPS values, suggesting that the inclusion of time-varying parameters in the MIDAS framework does not significantly enhance density forecast accuracy relative to the time-invariant case.

Given the significance of the proposed framework's high-frequency predictors in forecasting and identifying future recessions in empirical studies, we delve deeper into the forecasting performance of our six model specifications concerning the left and right tails of the density forecasts. Drawing from the methodology outlined by Gneiting and Ranjan (2011), we calculate the predictive quantile score for a given quantile τ, expressed as:

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

Here $QS_{\tau,t}$ represents the predictive quantile for US Real GDP, where $\mathbb{I}\{y_t \leq Q_{\tau,t}\}$ takes the value 1 if the realized value is at or below the predictive quantile and 0 otherwise. We assess the performance of the quantile score in both left (10%) and right (90%) tails by setting $\tau = 0.1$ and $\tau = 0.9$, respectively. We present the average quantile scores for the left and right tails for each model relative to the AR(1) model in Table 5. Significantly, our analysis reveals that the time-varying parameter MIDAS model outperforms the time-invariant case in forecasting the left tail, particularly evident in the Almon Lag case. Conversely, for right tail risk, the static MIDAS model specifications exhibit superior performance compared to the time-varying case. Both these results are visually illustrated in Figure 1, which depicts the rolling average of CRPS and quantile scores for $h = 0$ over time.

Table 4: The Average CRPS for the six models benchmarked against AR(1) model

Forecast Horizon	MIDAS	MIDAS-SV	MIDAS-SV-t	TVP-MIDAS	TVP-MIDAS-SV	TVP-MIDAS-SV-t
			Fourier Series			
$h = 2/3$	0.96	0.91	0.91	1.06	0.95	0.94
$h = 1/3$	0.97	0.92	0.91	1.11	0.95	0.94
$h = 0$	0.97	0.92	0.92	1.08	0.95	0.92
Average	0.97	0.92	0.91	1.08	0.95	0.93
			Almon Lag			
$h = 2/3$	0.99	0.93	0.93	1.03	0.95	0.94
$h = 1/3$	0.99	0.93	0.93	1.04	0.96	0.96
$h = 0$	0.99	0.93	0.93	1.04	0.96	0.97
Average	0.99	0.93	0.93	1.04	0.96	0.96

Notes: *,**,*** denotes the 10, 5, and 1 percent significant level of the Diebold-Mariano predictability test.

Table 5: The Average Quantile Scores for the left (10%) and right (90%) tail for the six models relative to the AR(1) model

Forecast Horizon	MIDAS	MIDAS-SV	MIDAS-SV-t	TVP-MIDAS	TVP-MIDAS-SV	TVP-MIDAS-SV-t
10 percent Quantile Scores (Left Tail)						
Fourier Series						
$h = 2/3$	0.98	1.04***	1.04***	1.01***	0.93***	0.93***
$h = 1/3$	0.99	1.05*	1.04*	1.00**	0.93***	0.93***
$h = 0$	0.99	1.05	1.04	0.99***	0.93***	0.93***
Average	0.99	1.05	1.04	1.00	0.93	0.93
Almon Lag						
$h = 2/3$	0.98**	1.05***	1.05***	0.91***	0.92***	0.92***
$h = 1/3$	0.98**	1.05***	1.04***	0.92***	0.93***	0.92***
$h = 0$	0.98***	1.05***	1.04***	0.92***	0.93***	0.92***
Average	0.98	1.05	1.04	0.92	0.93	0.92
90 percent Quantile Scores (Right Tail)						
Fourier Series						
$h = 2/3$	0.92***	0.92*	0.91*	1.12	1.07	1.07
$h = 1/3$	0.92***	0.92**	0.91**	1.19	1.08	1.07
$h = 0$	0.92***	0.92	0.92	1.14	1.08	1.06
Average	0.92	0.92	0.91	1.15	1.08	1.06
Almon Lag						
$h = 2/3$	0.94***	0.94**	0.92**	1.13	1.12	1.10
$h = 1/3$	0.94***	0.93**	0.93**	1.13	1.12	1.10
$h = 0$	0.95***	0.94	0.93	1.13	1.12	1.13
Average	0.94	0.94	0.93	1.13	1.12	1.11

Notes: *,**,*** denotes the 10, 5, and 1 percent significant level of the Diebold-Mariano predictability test.

