Received: 13 October 2020 | Revised: 10 August 2022 | Accepted: 11 August 2022

DOI: 10.1002/jae.2931

WILEY Journal of APPLIED ECONOMETRICS

RESEARCH ARTICLE

Forecasting low-frequency macroeconomic events with high-frequency data

Ana Beatriz Galvão[1] | **Michael Owyang**[2]

[1]Warwick Business School, University of Warwick, Coventry, UK

[2]Research Division, Federal Reserve Bank of St. Louis, St. Louis, Missouri, USA

Correspondence
Ana Beatriz Galvão, Warwick Business School, University of Warwick, Scarman Rd, Coventry CV4 7AL, UK.
Email: ana.galvao@wbs.ac.uk

Summary

High-frequency financial and economic indicators are usually time-aggregated before computing forecasts of macroeconomic events, such as recessions. We propose a mixed-frequency alternative that delivers high-frequency probability forecasts (including their confidence bands) for low-frequency events. The new approach is compared with single-frequency alternatives using loss functions for rare-event forecasting. We find (i) the weekly-sampled term spread improves over the monthly-sampled to predict NBER recessions, (ii) the predictive content of financial variables is supplementary to economic activity for forecasts of vulnerability events, and (iii) a weekly activity index can date the 2020 business cycle peak in real-time using a mixed-frequency filtering.

KEYWORDS

event probability forecasting, financial indicators, mixed frequency models, recession, weekly activity index

1 | INTRODUCTION

One way that forecasting helps decision making is by supplying predicted probabilities of critical future events. Examples of great consequence macroeconomic events include recessions (Bauer & Mertens, 2018; Chauvet & Potter, 2005; Estrella & Mishkin, 1998; Liu & Moench, 2016), sovereign defaults (Freitag, 2014; Manasse & Roubini, 2009), and periods of vulnerable growth (Adrian et al., 2019). Financial variables, which are often used as predictors for these events, are sampled at a higher frequency (daily) than most economic indicators and events (monthly, quarterly) and, consequently, are time-aggregated before they are used to estimate the forecast model.[1] In this paper, we propose a mixed-frequency strategy to exploit the predictive content of high-frequency economic and financial variables for low-frequency binary events.[2]

MIxed **DA**ta **S**ampling (MIDAS) regressions are typically estimated using (nonlinear) least squares (Clements & Galvão, 2008; Ghysels et al., 2007, 2006). For binary dependent variables, likelihood-based estimators are generally applied (Audrino et al., 2019; Freitag, 2014). Bayesian methods have been developed to accommodate MIDAS in both the

[1]Andreou et al. (2013), Galvão (2013), and Pettenuzzo et al. (2016) are exceptions where time aggregation is not applied to financial variables to predict economic activity.

[2]The modeling strategy is easily extended to any event with discrete outcomes. While our applications are primarily forecasting experiment, one could apply the basic principles in this model to measure dynamic causal effects on event probabilities from exogenous changes in a variable of interest sampled at a higher frequency than the observed event.

conditional mean and the conditional variance (Pettenuzzo et al., 2016) and to deal with a large number of predictors (Mogliani & Simoni, 2021). Our Bayesian estimation strategy combines the Gibbs sampler developed for probit models (Albert & Chib, 1993) with a Metropolis step to draw the parameters of the beta function that parsimoniously describes the aggregating weights.[3] Our choice of a beta function to aggregate high-frequency data accommodates applications with large numbers of leads and lags of the high-frequency variable and is compatible with cases where the binary dependent variable is available at quarterly frequency and the regressor at the daily frequency as in Galvão (2013) and Ghysels et al. (2019). Our approach differs from the use of Almon lags in Pettenuzzo et al. (2016) and the use of unrestricted weighting schemes in Carriero et al. (2015).

The MIDAS-probit model is applied to answer three empirical research questions regarding the use of high-frequency variables to predict low-frequency events. First, we revisit the predictive content of the spread between long-term and short-term interest rates sampled daily/weekly for US recession phases sampled monthly (Bauer & Mertens, 2018; Chauvet & Potter, 2005; Liu & Moench, 2016).[4] Second, we evaluate whether financial variables sampled weekly have additional predictive ability over economic activity variables sampled monthly/quarterly for vulnerable gross domestic product (GDP) periods with elevated downside risks as in Adrian et al. (2019) and Plagborg-Møller et al. (2020). As an alternative, we evaluate the high-frequency financial variables' predictive content by comparing the accuracy of MIDAS-probit models with the Survey of Professional Forecasters (SPF) consensus predictions for the probability of negative GDP growth. Third, we assess whether we could have anticipated the June 2020 NBER announcement of a peak in February 2020 using the weekly-sampled economic index, proposed by Lewis et al. (2020).

An important characteristic of mixed-frequency models is that forecasts for low-frequency events can be updated each time new observations of the high-frequency data become available. We evaluate whether these updates improve forecasting performance and find improved accuracy when using the weekly-sampled spread to predict quarterly GDP contractions.

The forecasts are evaluated by the relative losses incurred by false negatives versus false positives using three loss functions. The first two loss functions—the area under the receiver operating characteristic curve (AUROC; as applied in Berge & Jorda, 2011, and Liu & Moench, 2016), and the diagonal of the elementary score (Bouallègue et al., 2018)—measure the forecasting model's ability to accurately classify the event. We also consider a logarithm score, more in line with the quantile forecasting evaluation and the predictive scores in Adrian et al. (2019) and Plagborg-Møller et al. (2020). We compare the MIDAS-probit to simple forecasting rules. As in the climatology literature (Bouallègue et al., 2019), we measure the relative performance of our model with respect to the unconditional event probability using skill scores.

We describe our mixed-frequency approach to model binary dependent variables in Section 2. Details of our econometric implementation, including the Bayesian estimation strategy and probability forecasting loss functions are described in Section 3. Section 4 provides empirical results, discussion of our three empirical applications, and a comparison of our results with SPF probability forecasts of negative growth. Section 5 concludes with a summary assessment of the proposed forecasting model and implications of the empirical results for macroeconomic forecasting.

2 | THE MIDAS-PROBIT

2.1 | Setup

Define a low-frequency binary variable, $S_t = \{0, 1\}$, where $S_t = 1$ indicates that the economic event of interest (e.g., a recession phase) occurs. Define a latent variable, y_t^*, such that $y_t^* \geq 0$ if $S_t = 1$ and $y_t^* < 0$ if $S_t = 0$. Then, a single-regressor probit model for h-period-ahead event prediction is:

$$\begin{aligned} \Pr[S_{t+h} = 1|\Omega_t] &= \Pr[y_{t+h}^* \geq 0|\Omega_t] \\ &= \Phi(\beta_{0,h} + \beta_{1,h} z_t), \end{aligned} \tag{1}$$

[3]Similar algorithms have been proposed by Casarin et al. (2018) and Foroni et al. (2019)

[4]There are a number of reasons that one might want to forecast recessions rather than output growth. Recessions are large, rare and critical events that may produce nonlinearities in macroeconomic relationships. For example, recessions has been shown to alter the effect of fiscal policy on the macroeconomy (Auerback & Gorodnichenko, 2012).

where $\Phi(.)$ is the cumulative density function (CDF) of a standard normal density and the information set, Ω_t, consists only of a single monthly variable, z_t.[5] Direct multi-step-ahead monthly forecasts can be obtained by estimating (1) for $h = 1, \ldots, H$ by changing the indicator variable as $S_{t+1}, \ldots, S_{t+H}$, implying that parameters $\left(\beta_{0,h}, \beta_{1,h}\right)$ change with the horizon.

While events of interest are observed either monthly or quarterly, many predictors—for example, financial variables—are available at a daily or higher frequency.[6] The standard probit in Equation (1) requires the aggregation of high-frequency regressors to match the sampling of S_t. For example, data sampled daily can either be averaged over the month or represented by the last day of the month to match a binary variable S_t observed monthly. This temporal aggregation assumes that the high-frequency timing of innovations to the predictors are unimportant and can result in the loss of this information. In the next few sections, we propose an alternative specification of a mixed-frequency probit that uses high-frequency data directly, preserving information about the high-frequency data.

2.2 | The single-predictor MIDAS-probit

For expositional simplicity, we continue to characterize the one-predictor case; the generalization to multiple predictors at possibly multiple frequencies appears below. Assume that the high-frequency variable is sampled m times more frequently than S_t (and, consequently, y_t^*); for example, if the low-frequency variable is monthly and the high-frequency predictor is daily, $m = 21$ trading days (on average over the year). Thus, for each realization of S_t, the information set, Ω_t, also includes high-frequency observations within the low-frequency period: $z_t^{(m)}, z_{t-\frac{1}{m}}^{(m)}, z_{t-\frac{2}{m}}^{(m)}, \ldots, z_{t-\frac{m-1}{m}}^{(m)}$.[7]

We can preserve the probit-type formulation for the h-period-ahead direct multi-step forecast by writing:

$$\Pr[y_{t+h}^* \geq 0|\Omega_t] = \Phi\left(\beta_{0,h} + \beta_{1,h}\sum_{k=1}^{K}\varpi\left(k;\boldsymbol{\theta}_h\right)z_{t-\frac{k-1}{m}}^{(m)}\right), \tag{2}$$

where K is number of lags at the sampling frequency of $z_t^{(m)}$. Thus, Ω_t may include high-frequency lags over a number of past low-frequency periods. The functions, $\varpi\left(k;\boldsymbol{\theta}_h\right)$, attributing weights to each of the (high-frequency) lags of $z_t^{(m)}$ up to K, are often referred to as MIDAS functions (Andreou et al., 2010; Clements & Galvão, 2008; Ghysels et al., 2006, 2007; Kuzin et al., 2013). If the number of high-frequency lags, K, is equal to m, then $\sum_{k=1}^{K}\varpi\left(k;\boldsymbol{\theta}_h\right)z_{t-(k-1)/m}^{(m)}$ is the predictor, aggregated to the frequency of the binary dependent variable, S_t. If $\varpi\left(k;\boldsymbol{\theta}_h\right) = 1/m$, the aggregation scheme would average the high-frequency data (e.g., daily) within the low-frequency period (e.g., a month).

To identify the slope parameter, $\beta_{1,h}$, the weights are constrained to sum to 1:

$$\varpi\left(j;\boldsymbol{\theta}_h\right) = \frac{f\left(k,\boldsymbol{\theta}_h\right)}{\sum_{k=1}^{K}f\left(k,\boldsymbol{\theta}_h\right)}.$$

While the weights can, in principle, take on a number of functional forms, we employ a beta function:

$$f\left(k,\boldsymbol{\theta}_h\right) = \frac{\kappa^{\theta_1-1}(1-\kappa)^{\theta_2-1}\Gamma\left(\theta_{1,h}+\theta_{2,h}\right)}{\Gamma\left(\theta_{1,h}\right)\Gamma\left(\theta_{2,h}\right)}; \kappa = k/(K+1), \tag{3}$$

where $\Gamma(.)$ is a gamma function and $\boldsymbol{\theta}_h$ is a vector of the two parameters that govern the shape of the weighting function (Andreou et al., 2010; Ghysels et al., 2007). The beta function with $\theta_{1,h} = \theta_{2,h} = 1$ collapses to the equal-weighting aggregation scheme if applied within a period ($K = m$). In typical empirical MIDAS applications, however, the number

[5]The model is easily generalized to include lags of the predictor and additional regressors. For example, Liu and Moench (2016) suggest adding both z_t and z_{t-6} to predict recessions using the term-structure spread.