Consequently, our results suggest that during periods of heightened volatility, the incorporation of time-varying parameters and SV is essential for accurately forecasting the left tail of US real GDP or recessionary events. This conclusion aligns with the findings of Adrian, Boyarchenko, and Giannone (2019) and Estrella and Hardouvelis (1991), who

emphasize the importance of financial conditions and the slope of the yield curve as predictors for future recessions in the economy. In contrast, during periods of tranquility, a static time-invariant MIDAS model suffices for forecasting the right tail of US real GDP or positive economic growth.

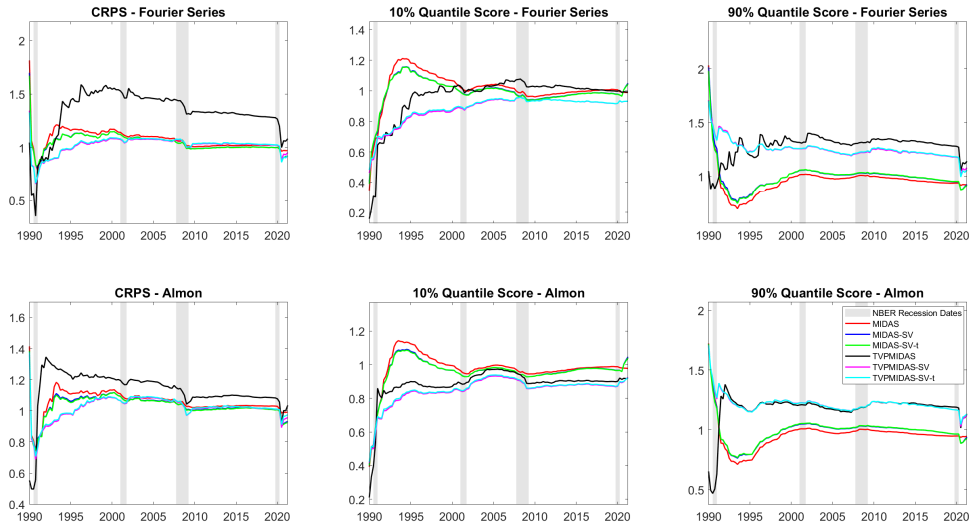

Figure 1: Plot of the Rolling Average CRPS and Quantile Scores for $h = 0$ across the six models, relative to the AR(1) model, over the forecast evaluation period.

5.3 Time-Varying Impact of NFCI and Slope of the Yield Curve

The proposed TVP-MIDAS framework offers a notable advantage by facilitating the evaluation of the impact of the NFCI and the slope of the yield curve on US real GDP across various time periods. Specifically, Figure 2 displays the posterior estimates of the time-varying parameters derived from the TVP-MIDAS-SV model with the Fourier expansion method, which showcases superior forecasting performance compared to alternative model specifications in the preceding real-time forecasting section. These estimates are derived from the final data vintage of 2021Q2, spanning from 1982Q2 to 2021Q1.