[6]Estrella and Mishkin (1998) demonstrate the predictive ability of the slope of the yield curve for US recessions using probit models, and Chauvet and Potter (2005), Kauppi and Saikkonen (2008), and Liu and Moench (2016) provide evidence of the predictive ability of the spread between the 10-year Treasury bond yield and the 3-month Treasury Bill rate.

[7]Notice that the integer timing t is still in the low-frequency variable. The high-frequency variable is observed in fractions of the low-frequency periods. Thus, $t - 1/m$ indexes one high-frequency period before the tth observation of the low-frequency variable.

of high-frequency lags of the predictor is set such that $K \geq m$ (Carriero et al., 2020; Galvão, 2013; Pettenuzzo et al., 2016). If $\theta_1 < \theta_2$, weights are decreasing across lags, so less weight is assigned to values further in the past.

The MIDAS approach is a parsimonious method of employing many lags of the high-frequency variable, as we only need to estimate three parameters $(\beta_{1,h}, \theta_{1,h}, \theta_{2,h})$ for each predictor. In contrast, Foroni et al. (2015) argue to use the UMIDAS specification that estimates the coefficient for each lag separately. Most applications of UMIDAS are for small differences in frequency ($m = 3$) as, in that case, parameter proliferation is limited (Foroni et al., 2015), or for cases where Bayesian methods are employed to deal with the large number of parameters (as, e.g., Carriero et al., 2020).

2.3 | The general MIDAS-probit model

The model introduced above extends to a mix of any number of variables sampled at multiple frequencies under the assumption that the binary dependent variable is sampled at the lowest common frequency. The general form of the MIDAS-probit can be estimated using an $(N + 1) \times 1$ vector of predictors that includes a constant and N same- or higher-frequency regressors.

Define $Z_{nt}(K_n, \boldsymbol{\theta}_n)$ as the weighted sum of K_n lags of the high-frequency variable, $z_{nt}^{(m)}$, using a beta weighting function with parameters $\boldsymbol{\theta}_n$—that is, $Z_{nt}(K_n, \boldsymbol{\theta}_n) = \sum_{k=1}^{K_n} \varpi(k; \boldsymbol{\theta}_i)\, z_{n,t-(k-1)/m}^{(m)}$. Then, define

$$\mathbf{Z}_t(\Theta) = [1, Z_{1t}(K_1, \boldsymbol{\theta}_1), \ldots, Z_{Nt}(K_N, \boldsymbol{\theta}_N)]', \tag{4}$$

where $\Theta = (\boldsymbol{\theta}_1', \ldots, \boldsymbol{\theta}_N')$. The general form of the MIDAS-probit is

$$\Pr\left[y_{t+h}^* \geq 0 | \Omega_t\right] = \Phi[\mathbf{Z}_t(\Theta_h)' \boldsymbol{\beta}_h], \tag{5}$$

where $y_{t+h}^* \geq 0$ if $S_{t+h} = 1$, $y_{t+h}^* < 0$ if $S_{t+h} = 0$ for $t = 1, \ldots, T-h$, $\boldsymbol{\beta}_h$ is a $(N+1)\times 1$ vector of slopes, and the parameters are indexed by horizon to produce a direct multi-step forecast.

Note that the MIDAS-probit can be written as a regression model:

$$y_{t+h}^* = \mathbf{Z}_t(\Theta_h)' \boldsymbol{\beta}_h + u_{t+h}, \tag{6}$$

where $u_{t+h} \sim N(0, 1)$. Changing the distributional assumption for u_{t+h} can change the model from probit to logit, and so on.

3 | ECONOMETRIC IMPLEMENTATION

In this section, we describe the steps required to estimate the model. We also outline how we form and evaluate the forecasts.

3.1 | Estimation

MIDAS models are often estimated using nonlinear least squares as in Ghysels et al. (2006) and Clements and Galvão (2008) or—in the case of binary dependent variables—by maximum likelihood (Audrino et al., 2019). As these techniques require numerical optimization, algorithms have been proposed to improve convergence (Ghysels & Qian, 2019). When MIDAS regressions are estimated using Bayesian methods, weighting schemes are often chosen to obtain linearity in the parameters—for example, the Almon weighting function (Pettenuzzo et al., 2016) and the unrestricted function (Carriero et al., 2020).[8]

We estimate the model using a Gibbs sampler (Casella & George, 1992; Gelfand & Smith, 1990) with a Metropolis-in-Gibbs step (Chib & Greenberg, 1995) to sample the MIDAS hyperparameters that govern the weights. As described in

[8]An advantage of the Bayesian estimation is that we are able to fully characterize uncertainty in the MIDAS weights and slopes, facilitating inference regarding the relevance of the use of mixed frequencies.

Greenberg (2013, chapter 8), Gibbs sampling is the standard method to estimate probit models using Bayesian methods and should be equivalent to maximum likelihood under some conditions.

To start, we require a prior for the slopes, $\boldsymbol{\beta}_h$, and the MIDAS hyperparameters, Θ_h. Conditional on known MIDAS weights, the latent variable in the probit can be written as a linear regression. Thus, we assume a standard conjugate zero-mean normal prior on the slope coefficients. We further assume that the MIDAS hyperparameters have (restricted) Gamma priors and are centered around the belief that the high-frequency data is equally weighted. For most of our specifications, we impose a joint restriction across each predictor's MIDAS hyperparameters, $\theta_{2,n,h} \geq \theta_{1,n,h}$, so that older data are given less weight. Because the model is based on the standard probit, the variance of the latent variable is assumed to be fixed at 1.

Assume that the estimation period is up to $t = \tau$. The sampler is decomposed into three blocks: (i) the slope parameters, $\boldsymbol{\beta}_h|\Theta_h, \left\{y^*_{t+h}\right\}_{t=1}^{\tau-h}$; (ii) the MIDAS weight hyperparameters, $\Theta_h|\boldsymbol{\beta}_h, \left\{y^*_{t+h}\right\}_{t=1}^{\tau-h}$; and (iii) the latent variable, $y^*_{t+h}|\boldsymbol{\beta}_h, \Theta_h$ for $t = 1, \ldots, \tau - h$. As mentioned above, the slope parameters have conjugate normal posteriors. The MIDAS hyperparameters are drawn via the MH-in-Gibbs step, assuming a Gamma proposal density truncated so that the aforementioned restriction, $\theta_{2,n,h} \geq \theta_{1,n,h}$, holds. The latent variable is drawn sequentially from independent truncated normal densities, where the direction of the truncation depends on the value of the observed binary indicator.[9]

We compute intervals for the out-of-sample event probabilities, $\Pr[y^*_{\tau+h} \geq 0|\Omega_\tau]$, for each low-frequency forecast origin, $\tau = L+1, \ldots, T-h$ (where L is the number of observations in the initial in-sample period) as follows: We use the sampler draws of $\boldsymbol{\beta}_h$ and Θ_h obtained as described earlier with data up to τ to compute a direct forecast of $\Pr[y^*_{\tau+h} \geq 0|\Omega_\tau]$ at each draw of the sampler. Then, for each τ, we compute the predicted probability as the mean and use quantiles (the 16^{th} and 84^{th} quantiles) to compute intervals. For in-sample predicted probabilities with data up to τ, we use the sampler draws as before but compute predictions by direct forecasting for $\Pr[y^*_{t+h} \geq 0|\Omega_t]$, where $t = 1, \ldots, \tau - h$.

The sampler consists of 10,000 draws: 5000 burn draws and 5000 to form the posterior distributions.[10]

3.2 | Forecasting low-frequency variables using high-frequency predictors

In the previous section, we assumed that the forecast origin coincided with the observation of the low-frequency variable—that is, for each $\tau = L+1, \ldots, T-h$ (at low-frequency), we observe the event indicator, S, and the conditioning information set, Ω, up to τ:

$$\Omega_\tau = \left\{z^{(m)}_\tau, z^{(m)}_{\tau-1/m}, \ldots, z^{(m)}_{\tau-(m-1)/m}, z^{(m)}_{\tau-1}, z^{(m)}_{\tau-1-(1/m)}, \ldots, z^{(m)}_{\tau-K/m}\right\},$$

if $K > m$. An advantage of MIDAS models is that forecasts can be updated between observations of the dependent variable, at the highest frequency available.[11]

We apply a strategy based on the approach in Ghysels et al. (2019): We generate a forecast at each intra-period observation of the high-frequency variable using the model parameters estimated up to τ.[12] Define the intra-period information set at a forecast origin j high-frequency periods after τ, available during current quarter $\tau + 1$:

$$\Omega^{[j]}_\tau = \left\{z^{(m)}_{\tau+\frac{j}{m}}, \ldots, z^{(m)}_{\tau+\frac{1}{m}}, z^{(m)}_\tau, \ldots, z^{(m)}_{\tau-\frac{K-j}{m}}\right\},$$

where $j = 1, \ldots, m-1$. We then generate a sequence of probabilistic forecasts, $\Pr[y^*_{\tau+h} \geq 0|\Omega^{[j]}_\tau]$, for $j = 1, \ldots, m-1$, all targeting the outcome $S_{\tau+h}$. For each sampler draw, high-frequency predictions are computed between two low-frequency periods, τ and $\tau + 1$, as follows:

[9] Appendix A describes the algorithm and the harmonic mean estimator used to compute the marginal likelihood. The approach employs the Chib (1995) estimator for the probit marginal likelihood as an input. Estimates of the marginal likelihood are used to compare MIDAS-probit specifications with different numbers of lags and weighting function specifications.

[10] See Appendix S1 for convergence analysis.

[11] Clements and Galvão (2008) show how MIDAS regressions can employ intra-quarter monthly data to nowcast quarterly GDP growth. They estimate separate MIDAS regression models for each monthly horizon.

[12] We compare this approach with the alternative of re-estimating the forecasting model every high-frequency period in Appendix S2.

$$\Pr\left[y^*_{\tau+h} \geq 0|\Omega^{[j]}_{\tau}\right] = \Phi\left(\beta_{0,h} + \beta_{1,h}\sum_{k=1}^{K}\varpi\left(k;\theta_h\right)z^{(m)}_{\tau+(j/m)-(k-1)/m}\right), \tag{7}$$

where the forecasting horizon, h, is measured in the low frequency. Using the draws of $\left(\beta_h, \Theta_h\right)$, we compute a sequence of m probability forecasts by conditioning on a sequence of high-frequency informations sets: $\Omega_\tau, \Omega^{[1]}_\tau, \ldots, \Omega^{[m-1]}_\tau$. These high-frequency event probability forecasts allow us to evaluate the contribution of high-frequency information to low-frequency events, similar to the analyses in the nowcasting literature (e.g., Bańbura et al., 2013). The MIDAS-probit can then be interpreted as a way to filter high-frequency data to obtain predicted probabilities of low-frequency events.