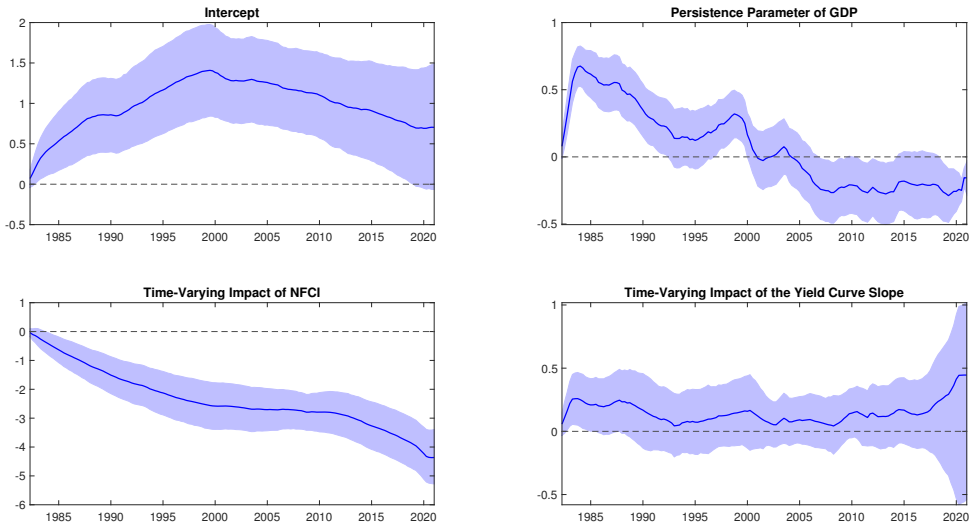

Notes: The thick blue line represents the posterior mean estimates of the time-varying parameters derived from the TVP-MIDAS-SV model. The shaded blue area indicates the corresponding 68 percent credible interval.

Figure 2: Posterior estimates of the Time-Varying Parameters from the TVP-MIDAS-SV model

The top panel of Figure 2 illustrates the time-varying intercept and autoregressive persistence parameter of US real GDP. Remarkably, the persistence parameter of US real GDP exhibits a declining trend over time, indicating a diminishing significance of past US real GDP values in forecasting future trends. In the second panel of Figure 2, we observe the evolving impact of the NFCI and the slope of the yield curve over time. Significantly, the impact of NFCI demonstrates a notable negative trend, suggesting an increasingly negative correlation between US real GDP and NFCI over time, consistent with the findings of Adrian, Boyarchenko, and Giannone (2019). However, the impact of the yield curve slope appears to have minimal effect on US Real GDP over time, as evidenced by the majority of uncertainty bands encompassing zero for much of the observed period.

6 Conclusion

In this study, we introduce a novel TVP-MIDAS framework developed with precision-based methods. Through a comprehensive assessment, we evaluate the effectiveness of the

proposed framework in real-time forecasting applications for US real GDP. Leveraging two high-frequency predictors, the daily interest rate spread and a weekly NFCI, our analysis demonstrates that TVP-MIDAS model specifications incorporating stochastic volatility consistently outperform their static counterparts. Specifically, our findings reveal that our TVP-MIDAS framework yields superior forecasts, particularly in capturing the left tail risk of US real GDP. Finally, our in-sample analysis unveils a significant negative trend in the influence of the NFCI, implying a progressively adverse correlation between US real GDP and NFCI over time.

References

Adrian, T., N. Boyarchenko, and D. Giannone (2019): "Vulnerable growth," *American Economic Review*, 109(4), 1263–1289.

Amburgey, A. J., and M. W. McCracken (2023): "On the real-time predictive content of financial condition indices for growth," *Journal of Applied Econometrics*, 38(2), 137–163.

Andreou, E. (2016): "On the use of high frequency measures of volatility in MIDAS regressions," *Journal of econometrics*, 193(2), 367–389.

Barnett, A., H. Mumtaz, and K. Theodoridis (2014): "Forecasting UK GDP growth and inflation under structural change. A comparison of models with time-varying parameters," *International Journal of Forecasting*, 30(1), 129–143.

Bobeica, E., and B. Hartwig (2023): "The COVID-19 shock and challenges for inflation modelling," *International Journal of Forecasting*, 39(1), 519–539.

Carriero, A., T. E. Clark, and M. Marcellino (2015): "Realtime nowcasting with a Bayesian mixed frequency model with stochastic volatility," *Journal of the Royal Statistical Society Series A: Statistics in Society*, 178(4), 837–862.

Carriero, A., T. E. Clark, M. Marcellino, and E. Mertens (2022): "Addressing COVID-19 outliers in BVARs with stochastic volatility," *The Review of Economics and Statistics*, Forthcoming.

Chan, J. C., and E. Eisenstat (2018): "Bayesian model comparison for time-varying parameter VARs with stochastic volatility," *Journal of applied econometrics*, 33(4), 509–532.

Chan, J. C., and C. Y. Hsiao (2014): "Estimation of stochastic volatility models with heavy tails and serial dependence," *Bayesian inference in the social sciences*, pp. 155–176.

Chan, J. C., and I. Jeliazkov (2009): "Efficient simulation and integrated likelihood estimation in state space models," *International Journal of Mathematical Modelling and Numerical Optimisation*, 1(1-2), 101–120.