3.3 | Evaluation of event probability forecasts

Recession probability forecasts have been evaluated using different methods. The pseudo-R^2 (Estrella & Mishkin, 1998) compares the log-likelihood function of the model with predictors to a model that includes only an intercept and is typically applied as an in-sample measure of fit.[13] An alternative is the AUROC (see Berge and Jorda, 2011, and Liu and Moench, 2016), which does not require forecasted probabilities to be converted into binary events, does not rely on a specific loss function, and has become the measure of choice for classification problems (Berge & Jorda, 2011). The AUROC, however, is not a proper score: Deviations from true probabilities can improve the score in some circumstances. Bouallègue et al. (2019) argue that the AUROC is a measure of potential skill in classification of binary events, but suggest alternative proper scores for rare events: the logarithm (or ignorance) score (LS) and the diagonal elementary score (DES).

Assume we observe predicted probabilities, $P_{\tau+h} = \Pr[y^*_{\tau+h} \geq 0|\Omega_\tau]$, computed over the out-of-sample period $\tau = L+1, \ldots, T-h$, where the number of forecasts over the out-of-sample is $R = T-h-L$. Then the out-of-sample LS score is

$$LS(h) = \frac{1}{R}\sum\nolimits_{\tau=L+1}^{T-h} - \ln\left|1 - S_{\tau+h} - P_{\tau+h}\right|.$$

The DES assumes that the objective is to maximize the classification of events and non-events (Bouallègue et al., 2018).[14] Bouallègue et al. (2019) recommend the DES over the LS when binary events are rare but have high-impact consequences and false positives do not cause large losses. The out-of-sample DES is

$$DES(h) = \frac{1}{R}\sum\nolimits_{\tau=L+1}^{T-h} \pi\mathbf{I}[P_{\tau+h} > \pi](1 - S_{\tau+h}) + (1-\pi)\mathbf{I}[P_{\tau+h} \leq \pi]S_{\tau+h},$$

where $\mathbf{I}[.]$ is an indicator function and $\pi = \frac{1}{R}\sum_{\tau=L+1}^{T-h} S_{\tau+h}$.

Following Bouallègue et al. (2019), we convert three of the metrics above to skill scores that compare the performance between the probability forecasts, $P_{\tau+h}$, and a reference forecast equal to the constant probability forecast using π (or the unconditional forecast). Skill scores measure the reliability and resolution of the probabilistic forecasts when using these proper score functions (Bouallègue et al., 2018) used to evaluate the performance of economic forecasters (Galbraith & van Norden, 2012). The skill score measures for the LS and DES are

$$LSS(h) = 1 - \frac{LS(h)}{LS^{unc}(h)},$$

$$DESS(h) = 1 - \frac{DES(h)}{DES^{unc}(h)}.$$

The LSS can be negative, where the $DESS$ is always positive. In both cases, these skill scores measure gains with respect to the unconditional probability forecast. One can convert the AUROC to a skill score:

$$ROCS = 2AUROC - 1.$$

[13]Chauvet and Potter (2005) uses the out-of-sample Briers score. However, Benedetti (2010) argues that it is not suitable for rare events. Because NBER recessions occur in only 10.8% of the out-of-sample months since 1977M1, we forgo analysis with the Briers score.

[14]The score is equivalent to the Kuipers score when the threshold is equal to the unconditional probability of the event π.

TABLE 1 Measures of fit and out-of-Sample performance of models for 1-year-ahead predictions of monthly NBER recessions using the yield curve spread

			Full sample (62M1–20M4)	Out-of-sample (77M4–20M4)		
Model	**Spec**	**n. lags**	**Marg.Lik.**	**ROC_S**	**LS_S**	**DES_S**
Probit	M	1	−193.56			
Probit	M	1, 6	−185.54	0.698	0.196	0.623
MIDAS-probit	Beta weig., W	24 weeks	−184.25			
MIDAS-probit	Beta weig., W	32 weeks	−183.59	0.770	0.288[a]	0.647
MIDAS-probit	Beta weig., W	50 weeks	−183.85			
MIDAS-probit	1st ord. Almon weig., W	32 weeks	−193.81			
MIDAS-probit	2nd ord. Almon weig., W	32 weeks	−200.71			
MIDAS-probit	Unrestricted, Shrink. Prior, W	32 weeks	−176.68	0.706	0.200	0.625

Note: The full sample marginal likelihood is computed using a modified harmonic estimator for **MI**xed **DA**ta **S**ampling (MIDAS)-probit specifications. The marginal likelihood was computed using 5000 draws of the posterior distribution of the parameters after removing the first 5000. Models are re-estimated every year with increasing samples to obtain predicted probability over the 516 monthly forecast origins from 1977M4 to 2019M4. ROC_S, LS_S, and DES_S are skill measures; that is, they measure how the predicted probabilities fare in comparison to a model that the predicted probability is always equal to the unconditional mean. Positive values indicate that models have a better resolution than the unconditional forecast. ROC_S is a skill based on the AUROC, LS_S is based on the logarithm score, and DES_S is based on the diagonal of the elementary score.

Abbreviation: MIDAS, **MI**xed **DA**ta **S**ampling.

[a]The model is significantly more accurate than the probit (at 10% level) using an equal-accuracy test modified for short samples and corrected for serial correlation and heteroscedasticity on the loss function differentials.

We apply these measures of the accuracy of event probabilistic forecasts in the next section.

4 | EMPIRICAL APPLICATIONS

In this section, we apply the MIDAS-probit model to three empirical macroeconomic research questions. The empirical exercises exploit how the MIDAS-probit model contributes to event probability forecasting in empirical macroeconomic research questions.

4.1 | NBER recession probabilities using the spread

As in Estrella and Mishkin (1998), Chauvet and Potter (2005), and Kauppi and Saikkonen (2008), we use the NBER chronology of business cycles, where the peak defines the last month of an expansion phase and the trough defines the last month of a recession phase. Define $S_t = 1$ as the month after the peak through the month at the trough. We use the spread between the 10-year Treasury bond and the 3-month Treasury Bill to construct 12-month-ahead forecasts. This exercise highlights the use of weekly-sampled spread data using Equation (2) compared with the monthly-sampled spread data using Equation (1) as a means of improving event forecasting accuracy.[15]

Our baseline MIDAS-probit specification sets $K = 32$, equivalent to about 7 months of weekly lags; we also consider alternatives with $K = 24$ and $K = 50$. For the probit using only monthly data on the spread, we consider both one monthly lag, as in Chauvet and Potter (2005), and both z_t and z_{t-6}, as in Liu and Moench (2016).[16] We also compare the MIDAS-probit with weighting function as in (3) with two alternatives: (i) Almon polynomials of first and second orders (see Mogliani & Simoni, 2021; Pettenuzzo et al., 2016) and an unrestricted specification with all $K = 32$ coefficients estimated using the Bayesian shrinkage priors (see Carriero et al., 2015).[17]

[15]As our aim is to evaluate possible gains from disaggregation (as in Andreou et al., 2010), we do not consider intra-month information sets in this exercise. Our main results are for the weekly-sampled spread. If we use the daily spread to forecast output growth, as suggested by Andreou et al. (2013), the AUROC is virtually the same as with weekly data. Thus, as in Galvão (2013), we prefer to use weekly—instead of daily—data.

[16]Chauvet and Potter (2005) and Kauppi and Saikkonen (2008) suggest the dynamic probit that includes lags of the latent variable y_t^* as predictors produces more accurate recession forecasts at short-horizons. While the dynamic probit accounts for serial correlation in the business cycle phases, it complicates multi-step-ahead forecasting. Serial correlation can be accommodated by adding economic activity (as shown in our second exercise). We leave study of the mixed-frequency dynamic probit for future research.

[17]For these two specifications, the Metropolis-in-Gibbs step described in Appendix A is not required, and equivalent simplification of the sampler is applied when computing the marginal likelihood.

We consider both full in-sample (1962M1 to 2020M4) and out-of-sample performance (1977M5 to 2020M4, $R = 516$). The marginal likelihood, computed as described in Appendix A, is our measure of full in-sample performance. The three skill measures described in Section 3.3 are used to compare out-of-sample performance. The full in-sample results in Table 1 suggest that the mixed-frequency specifications improve fit in comparison with the best monthly specification (with both z_t and z_{t-6}), except when using Almon polynomials. The specification with unrestricted slope coefficients fits better, and, for the specifications with beta weighting functions, $K = 32$ is the best option.

Based on these full-sample results, we consider three specifications in an out-of-sample exercise: (i) the probit model with both z_t and z_{t-6}; (ii) the MIDAS-probit with $K = 32$ unrestricted slope coefficients; and the MIDAS-probit with beta weighting polynomials and $K = 32$. Each specification is re-estimated every year during the out-of-sample period, and one-year-ahead forecasts are computed at each monthly forecast origin. The out-of-sample skill scores in Table 1 suggest that the spread improves the reliability and resolution of recession forecasts at a one-year-horizon, as all these scores are positive. The MIDAS-probit with the beta weighting function performs clearly better, and we find statistical improvements using the logscore by applying the modified Diebold and Mariano test as in Harvey et al. (2017).[18]

4.2 | Evaluation of forecasts for quarterly vulnerable growth events

As Adrian et al. (2019) has suggested, the identification of periods when GDP growth is " vulnerable" is a key input for stabilization policy, as adverse shocks may lead to future economic contractions. Identifying these periods of economic vulnerability is important to forward-looking policymakers. Vulnerable periods are based on the currently observed, quarterly-sampled year-on-year U.S. GDP growth ($g_t = 100[(GDP_t/GDP_{t-4}) - 1]$). If $g_t < 0.5$, we set $S_t = 1$ and interpret it as a vulnerable quarter. The threshold defining vulnerability is slightly larger than zero (0.5%) to encompass all NBER recession periods.[19]

Vulnerable event probabilities can be directly computed using a MIDAS-probit model applied to the vulnerable event just defined. An alternative method is to compute the event probabilities using density forecasts for year-on-year GDP growth. We compare MIDAS-probit event probabilities with equivalent predictions from the Double MIDAS model proposed by Pettenuzzo et al. (2016).[20] High-frequency financial variables enter as predictors for both the conditional mean and the conditional variance; Pettenuzzo et al. (2016) provide evidence that the Double MIDAS approach is particularly useful for density forecasting.

During the out-of-sample period (from 1985Q2 up to 2020Q1), the unconditional probability of a vulnerable growth quarter is 7.8%: even rarer than the monthly NBER recession and not far from the 5% growth-at-risk quantile in Adrian et al. (2019). Consequently, we consider skill measures based on the LSS and DES scores, in addition to the AUROC, (described in Section 3.3) to compare probabilistic forecasts for the contraction event.

4.2.1 | Vulnerability forecasting models

The literature suggests both the yield spread and the Chicago Fed Financial Condition Index (NFCI) as predictors of vulnerable growth events. The NFCI anticipates periods when growth is at risk (Adrian et al., 2019), although not

[18]Figure SA2 shows the slope estimated with the increasing samples over the out-of-sample period for the three specifications. For the unrestricted MIDAS-probit, we present the sum of all slope coefficients, including 68% credible intervals. The differences between the unrestricted and the beta-weighting specifications are more pronounced in the earlier part of the sample, suggesting that gains from the parsimony are more important when the estimation sample is short.

[19]Because the vulnerability event is constructed from a moving average of the past year's data, its timings are not simultaneous to the NBER business cycle phases. While there is a vulnerability event for all of the NBER recession phases, their timings are, in general, delayed.

[20]The Double MIDAS approach has two disadvantages in comparison with the MIDAS-probit approach for event probability forecasting. First, we are not able to easily compute credible intervals for the predicted probabilities. Second, the model needs to be re-estimated for each high-frequency horizon if the aim is to compute high-frequency forecasts of low-frequency events as described in Section 3.2.

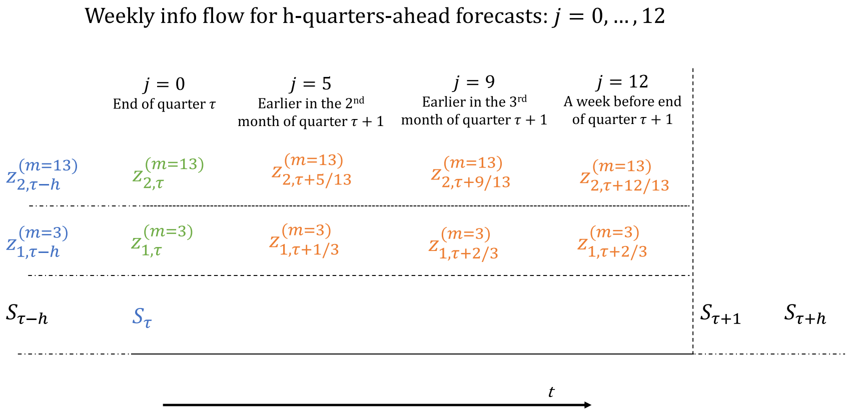

FIGURE 1 Flow of information. *Note:* Flow for a quarterly binary dependent variable S_t, and predictors of different frequencies: z_1 is monthly and z_2 is weekly. Forecasting target is $Prob(S_{\tau+h} = 1)$. When estimating the parameters at τ, values in blue are the last values including in the estimation. Then values in green are the conditioning information set at τ if only information up to the end of the quarter is considered. Values in orange are included in alternative conditioning sets using intra-quarter information for monthly ($m = 3$) and weekly series ($m = 13$). For comparison with forecasts using weekly information sets, conditioning monthly info indicated end of quarter τ is repeated up to $j = 4$, the info for $j = 5$ is repeated up to $j = 8$, then the value for $j = 9$ is repeated up to $j = 12$.

all of these periods materialize as vulnerable phases. The yield spread is often negative before recessions (Chauvet & Potter, 2005), so it may also lead vulnerability events. Both the spread and the NFCI are available weekly; thus, we use high-frequency data on both of these financial variables when computing contraction probabilities using the MIDAS-probit specification:

$$\Pr[y^*_{t+h} \geq 0|\Omega_t] = \Phi\left(\beta_{0,h} + \beta_{1,h}\sum_{k=1}^{K_1}\varpi\left(k;\theta_{1,h}\right)z^{(m=3)}_{1,t-\frac{k-1}{3}} + \beta_{2,h}\sum_{k=1}^{K_2}\varpi\left(k;\theta_{2,h}\right)z^{(m=13)}_{2,t-\frac{k-1}{13}}\right), \tag{8}$$

where z_1 is the Chicago Fed National Activity Index (CFNAI) sampled monthly and z_2 is one of the financial variables sampled weekly. To address whether high-frequency financial variables improve forecasting performance when the target is a low-frequency event (see also Plagborg-Møller et al., 2020), we include $z^{(m=3)}_{1,t-\frac{k-1}{3}}$, the monthly CFNAI, as a lower-frequency predictor measuring economic activity.[21]

Figure 1 describes how the intra-quarter information for both the weekly, $z^{(m=13)}_{2\tau}$, and the monthly, $z^{(m=3)}_{1\tau}$, regressors are employed to compute $\Pr[y^*_{\tau+h} \geq 0|\Omega^{[j]}_\tau]$ for $j = 0, \ldots, 12$. The last observations available to estimate (β_h, Θ_h) are in blue, the last observations in the conditioning information set at the end of the quarter τ in green, and the intra-quarter information sets in orange. The available information on $z^{(m=3)}_{1\tau}$ does not change every week as $z^{(m=13)}_{2\tau}$.

We compute one-quarter-ahead and 1-year-ahead forecasts for a benchmark probit specification using quarterly CFNAI. We compare that with the MIDAS-probit specification with: (i) only the monthly-sampled CFNAI, (ii) the monthly-sampled CFNAI plus the weekly-sampled term spread, and (iii) the monthly-sampled CFNAI plus weekly-sampled NFCI.[22] We also compute forecasts from Double MIDAS specifications using the latter two configurations to evaluate possible gains from directly estimating the event probabilities instead of extracting them from a predictive density. The Double MIDAS model uses a first-order Almon weighting function and is estimated using the Gibbs sampler and priors described in Section 3.2 of Pettenuzzo et al. (2016).[23]

[21] Berge and Jorda (2011) provide evidence that the monthly CFNAI is informative about US recession phases. This analysis extends the empirical results for GDP growth in Andreou et al. (2013) and Galvão (2013).

[22] We include 12 monthly lags of the CFNAI and either 32 weeks (approximately 7 months) of the term spread or 16 weeks of the NFCI. We experimented with alternative lag lengths for all predictors. Preliminary evidence on forecast performance support these choices.

[23] For the conditional mean, we use both the monthly-sampled CFNAI and the weekly-sampled financial predictors as the MIDAS-probit specification in (8). For the conditional variance, we use the square of the weekly-sampled financial predictor, in addition to the low-frequency autoregressive component.

TABLE 2 Out-of-sample performance of specifications to predict vulnerable growth quarters from 1985Q2 to 2020Q1 at one-quarter and four-quarter horizons

			$h = 1$			$h = 4$		
Model	**Predictors**	**Info time**	**ROC_S**	**LS_S**	**DES_S**	**ROC_S**	**LS_S**	**DES_S**
Probit	CFNAI, Q	τ	0.839	0.275	0.512	0.683	−0.059	0.008
MIDAS-probit	CFNAI, M	τ	0.900	0.333	0.512	0.296	−0.142	0.000
MIDAS-probit	CFNAI, M; Spread, W	τ	0.903	0.412	0.674	0.357	−0.095	0.251
Double MIDAS	CFNAI, M; Spread, W	τ	0.957	0.558	0.760	0.650	0.154	0.506
MIDAS-probit	CFNAI, M; NFCI, W	τ	0.894	0.545	0.723	0.285	0.045	0.199
Double MIDAS	CFNAI, M; NFCI, W	τ	0.940	0.483	0.457	0.376	−0.076	0.178
MIDAS-probit	CFNAI, M	$\tau+_{(5/13)}$	0.891	0.307	0.519	0.335	−0.133	0.000
MIDAS-probit	CFNAI, M; Spread, W	$\tau+_{(5/13)}$	0.885	0.370	0.667	0.257	−0.176	0.235
MIDAS-probit	CFNAI, M; NFCI, W	$\tau+_{(5/13)}$	0.889	0.527	0.739	0.271	0.068	0.206
MIDAS-probit	CFNAI, M	$\tau+_{(9/13)}$	0.879	0.266	0.504	0.413	−0.122	0.000
MIDAS-probit	CFNAI, M; Spread, W	$\tau+_{(9/13)}$	0.835	0.309	0.576	0.204	−0.239	0.228
MIDAS-probit	CFNAI, M; NFCI, W	$\tau+_{(9/13)}$	0.901	0.501	0.746	0.264	0.081	0.321

Note: Probit specifications are re-estimated at each 10 quarters over the out-of-sample period, but predicted probabilities are computed for 140 forecast origins. The number of lags is 12 months for CFNAI, 32 weeks for spread, and 16 weeks for NFCI. When employing the Double MIDAS specification, the squared of the weekly variable also enters the volatility equation. Gibbs algorithm and priors to estimate the Double MIDAS specification follow Pettenuzzo et al (2016).

Abbreviations: CFNAI, Chicago Fed National Activity Index; MIDAS, **MI**xed **DA**ta **S**ampling; NFCI, Chicago Fed Financial Condition Index.

4.2.2 | Out-of-sample comparison

We compute the out-of-sample probability forecasts for vulnerable growth events from 1985Q2 through 2020Q1 ($R = 140$).[24] Because the coefficients, $\boldsymbol{\beta}_h$ and Θ_h, are relatively stable, we re-estimate the MIDAS-probit models every 10 quarters with increasing samples.[25] In order to get updated forecasts for the conditional variance, we re-estimate the Double MIDAS models at each new forecast origin.

Table 2 presents the out-of-sample skill scores for one- and four-quarter-ahead forecasts of vulnerable growth quarters for the models described above. The skill scores for the first panel show results using the information set available at τ using estimates obtained with data up to τ. These results are available for all specifications. The second panel shows results using intra-quarter information as in Equation (7) using estimates up to τ. These are the information sets up to earlier in the second month ($j = 5$) and in the third month ($j = 9$) of the current quarter $\tau + 1$. These results are only available for the MIDAS-probit specifications. The results in Table 2 suggest that the inclusion of one of the weekly financial variables improves the reliability of the vulnerability forecasts for both horizons. Whether the NFCI outperforms the spread depends on the horizon, the loss function, and the information set available. The most accurate forecasts, however, are obtained by first computing density forecasts for output growth using the spread as the predictor in the Double MIDAS model and then computing the vulnerability event probabilities.

Figure 2 displays the predicted probabilities sampled quarterly for both the MIDAS-probit and the Double MIDAS for both horizons using the same set of mixed frequency predictors (CFNAI and either the spread or the NFCI). When using the spread, the Double MIDAS improves vulnerable growth quarters classification by reducing the number of false alarms. When using NFCI, the Double MIDAS performance relative to the MIDAS-probit depends on skill measure. The second panel of Table 2 also suggests that the performance of the MIDAS-probit using the spread deteriorates by updating forecasts with the current quarter information.

4.2.3 | Intra-quarter (weekly) out-of-sample forecast comparison

We now explore the effects of varying the information set each week within a quarter by constructing forecasts using the posterior mean draws for $\boldsymbol{\beta}_h$ and Θ_h. For each quarterly forecast origin, τ, we compute the intra-quarter probability

[24] The in-sample period employed to estimate each forecasting model varies with the availability of the predictors. The CFNAI data is available from 1968, the NFCI from 1971, and the spread from 1962.

[25] This strategy has no negative impact on forecasting performance compared with estimating at each new quarterly origin.