Clark, T. E. (2011): "Real-time density forecasts from Bayesian vector autoregressions with stochastic volatility," *Journal of Business and Economic Statistics*, 29(3), 327–341.

Clark, T. E., and F. Ravazzolo (2015): "Macroeconomic forecasting performance under alternative specifications of time-varying volatility," *Journal of Applied Econometrics*, 30(4), 551–575.

Cong, Y., B. Chen, and M. Zhou (2017): "Fast simulation of hyperplane-truncated multivariate normal distributions," *Bayesian Analysis*, 12(4), 1017–1037.

Cross, J., C. Hou, G. Koop, and A. Poon (2023): "Large stochastic volatility in mean VARs," *Journal of Econometrics*, 236(1), 105469.

Cross, J., and A. Poon (2016): "Forecasting structural change and fat-tailed events in Australian macroeconomic variables," *Economic Modelling*, 58, 34–51.

D'Agostino, A., L. Gambetti, and D. Giannone (2013): "Macroeconomic forecasting and structural change," *Journal of applied econometrics*, 28(1), 82–101.

Estrella, A., and G. A. Hardouvelis (1991): "The term structure as a predictor of real economic activity," *The journal of Finance*, 46(2), 555–576.

Estrella, A., A. P. Rodrigues, and S. Schich (2003): "How stable is the predictive power of the yield curve? Evidence from Germany and the United States," *Review of Economics and Statistics*, 85(3), 629–644.

Foroni, C., and M. Marcellino (2014): "A comparison of mixed frequency approaches for nowcasting Euro area macroeconomic aggregates," *International Journal of Forecasting*, 30(3), 554–568.

Ghysels, E., P. Santa-Clara, and R. Valkanov (2005): "There is a risk-return trade-off after all," *Journal of financial economics*, 76(3), 509–548.

Ghysels, E., A. Sinko, and R. Valkanov (2007): "MIDAS regressions: Further results and new directions," *Econometric reviews*, 26(1), 53–90.

Gneiting, T., and R. Ranjan (2011): "Comparing density forecasts using threshold- and quantile-weighted scoring rules," *Journal of Business & Economic Statistics*, 29(3), 411–422.

Guérin, P., and M. Marcellino (2013): "Markov-switching MIDAS models," *Journal of Business & Economic Statistics*, 31(1), 45–56.

Kim, S., N. Shephard, and S. Chib (1998): "Stochastic Volatility: Likelihood Inference and Comparison with ARCH Models," *Review of Economic Studies*, 65(3), 361–393.

Koop, G., and D. Korobilis (2013): "Large time-varying parameter VARs," *Journal of Econometrics*, 177(2), 185–198.

Kuzin, V., M. Marcellino, and C. Schumacher (2011): "MIDAS vs. mixed-frequency VAR: Nowcasting GDP in the euro area," *International Journal of Forecasting*, 27(2), 529–542.

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

Mogliani, M., and A. Simoni (2021): "Bayesian MIDAS penalized regressions: estimation, selection, and prediction," *Journal of Econometrics*, 222(1), 833–860.

Pettenuzzo, D., A. Timmermann, and R. Valkanov (2016): "A MIDAS approach to modeling first and second moment dynamics," *Journal of Econometrics*, 193(2), 315–334.

Poon, A., and D. Zhu (2024): "Do Recessions and Bear Markets Occur Concurrently Across Countries? A Multinomial Logistic Approach," *Journal of Financial Econometrics*.

Potjagailo, G., and D. Kohns (2023): "Flexible Bayesian MIDAS: time-variation, group-shrinkage and sparsity," .

Primiceri, G. E. (2005): "Time varying structural vector autoregressions and monetary policy," *The Review of Economic Studies*, 72(3), 821–852.

Rudebusch, G. D., and J. C. Williams (2009): "Forecasting recessions: The puzzle of the enduring power of the yield curve," *Journal of Business & Economic Statistics*, 27(4), 492–503.

Rue, H., and L. Held (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," .

Stock, J. H., and M. W. Watson (2016): "Core inflation and trend inflation," *Review of Economics and Statistics*, 98(4), 770–784.

Wang, L., F. Ma, J. Liu, and L. Yang (2020): "Forecasting stock price volatility: New evidence from the GARCH-MIDAS model," *International Journal of Forecasting*, 36(2), 684–694.