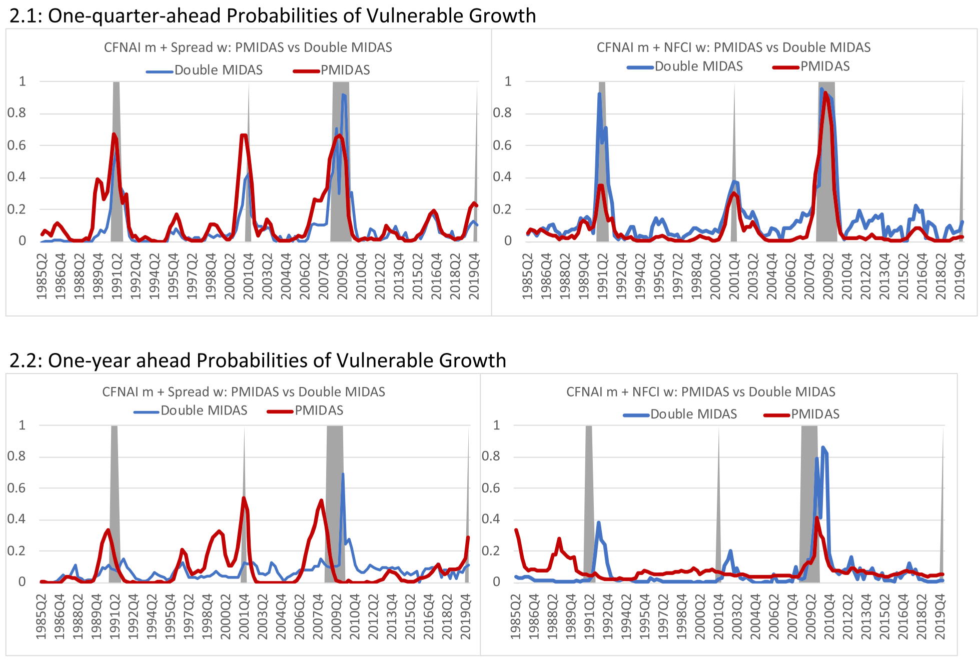

FIGURE 2 Out-of-sample vulnerability event probabilities: **MI**xed **DA**ta **S**ampling (MIDAS)-probit (PMIDAS) and Double MIDAS. *Note:* Shaded areas describe the vulnerability events and dates refer to their timings. Predicted probabilities were computed using end-of-quarter information sets either from one-quarter (upper panel) or 1-year (lower panel) before the indicated dates. The values are posterior mean estimates of the predicted probabilities for the MIDAS-probit. For the Double MIDAS, event probabilities are inferred from the predictive density of gross domestic product (GDP) growth. For both specifications, the number of lags is 12 months for Chicago Fed National Activity Index (CFNAI), 32 weeks for spread, and 16 weeks for NFCI. When employing the Double MIDAS specification, the squared of the weekly variable also enters the volatility equation. Gibbs algorithm and priors to estimate the Double MIDAS specification follow Pettenuzzo et al. (2016).

forecasts with $\Omega_\tau, \Omega_\tau^{[1]}, \ldots, \Omega_\tau^{[12]}$ described in (7). Figure 3 shows the means of the posterior probabilities for each weekly forecast origin $\tau+(j/m)$; shaded areas represent the realizations observed at $\tau+h$. Each forecast within a quarter predicts the outcome at $\tau+h$ for $\tau = L+1, \ldots, T-h$ as in Section 3.3. The predictive probabilities are computed for four models: (i) a probit model with quarterly-sampled CFNAI, (ii) a mixed-frequency version with only the monthly-sampled CFNAI, (iii) the monthly-sampled CFNAI and weekly-sampled spread, and (iv) the monthly-sampled CFNAI and the weekly-sampled NFCI. The top panel shows results for $h = 1$ and the bottom panel shows results for $h = 4$.

The inclusion of financial variables sharpens the identification of vulnerable growth quarters for both horizons, as suggested by comparing MIDAS-probit forecasts with CFNAI only to those that include financial variables. For the longer horizon, the inclusion of the spread as a weekly-sampled predictor leads to stronger signals of future vulnerable growth compared with alternatives. These signals, however, tend to lead to false positives in many occasions.

Figure 4 presents three different skill scores (ROCS, top row; LSS, center row; and DESS, bottom row) computed for the out-of-sample predicted probabilities in Figure 3 split for each weekly information set ($j = 0, \ldots, 13$ with R observations for each j). The values in the left column are for $h = 1$ and the ones in the right column for $h = 4$. In all cases, more accurate forecasts lead to higher skill values.

The updating of forecasts using intra-quarter weekly information sets does not generally improve forecasting performance, except when the NFCI is one of the predictors. Using more recent information sets, MIDAS-probit with NFCI performs best. Using only information up to the end of the quarter, τ, the relative performance between NFCI and the spread depends on the loss function and the horizon.

4.3 | Comparison with the SPF probabilities of negative GDP growth

The Survey of Professional Forecasters, published regularly by the Philadelphia Fed, asks respondents their prediction for the probability of negative quarterly GDP growth for the current quarter and the next four quarters. We re-estimate the MIDAS-probit specifications considered in 4.2 using $S_t = 1$ if the quarterly growth rate ($y_t = 100[(GDP_t/GDP_{t-1}) - 1]$)

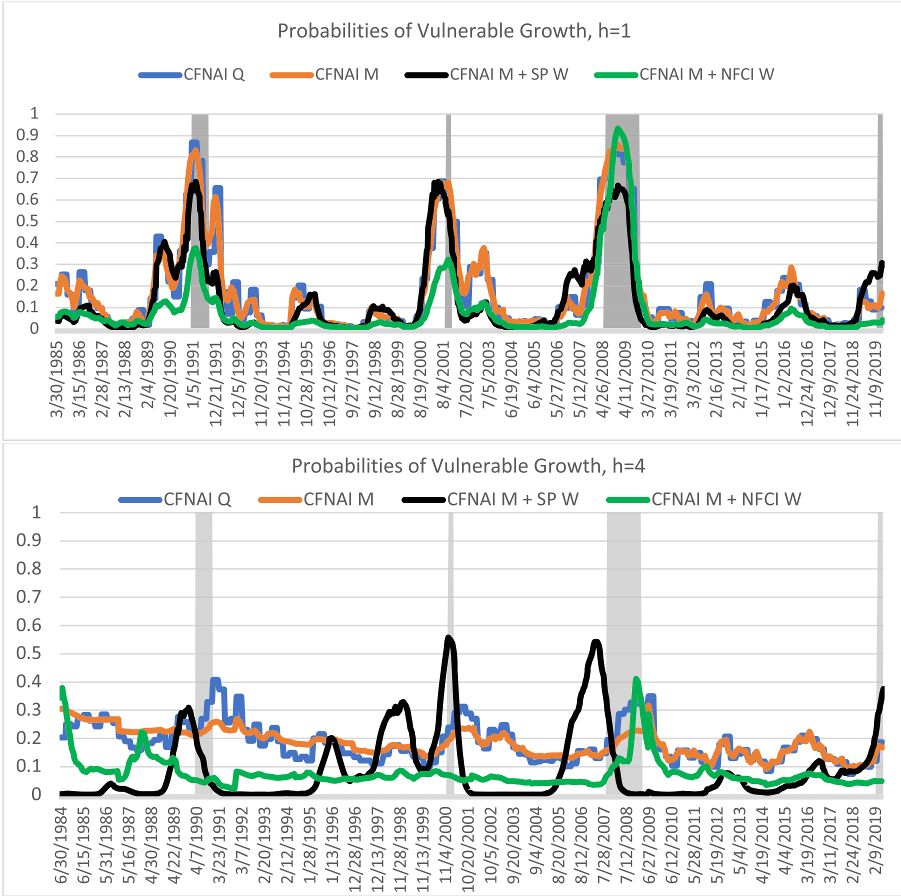

FIGURE 3 Posterior mean estimates of out-of-sample vulnerability event probabilities using the Chicago Fed National Activity Index (CFNAI), the spread, and the NFCI for weekly-updated information sets. *Note:* Dates refer to *last weekly information set employed to compute the probability forecast*. Shaded dates indicate periods of contraction for the outcome variable observed at $\tau + 1$ (top plot) or $\tau + 4$ (bottom plot) (from 1985Q2 to 2020Q1 for $h = 4$). These are computed with 5000 draws (after initial 5000 are removed). The model for the quarterly CFNAI (CFNAI) is a probit, all other predicted probabilities are computed using MIDAS-probit specifications. CFNAI M includes 12 monthly lags of CFNAI. The other two specifications follow Equation (8) using either the weekly spread (SP W) or the NFCI (NFCI W).

is negative. We are then able to compare the MIDAS-probit predicted probabilities for $h = 1$ and $h = 4$ over the sample out-of-sample period as in Table 2 with the SPF predictions. This provides an external evaluation of the forecasts obtained with the MIDAS-probit model.

Table 3 shows results for the same skill scores and period as Table 2 but assuming that the target is the probability of negative growth. In addition to the probit and the MIDAS-probit specifications using information up to τ (end of quarter), we show results for the SPF mean probabilities of negative growth. These results confirm that financial variables improve accuracy in comparison with models with only CFNAI, but they also show that SPF forecasts are more accurate for the shortest horizon. For 1-year-ahead forecasts, only the MIDAS-probit specification with the spread presents any predictive content for negative growth events. In summary, MIDAS-probit forecasts using financial variables are in particular useful for 1-year-ahead forecasts.

4.4 | Recession probabilities using the weekly economic index

Lewis et al. (2020) propose the weekly economic index (WEI) to monitor the US economy in real time. Indeed, since 21 March 2020, the weekly measure of economic activity has been updated and released every week. The index is computed using 10 different time series sampled weekly, including some traditional business cycle indicator variables, such as initial claims of unemployment insurance. As the NBER turning points chronologies are available at monthly frequency, the MIDAS-probit is a candidate approach to extract recession probabilities information from the WEI that may improve real-time classification of turning points.

FIGURE 4 Skill scores for out-of-sample vulnerability event probabilities using the Chicago Fed National Activity Index (CFNAI), the spread, and the NFCI for weekly-updated information sets. *Note:* The out-of-sample period is from 1985Q2 to 2020Q1 ($R = 140$). The horizonal axis describes j as in Figure 1, which describes how the information set is updated within the quarter. Details of the forecasting models are in the note to Figure 3 as these are skill scores computed for the predictions presented in Figure 3.

TABLE 3 Out-of-sample performance of SPF and MIDAS-probit to predict negative GDP growth quarters from 1985Q2 to 2020Q1 at one-quarter and four-quarter horizons

			$h = 1$			$h = 4$		
Forecaster	**Predictors**	**Info time**	**ROC_S**	**LS_S**	**DES_S**	**ROC_S**	**LS_S**	**DES_S**
SPF		τ	0.757	0.282	0.474	−0.001	−0.004	0.015
Probit	CFNAI, Q	τ	0.652	0.191	0.396	−0.201	−0.094	0.000
MIDAS-probit	CFNAI, M	τ	0.653	0.185	0.380	−0.481	−0.125	0.000
MIDAS-probit	CFNAI, M; Spread, W	τ	0.562	0.147	0.462	0.327	0.015	0.245
MIDAS-probit	CFNAI, M; NFCI, W	τ	0.639	0.204	0.407	−0.129	−0.006	−0.075

Note: Probit and MIDAS-probit specification as in Table 2, but the target event is negative quarterly GDP growth instead. SPF $h = 1$ forecasts are current quarter forecasts (RECESS1).
Abbreviations: CFNAI, Chicago Fed National Activity Index; GDP, gross domestic product; MIDAS, **MI**xed **DA**ta **S**ampling; NFCI, Chicago Fed Financial Condition Index.

4.4.1 | The MIDAS-probit specification

To obtain accurate real-time recession probabilities from the weekly WEI, we modify the MIDAS-probit specification in (2) to accommodate two relevant features of the data. First, the NBER Business Cycle Dating Committee calls peaks and troughs with a delay. On 9 June 2020 the committee called a peak for February 2020, implying a recession starting in March 2020. As a consequence, March, April, and May are unclassified when computing recession probabilities forecasts during these months. Thus, we estimate the model using the information on the binary variable only up through 2019, implying that the model outputs real-time classifications, varying only with changes in the availability of the information set as described in Section 3.3.

Second, the WEI may lag business cycle phases because it includes variables related to unemployment claims, which is classified as lagging the reference cycle in Stock and Watson (1999). Thus, we use both past and future information on

TABLE 4 Measuring the fit of MIDAS-probit specifications for using the WEI to predict NBER recessions over 2008M4–2019M12

Only lags	Marg. Lik.	Lags + leads	Marg. Lik.
$K_b = 5$	−23.178	$K_b = 5, K_f = 4$	−17.677
$K_b = 9$	−23.989	$K_b = 9, K_f = 8$	−15.337
$K_b = 14$	−25.428	$K_b = 14, K_f = 13$	−13.909

Note: The marginal likelihood was computed using 5000 draws of the posterior distribution of the parameters after removing the first 5000.
Abbreviations: MIDAS, **MI**xed **DA**ta **S**ampling; WEI, weekly economic index.

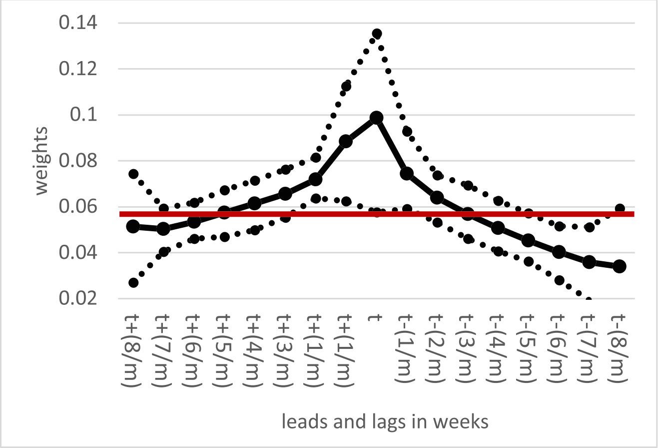

FIGURE 5 Posterior estimates of the weighting functions for **MI**xed **DA**ta **S**ampling (MIDAS)-probit with weekly economic index (WEI). Note: The above link two beta-weighting functions: one for the leads (+(8/m) up to +(1/m)) and the other for contemporaneous and lags (0 up to −(8/m)). For this specification, m is the number of weeks in a month (4.33). Weights are normalized to sum up to 1. Dotted lines are 68% bands. Sample period: 2008M1–2019M12. The red line indicates equal-weighting values.

weekly WEI to compute the probability of being in recession in month t. As our aim is to identify turning points that are published with a delay of a few months with an economic activity variable available with a delay of days, using available " future" information from leads (as proposed in Andreou et al., 2013) is feasible in real-time and may improve accuracy.

The MIDAS-probit model to extract information from leads and lags of weekly WEI, $z_t^{(m)}$, to compute pseudo-real-time recession probabilities is:

$$P\left[y_t^* \geq 0|\Omega_t^{[K_f]}\right] = \Phi\left(\beta_0 + \beta_f \sum_{k=1}^{K_f} \varpi\left(k;\boldsymbol{\theta}_f\right) z_{t+(k/m)}^{(m)} + \beta_b \sum_{j=0}^{K_b-1} \varpi\left(k;\boldsymbol{\theta}_b\right) z_{t-(k/m)}^{(m)}\right), \tag{9}$$

where $\varpi\left(k;\boldsymbol{\theta}_b\right)$ weights contemporaneous and lag weekly values of $z_t^{(m)}$, and $\varpi\left(k;\boldsymbol{\theta}_f\right)$ weights K_f lead values. In the algorithm described in Section 3.1, we impose restrictions on the parameters of the beta functions such that larger weights are given for $x_t^{(m)}$ at weeks that are near t. The restrictions are such that the weighting function, $\varpi\left(k;\boldsymbol{\theta}_f\right)$, decreases with the lead/lag horizon.

Table 4 shows the marginal likelihood computed from 5000 draws from the posterior density of the parameters after the initial 5000 are discarded for six different specifications of MIDAS-probit using the weekly-sampled WEI index as a predictor for NBER recessions observed monthly. We consider recession events from 2008M4 to 2019M12 using the WEI initial vintage published in 16 April 2020. The table presents results for specifications with only lags of WEI, $K_b = 5, 9, 14$, that is, using 1, 2, and 3 months of past information, and specifications that include $K_f = 4, 8, 13$ leads. In-sample fit

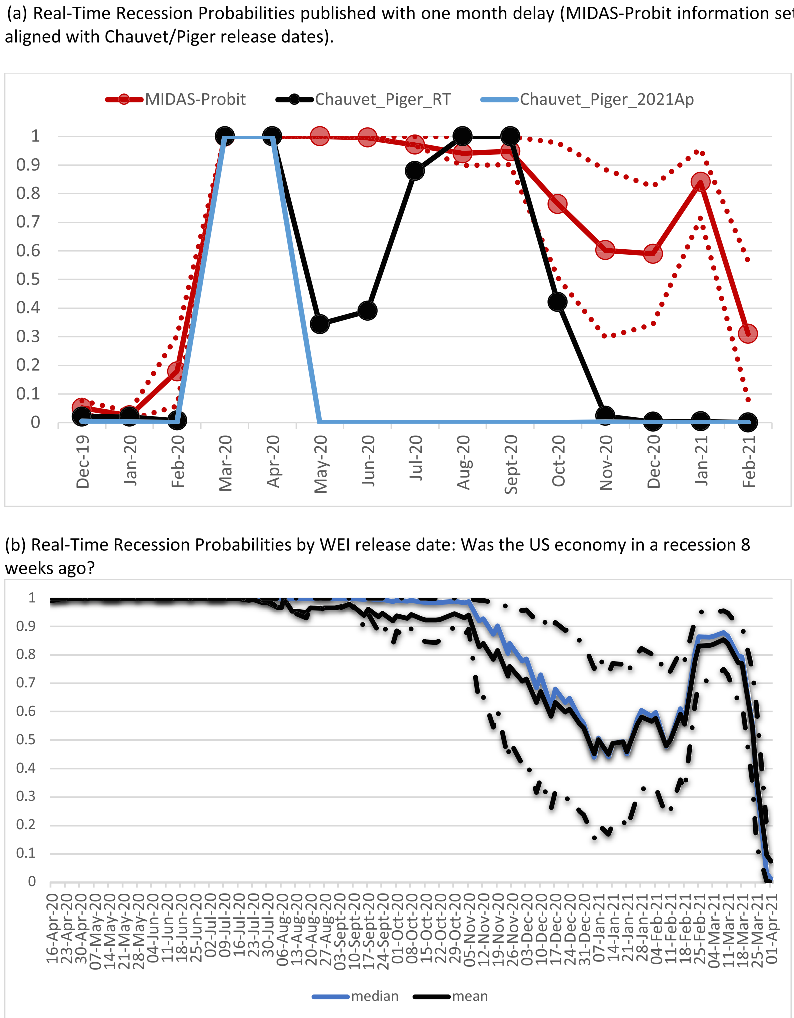

FIGURE 6 Real-time NBER recession probabilities using the **MI**xed **DA**ta **S**ampling (MIDAS)-probit with weekly economic index (WEI) and Chauvet and Piger Alternatives. *Note:* Dotted lines are 68% bands. MIDAS-probit parameters were estimated with data up to 2019M12. These are computed with 5000 draws (after initial 5000 are removed).

results clearly support the use of leads. For real-time predictions of recession probabilities, we chose the specification with $K_b = 9$ and $K_f = 8$.[26]

Figure 5 shows posterior mean estimates and 68% coverage bands for both weighting functions (lead and lag) with the weights normalized to sum up to 1. The first lead $\left(t+\frac{1}{m}\right)$ is assigned proportionately more weight than the first lag $\left(t-\frac{1}{m}\right)$, consistent with unemployment claims being a lagging indicator and requiring future values to identify turning points.

4.4.2 | Real-time recession probabilities and turning points

We consider two real-time recession prediction exercises to assess the usefulness of the MIDAS-probit model to extract information on recession probabilities out of the WEI.

The first one compares the MIDAS-probit with a Dynamic Factor Model with Markov-Switching, the model employed by Chauvet and Piger (2008) to publish their smoothed recession probabilities. Specifically, we employ the real-time vintages of Chauvet and Piger (2008) smoothed recession probabilities available at ALFRED (St. Louis Fed). We consider

[26]With this specification, we need to wait 30 days after the end of the reference month to compute its first recession probability. This "publication" delay is in line with the real-time competitor used below and considers the real-time availability of the WEI. In terms of fit, the data would have supported specifications with additional leads, as suggested in Table 3.

vintages from February 2020 to April 2021. For these vintages, the most recent probabilities are for the period between December 2019 and February 2021. Figure 6a presents these real-time smoothed recession probabilities and the full time series of the April 2021 smoothed probabilities for the months between 19 December and 21 February.[27]

For a fair comparison with the real-time smoothed probabilities, we use the MIDAS-probit parameters' posterior densities obtained using the first vintage of WEI and with observations to 2019M12. As discussed earlier, the MIDAS-probit specification is set as in (9) and has $K_b = 9$ and $K_f = 8$. Using the same posterior densities, we compute recession probabilities using monthly information sets. We use weekly data from the last vintage in a given month. To be compatible with a month delay, we use the vintage from the month that follows the reference month displayed in Figure 6a, which is equivalent to publication of recession probabilities with a month delay. Because the WEI was first published in April 2020, the probability predictions for 19 December to 20 February are not truly real-time but the subsequent values up to February 2021 are. Figure 6a displays the mean of the predicted probabilities and 68% credible intervals.

The MIDAS-probit is able to identify March as the first recession month, and recession probabilities are below 90% from October 2020 onward. The uncertainty on the predicted recession probabilities is substantial for the 20 October to 21 February period. These real-time probabilities suggest that, even though the date of a peak for the 2020 recession was easy to identify, the trough date is not clear using the weekly-sampled WEI. This exercise shows the usefulness of the MIDAS-probit approach in real-time, including the relevance of credible intervals for recession probabilities.

Our second real-time exercise uses all real-time vintages of WEI—published twice a week—from 16 April 2020 to 1 April 2021 to update estimates of whether the US economy was in a recession 8 weeks earlier. Figure 6b presents the predicted recession probabilities using the MIDAS-probit for each WEI real-time vintage. The figure displays the mean and median predicted probabilities and includes 68% credible bands. A recession is clearly defined using the predicted probability until the 5 November 2020 vintage. After that, predicted probabilities of recession are uncertain, with median values around 50% earlier than January 2021, followed by recession probabilities of 80% in March. This suggests the possibility of a recession phase until the end of 2020. This is longer than the recession phase dated by Chauvet and Piger (2008) and the NBER that ends in April 2020, and may reflect issues in using the WEI to identify the trough.

5 | CONCLUSIONS

In this paper, we propose a new tool for macroeconomic forecasting of critical low-frequency events. The MIDAS-probit model effectively delivers high-frequency probability forecasts for macroeconomic events by exploiting the predictive content of high-frequency financial and economic indicators. We provide empirical evidence that weekly-sampled financial variables help predict vulnerable growth events at a one-year horizon. The MIDAS-probit model with the yield curve spread sampled weekly improves one-year-ahead predictions of negative quarterly GDP growth compared to the SPF consensus and of NBER recessions compared to models with monthly-sampled spread. We also show how to filter the information of the WEI to obtain weekly-updated NBER recession probabilities.

We consider three different loss functions when evaluating the additional predictive content of financial variables to predict contractions. One of them, the diagonal of the elementary score, has been recently proposed in the climatology literature by Bouallègue et al. (2018) as a proper score to evaluate the classification ability of alternative forecasts of rare events. The diagonal score heavily penalizes false negatives (misses) and suggests that the MIDAS-probit model that includes the financial condition index improves the classification of future vulnerable growth quarters as high-frequency current quarter information is made available.

ACKNOWLEDGEMENTS

This paper builds on a draft written and circulated by Michael Owyang in 2011. M. Owyang benefitted from conversations with Eric Ghysels, Jim Hamilton, and Mike McCracken. Kristie Engemann, Kate Vermann, Hannah Shell, and Julie K. Bennett provided research assistance. The views expressed here are the authors' alone do not reflect those of the Federal Reserve Bank of St. Louis or the Federal Reserve System.

[27]The smoothed probabilities were heavily revised in the 21 January vintage when the authors re-estimated their model to indicate a peak in February 2020 and a trough in April 2020. The NBER has dated a peak in 2020 in June 2020 but at time of computation of these predictions (May 2021), it has not dated a trough yet. The trough was eventually called in 19 July 2021.

CONFLICT OF INTEREST

No conflict of interest has been declared by the authors.

OPEN RESEARCH BADGES

This article has earned an Open Data Badge for making publicly available the digitally-shareable data necessary to reproduce the reported results.

DATA AVAILABILITY STATEMENT

Data to replicate all empirical exercises in this paper are available for download from the JAE data archive (http://qed.econ.queensu.ca/jae/datasets/galvao003/).

ORCID

Ana Beatriz Galvão https://orcid.org/0000-0003-3263-9450
Michael Owyang https://orcid.org/0000-0002-2109-3432

REFERENCES

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

Albert, J. H., & Chib, S. (1993). Bayesian analysis of binary and polychotomous response data. *Journal of the American Statistical Association*, *88*(422), 669–679.

Andreou, E., Ghysels, E., & Kourtellos, A. (2010). Regression models with mixed sampling frequencies. *Journal of Econometrics*, *158*(2), 246–261.

Andreou, E., Ghysels, E., & Kourtellos, A. (2013). Should macroeconomic forecasters use daily financial data and how? *Journal of Business and Economic Statistics*, *31*(2), 240–251.

Audrino, F., Kostrov, A., & Ortega, J.-P. (2019). Predicting U.S. bank failures with MIDAS logit models. *Journal of Financial and Quantitative Analysis*, *54*(6), 2575–2603.

Auerback, A. J., & Gorodnichenko, Y. (2012). Measuring the output responses to fiscal policy. *American Economic Journal: Economic Policy*, *4*, 1–27.

Bańbura, M., Giannone, D., Modugno, M., & Reichlin, L. (2013). Now-casting and the real-time data flow, *Handbook of economic forecasting, volume 2a*: Elsevier, pp. 195–237.

Bauer, M. D., & Mertens, T. M. (2018). Economic forecasts with the yield curve. *FRBSF Economic Letter*, 2018-7.

Benedetti, R. (2010). Scoring rules for forecast verification. *Monthly Weather Review*, *138*(1), 203–211.

Berge, J., & Jorda, O. (2011). Evaluating the classification of economic activity into recessions and expansions. *American Economic Journal: Macroeconomics*, *3*(2), 246–277.

Bouallègue, Z. B., Haiden, T., & Richardson, D. S. (2018). The diagonal score: Definitions, properties, and interpretations. *Quarterly Journal of the Royal Meteorological Society*, *144*(714), 1463–1473.

Bouallègue, Z. B., Magnusson, L., Haiden, T., & Richardson, D. S. (2019). Monitoring trends in ensemble forecast performance focusing on surface variables and high-impact events. *Quarterly Journal of the Royal Meteorological Society*, *145*(721), 1741–1755.

Carriero, A., Clark, T., & Macellino, M. (2015). Realtime nowcasting with a Bayesian mixed frequency model with stochastic volatility. *Journal of the Royal Statistical Society, Series A*, *178*, 837–862.

Carriero, A., Clark, T., & Marcellino, M. (2020). Nowcasting tail risks to economic activity with many indicators. (*Working Paper No. 20-13R2*): Federal Reserve Bank of Cleveland.

Casarin, R., Foroni, C., Marcellino, M., & Ravazzolo, F. (2018). Uncertainty through the lenses of a mixed-frequency Bayesian panel Markov-switching model. *The Annals of Applied Statistics*, *12*(4), 2559–2586.

Casella, G., & George, E. I. (1992). Explaining the Gibbs sampler. *The American Statistician*, *46*(3), 167–174.

Chauvet, M., & Piger, J. (2008). A comparison of the real-time performance of business cycle dating methods. *Journal of Business and Economic Statistics*, *26*(1), 42–49.

Chauvet, M., & Potter, S. (2005). Forecasting recessions using the yield curve. *Journal of Forecasting*, *24*(2), 77–103.

Chib, S. (1995). Marginal likelihood from Gibbs output. *Journal of the American Statistical Association*, *90*(432), 1313–1321.

Chib, S., & Greenberg, E. (1995). Understanding the Metropolis-Hastings algorithm. *The American Statistician*, *49*(4), 327–335.

Clements, M. P., & Galvão, A. B. (2008). Macroeconomic forecasting with mixed-frequency data: Forecasting output growth in the United States. *Journal of Business and Economic Statistics*, *26*(4), 546–554.

Estrella, A., & Mishkin, F. S. (1998). Predicting U.S. recessions: Financial variables as leading indicators. *The Review of Economics and Statistics*, *80*(1), 45–61.

Foroni, C., Marcellino, M., & Schumacher, C. (2015). Unrestricted mixed data sampling (MIDAS): MIDAS regressions with unrestricted lag polynomials. *Journal of the Royal Statistical Society, Series A*, *178*, 57–82.

Foroni, C., Ravazzolo, F., & Rossini, L. (2019). Forecasting Daily Eletricity Prices with monthly macroeconomic variables. ECB Working Paper Series n. 2250.
Freitag, L. (2014). Default probabilities, CDS premiums and downgrades : A probit-MIDAS analysis. (*Research Memorandum 038*): Maastricht University.
Galbraith, J. W., & van Norden, S. (2012). Assessing gross domestic product and inflation probability forecasts derived from bank of England fan charts. *Journal of the Royal Statistical Society, series A*, *175*(3), 713–727.
Galvão, A. B. (2013). Changes in predictive ability with mixed frequency data. *International Journal of Forecasting*, *29*(3), 395–410.
Gelfand, A. E., & Smith, A. F. M. (1990). Sampling-based approaches to calculating marginal densities. *Journal of the American Statistical Association*, *85*(410), 398–409.
Geweke, J. (1999). Using simulation methods for bayesian econometric models: Inference, development,and communication. *Econometric Reviews*, *18*(1), 1–73.
Ghysels, E., Plazzi, A., Valkanov, R., Rubia, A., & Dossani, A. (2019). Direct versus iterated multi-period volatility forecasts: Why MIDAS is king. *Annual Review of Financial Economics*, *11*, 173–195.
Ghysels, E., & Qian, H. (2019). Estimating MIDAS regressions via OLS with polynomial parameter profiling. *Econometrics and Statistics*, *9*, 1–16.
Ghysels, E., Santa-Clara, P., & Valkanov, R. (2006). Predicting volatility: Getting the most out of return data sampled at different frequencies. *Journal of Econometrics*, *131*, 59–95.
Ghysels, E., Sinko, A., & Valkanov, R. (2007). MIDAS regressions: Further results and new directions. *Econometric Reviews*, *26*(1), 53–90.
Greenberg, E. (2013). *Introduction to Bayesian econometrics* (2nd ed.).: Cambridge University Press.
Harvey, D. I., Leybourne, S. J., & Whitehouse, E. J. (2017). Forecast evaluation tests and negative long-run variance estimates in small samples. *International Journal of Forecasting*, *33*(4), 833–847.
Kauppi, H., & Saikkonen, P. (2008). Predicting U.S. recessions with dynamic binary response models. *Review of Economic and Statistics*, *90*(4), 777–791.
Kuzin, V., Marcellino, M., & Schumacher, C. (2013). Pooling versus model selection for nowcasting GDP with many predictors: Empirical evidence for six industrialized countries. *Journal of Applied Econometrics*, *28*(3), 392–411.
Lewis, D., Mertens, K., Stock, J., & Trivedi, M. (2020). Measuring real activity using a weekly economic index. (*Staff Reports, No. 920*): Federal Reserve Bank of New York.
Liu, W., & Moench, E. (2016). What predicts US recessions? *International Journal of Forecasting*, *32*(4), 1138–1150.
Manasse, P., & Roubini, N. (2009). “Rules of thumb” for sovereign debt crisis. *Journal of International Economics*, *78*, 192–205.
Mogliani, M., & Simoni, A. (2021). Bayesian MIDAS penalized regressions: Estimation, selection and prediction. *Journal of Econometrics*, *222*(1), 833–860.
Pettenuzzo, D., Timmermann, A., & Valkanov, R. (2016). A MIDAS approach to modelling first and second moment dynamics. *Journal of Econometrics*, *193*(2), 315–334.
Plagborg-Møller, M., Reichlin, L., Ricco, G., & Hasenzagl, T. (2020). When is growth at risk? *Brooking Papers on Economic Activity*, Spring, 167–229.
Stock, J. H., & Watson, M. W. (1999). Business cycle fluctuations in US macroeconomic time series. In: *Handbook of macroeconomics* (Vol. 1, pp. 3–64). Elsevier.

SUPPORTING INFORMATION

Additional supporting information can be found online in the Supporting Information section at the end of the article.

How to cite this article: Galvão A. B., Owyang M. Forecasting low-frequency macroeconomic events with high-frequency data. *Journal of Applied Econometrics*. 2022;1-20. https://doi.org/10.1002/jae.2931

APPENDIX A: BAYESIAN ESTIMATION OF MIDAS-PROBIT MODELS

A.1 | The Metropolis-in-Gibbs algorithm

The MCMC sampler for the MIDAS-probit in (5) can be broken down into blocks: the block for the slope coefficients, $\boldsymbol{\beta}_h$; the beta weighting function parameters, Θ_h; and the latent data, $\left\{\mathbf{y}^* = \mathbf{y}_\tau^{*h} = y^*_{1+h}, y^*_{2+h}, \dots, y^*_{\tau-h+h}\right\}$.

As described in Table A1, we adopt the standard normal prior for the slope coefficients and the constant, and we make use of the identification restriction that $var(u_{t+h}) = 1$ as indicated in (6). The priors for the parameters of the weighting function are gamma distributed (as these parameters should be positive) and constructed to center around the belief that the high-frequency data is equally weighted.[28] Table A1 also includes the hyperparameter Δ_i that is designed to control the acceptance of the metropolis step.

TABLE A1 Priors for estimation

Parameter	Prior distribution	Hyperparameters
β_h	$N(\mathbf{m}_0, \mathbf{M}_0)$	$\mathbf{m}_0 = \mathbf{0}_{N+1}$; $\mathbf{M}_0 = \mathbf{I}_{N+1}$
$\theta_{i,1}, \theta_{i,2},$	$\Gamma(d_0, D_0)$	$d_0 = 1$; $D_0 = 1$; Δ_i

A.1.1 | Drawing $\boldsymbol{\beta}_h$ conditional on $\Theta_h, \mathbf{y}^*$

Conditional on Θ_h and $\mathbf{y}^*$, (6) is a linear regression. Let $\mathbf{Z}$ represent the $(\tau - h) \times (1 + N)$ matrix of stacked $\mathbf{Z}_t(\Theta_h)'$ vectors. Then, given the prior $N(\mathbf{m}_0, \mathbf{M}_0)$, a draw of $\boldsymbol{\beta}_h$ can be made from $\boldsymbol{\beta}_h | \Theta_h, \mathbf{y}^* \sim N(\mathbf{m}, \mathbf{M})$, where

$$\mathbf{M} = \left(\mathbf{M}_0^{-1} + \mathbf{Z}'\mathbf{Z}\right)^{-1}$$

and

$$\mathbf{m} = \mathbf{M}\left(\mathbf{M}_0^{-1}\mathbf{m}_0 + \mathbf{Z}'\mathbf{y}^*\right).$$

A.1.2 | Drawing Θ_h conditional on $\boldsymbol{\beta}, \mathbf{y}$

Obtaining a draw of $\boldsymbol{\theta}_i$ ($i = 1, \ldots, N$, where the h subscript is removed for simplicity here) can be accomplished using a Metropolis-in-Gibbs step (Chib & Greenberg, 1995) to sample from the nontractable posterior distribution. The Metropolis step requires a candidate draw from a proposal density, which is accepted with a probability that depends on both the likelihood and parameters' prior distribution.

We utilize a Gamma proposal density, whose hyperparameters depend on the previous accepted draw. In other words, for the j iteration, we draw a candidate $\boldsymbol{\theta}_i^{[j]} = \left(\theta_{i,1}^{[j]}, \theta_{i,2}^{[j]}\right)'$ from

$$\theta_{i,1}^{[j]} \sim \Gamma\left(\sqrt{\Delta_{i,1}}\theta_{i,1}^{[j-1]}, \left(\Delta_i \theta_{i,1}^{[j-1]}\right)^2\right) \tag{A1}$$

$$\theta_{i,2}^{[j]} \sim \Gamma\left(\sqrt{\Delta_{i,2}}\theta_{i,2}^{[j-1]}, \left(\Delta_i \theta_{i,2}^{[j-1]}\right)^2\right) \tag{A2}$$

where the superscript $j - 1$ represents the draw from the previous iteration. The hyperparameter Δ_i is a scaling factor that can be tuned to achieve a reasonable acceptance rate. As we have prior information on the most likely shape of the beta weighting function (increasing, decreasing, hump-shaped) for a given empirical application, then we use this information to accept only candidate draws compatible with our prior view on the shape of the beta function. For example, if we think the weighting function should be decreasing as in Figure 1, then we repeat (A1) and (A2) until we find a candidate draw that satisfies $\theta_{i,1}^{[j]} \leq \theta_{i,2}^{[j]}$. These restrictions help with identification of the weighting function and slope parameters and are a clear advantage of our estimation strategy.

The candidate draw is then accepted with probability $A = \min\{\alpha, 1\}$, where

$$\alpha = \frac{f\left(\mathbf{y}^* | \boldsymbol{\theta}_i^{[j]}\right)}{f\left(\mathbf{y}^* | \boldsymbol{\theta}_i^{[j-1]}\right)} \frac{dG\left(\boldsymbol{\theta}_i^{[j]} | \mathbf{d}_0, \mathbf{D}_0\right)}{dG\left(\boldsymbol{\theta}_i^{[j-1]} | \mathbf{d}_0, \mathbf{D}_0\right)} \frac{dG\left(\boldsymbol{\theta}_i^{[j-1]} | \sqrt{\Delta_i}\boldsymbol{\theta}_i^{[j]}, \left(\Delta_i \boldsymbol{\theta}_i^{[j]}\right)^2\right)}{dG\left(\boldsymbol{\theta}_i^{[j]} | \sqrt{\Delta_i}\boldsymbol{\theta}^{[j-1]}, \left(\Delta_i \boldsymbol{\theta}^{[j-1]}\right)^2\right)}, \tag{A3}$$

[28]We could adopt a diffuse prior over the $\boldsymbol{\theta}$ hyperparameters. This would be an improper prior and would be invalid for computation of marginal likelihoods.

where $f(.|.)$ reflects the conditional likelihood whose log is

$$\ln f\left(\mathbf{y}^*|\theta_i^{[j]}\right) = \sum_t \ln \phi\left[y^*_{t+h} - (\mathbf{Z}_t(\Theta_h^{[j]})'\boldsymbol{\beta}_h)\right]$$

and $dG(.|.) = \prod_{s=1}^{2}\Gamma\left(\theta_{i,s}^{[j]}|.\right)$ is the gamma pdf and $\phi()$ is the normal pdf. Note that these steps are designed to draw the vector $\boldsymbol{\theta}_i$, that is, the parameters of one weighting function at a time. This means that we draw $\boldsymbol{\theta}_i|\boldsymbol{\theta}_{\neq i}$. In the empirical applications covered in this paper, we either have $N = 1$ or $N = 2$.

A.1.3 | Drawing $\mathbf{y}^*$ conditional on $\boldsymbol{\beta}_h, \Theta_h$

Given the parameters $\boldsymbol{\beta}_h$ and Θ_h and the observables, we draw the vector $\mathbf{y}^*$ element by element, that is, $y^*_{1+h}, \ldots, y^*_\tau$ (as regression errors are assumed to be *iid*). Each value is drawn at each MCMC iteration from a truncated normal density, that is,

$$y^*_{t+h} \sim \begin{cases} TN_{(-\infty,0]}\left(\mathbf{Z}_t(\Theta_h)'\boldsymbol{\beta}_h, 1\right) & \text{if } S_{t+h} = 0 \\ TN_{(0,\infty)}\left(\mathbf{Z}_t(\Theta_h)'\boldsymbol{\beta}_h, 1\right) & \text{if } S_{t+h} = 1 \end{cases}$$

$$\text{for } t = 1, \ldots, \tau - h.$$

A.1.4 | The predicted probabilities

For each posterior draw of $\boldsymbol{\beta}_h, \Theta_h$ and $\mathbf{y}^*$, we compute predictive probabilities using:

$$P(S_{t+h} = 1) = P\left(y^*_{t+h} \geq 0\right) = \Phi[\mathbf{Z}_t(\Theta_h)'\boldsymbol{\beta}_h].$$

A.2 | Computation of the marginal likelihood

The marginal likelihood is useful to compare different MIDAS-probit specifications, including to select the lag order K_n.

Define the marginal likelihood as $p(\mathbf{Y})$, where $\mathbf{Y}$ includes all data, that is both $\{y_{t+h}\}_{t=1}^{t=\tau-h}$ and $\{\mathbf{Z}_t\}_{t=1}^{t=\tau-h}$. Then set $p(\mathbf{Y}|\Theta_h)$ as the marginal likelihood conditional on values for the beta weighting function parameters. $p(\mathbf{Y}|\Theta_h)$ can be obtained using the probit marginal likelihood proposed by Chib (1995), that is,

$$\ln(p(\mathbf{Y}|\Theta_h)) = \ln f(\mathbf{Y}|\hat{\boldsymbol{\beta}}_h, \Theta_h) + \ln \phi(\hat{\boldsymbol{\beta}}_h|\mathbf{m}_0, \mathbf{M}_0) - \ln\left(D^{-1}\sum_{d=1}^{D}\phi(\hat{\boldsymbol{\beta}}_h|\mathbf{m}^d, \mathbf{M})\right) \tag{A4}$$

where $\ln f(\mathbf{Y}|\hat{\boldsymbol{\beta}}_h, \Theta_h) = \sum_{t=1}^{\tau-h} y_t \ln \Phi(\mathbf{Z}_t(\Theta_h)'\hat{\boldsymbol{\beta}}_h) + (1 - y_t)\ln[1 - \Phi(\mathbf{Z}_t(\Theta_h)'\hat{\boldsymbol{\beta}}_h)]$. We keep D sampler draws after removing the first 5000 draws such that $\hat{\boldsymbol{\beta}}_h = D^{-1}\sum_{d=1}^{D}\boldsymbol{\beta}_h^d$.

The marginal likelihood is then computed using a Modified Harmonic Mean Estimator as

$$\widehat{p(\mathbf{Y})} = \left[\frac{1}{D}\sum_{d=1}^{D}\frac{f\left(\Theta_h^d\right)}{p\left(\mathbf{Y}|\Theta_h^d\right)p\left(\Theta_h^d|\mathbf{d}_0, \mathbf{D}_0\right)}\right]^{-1}$$

where $p(\Theta_h^d|\mathbf{d}_0, \mathbf{D}_0)$ evaluates the draw Θ_h^d at gamma prior (with parameters set to 1). The prior is evaluated at each parameter in Θ_h (which is a vector of dimension $2N$, where N is the number of high-frequency predictors) such that $p(\Theta_h^d|\mathbf{d}_0, \mathbf{D}_0)$ multiplies all these evaluations. The term $p(\mathbf{Y}|\Theta_h^d)$ is obtained using the conditional marginal density in (A4) evaluated at each draw Θ_h^d, as values of Θ_h^d affect the first and the third term of $\ln(p(\mathbf{Y}|\Theta_h))$. The function $f(\Theta_h^d)$ employs a multivariate Gaussian large sample approximation for empirical distribution of the draws Θ_h^d for $d = 1, \ldots, D$. The purpose of $f(\Theta_h^d)$ is to remove the impact of tail draws in the computation of the marginal likelihood to improve the stability of the marginal likelihood estimator as suggested by Geweke (1999). We set $f(\Theta_h^d)$ such that 5% of the draws Θ_h^d are chopped based on the Gaussian approximation.