7-2021

Forecast pooling or information pooling during crises? MIDAS forecasting of GDP in a small open economy

Hwee Kwan CHOW-TAN
Singapore Management University, hkchow@smu.edu.sg

Daniel HAN

SMU ECONOMICS & STATISTICS

Forecast Pooling or Information Pooling During Crises? MIDAS Forecasting of GDP in a Small Open Economy

Hwee Kwan Chow, Daniel Han

Jul 2021

Paper No. 06-2021

Forecast Pooling or Information Pooling During Crises? MIDAS Forecasting of GDP in a Small Open Economy

Hwee Kwan Chow* and Daniel Han

School of Economics
Singapore Management University

Abstract

This study compares two distinct approaches, pooling forecasts from single indicator MIDAS models versus pooling information from indicators into factor MIDAS models, for short-term Singapore GDP growth forecasting with a large ragged-edge mixed frequency dataset. We investigate their relative predictive performance in a pseudo-out-of-sample forecasting exercise from 2007Q4 to 2020Q3. In the stable growth non-crisis period, no substantial difference in predictive performance is found across forecast models. We find factor MIDAS models dominate both the quarterly benchmark model and the forecast pooling strategy by wide margins in the Global Financial Crisis and the Covid-19 crisis. Reflecting the small open nature of the economy, pooling single indicator forecasts from a small subgroup of foreign-related indicators beats the benchmark, offering a quick method to incorporate timely information for practitioners who have difficulty updating a large dataset. Nonetheless, the information pooling approach retains its superior ability at tracking rapid output changes during crises.

Keywords: *Forecast evaluation, Factor MIDAS, pooling GDP forecasts, global financial crisis, Covid-19 pandemic crisis*

JEL Classification: *C22, C53, C55*

1. Introduction

In a fast-evolving environment, such as in a crisis, the information content of higher frequency variables can be particularly useful to provide a more timely assessment of current and near-term future economic conditions. In any case, it is not uncommon for forecasters to draw information from a wide variety of economic and financial indicators sampled at different frequencies when forecasting the output growth of an economy. For instance, central bankers would use a host of domestic and foreign indicators—some of which are monthly variables—to produce short-term forecasts of quarterly GDP growth. The traditional approach to dealing with mixed frequency data is to time-aggregate the higher frequency variables to a lower frequency. However, temporal aggregation may lead to inefficiency and biased or inconsistent estimation of the parameters (Andreou, Ghysels and Kourtellos, 2010). Besides, time aggregation would also lead to a loss of high frequency information at the end of the indicator series.

There are various methods to exploit the information in higher frequency data to predict a lower frequency variable. In particular, the mixed-data sampling framework (MIDAS) due to Ghysels, Sinko and Valkanov (2007) directly relates mixed frequency variables in a highly parsimonious way. Since the seminal work of Clements and Galvao (2008), there has been a burgeoning literature on the application of MIDAS techniques to macroeconomic forecasting, particularly for larger economies like the US and the euro-area (see *inter alia* Armesto, Engemann and Owyang, 2010; Foroni and Marcellino, 2014). Applications to small open economies are found in more recent studies such as Kim and Swanson (2017) for Korea; Hueng and Yau (2019) for Taiwan; and Galli, Hepenstrick and Scheufele (2019) for Switzerland.

In addition to dealing with data of different frequencies, forecasters often need to consider what strategies to adopt to pool the predictive content from a large number of variables. After all, the traditional approach of selecting only a few indicators and performing forecasting using a small-scale model is problematic as the information content of individual indicators would tend to vary over time. The two common approaches to extract the predictive content from a large number of variables into GDP growth forecasts are pooling forecasts from many single indicator MIDAS models versus pooling a large amount of information into a few factors for inclusion into a single factor MIDAS model. The factors represent the underlying movement in the predictors and can be estimated using a state-space framework to account for the ragged edge of the dataset that reflects varying publication delays of the indicators.

It is well-recognised that pooling forecasts from different models helps to average out idiosyncratic errors arising from the misspecification of individual models. Theoretical results are found in Timmerman (2006), while Bates and Granger (1969), amongst others, provide empirical evidence. In comparison, pooling information from multiple predictors into a single model can average out the noise from individual predictors, as shown in Forni et al. (2003). The debate on which of these two strategies produces more accurate forecasts is fuelled by mixed evidence from empirical studies. Some in the literature, including Heinisch and Scheufele (2018), provide empirical evidence that the pooled single indicator forecasts strategy outperforms the factor-based information pooling approach in the case of Germany, while Kuzin, Marcellino and Schumacher (2013) show otherwise for six industrialised countries. A natural question that arises is which of these two pooling strategies is more beneficial for forecasting in the context of a small open economy.

This study focuses on the short-term forecasting of quarterly real GDP growth of Singapore, an archetypal small open economy where external factors play a prominent role in driving domestic fluctuations. We construct a large, ragged-edge panel dataset of 95 monthly variables comprising domestic and foreign indicators that reflect important aspects of the Singapore economy. To assess the forecast performance of the two approaches, we conduct a pseudo-out-of-sample GDP forecasting exercise by recursively generating GDP growth forecasts for the period 2007Q4 to 2020Q3. Within each approach, we generate and combine the forecasts across four variants of MIDAS models. The forecast evaluation period is split into the following three subperiods: a cycle which includes the global financial crisis (GFC), the COVID-19 pandemic crisis and the non-crisis period in between. Further, forecasting with subgroups of indicators is also investigated.

This paper makes three main contributions to the literature. Firstly, the comparison of the forecast pooling versus information pooling strategies adds to the evidence in the literature on their relative usefulness under different regimes, namely crisis versus non-crisis periods, for short-term forecasting of GDP growth. Secondly, we present evidence that in the context of a very small open economy, forecast pooling based on a small subgroup of foreign-related indicators offers a quick method to incorporate high frequency information into output growth forecast without having to maintain a large dataset. Thirdly, to the best of our knowledge, our study is the first to apply the MIDAS technique to a large mixed-frequency dataset for forecasting Singapore quarterly GDP growth.

Only a couple of past studies adopt the MIDAS approach to forecast Singapore GDP growth, and all of them use very few predictors.[1] For instance, Tsui, Xu and Zhang (2018) employs only daily stock prices in a MIDAS forecast model while Abeysinghe (1998) uses only the external trade variable in a non-linear dynamic regression model. Our study fills the gap by applying the MIDAS framework to forecast Singapore GDP growth with many indicators, which is important since the multitude of external influences means forecasters need to consider not only domestic indicators but also a myriad of foreign ones.

The rest of this paper proceeds as follows. The next section presents an overview of the MIDAS models used in this study, while Section 3 provides a description of the data and the empirical forecast procedure adopted. Section 4 reports and discusses the findings of the pseudo-out-of-sample forecasting exercise, while Section 5 investigates the use of subgroups in forecast pooling. Section 6 concludes.

2. An Overview of MIDAS Models

This section describes the models used in our study. We first offer brief descriptions of the four variants of MIDAS models and then discuss the factor MIDAS approach. To simplify the

[1] Other studies that use many indicators to forecast Singapore GDP growth do not employ the MIDAS technique. For instance, Chow and Choy (2009a) generates Singapore GDP growth forecasts with a single frequency dynamic factor model using a large panel dataset whereby higher frequency data are aggregated down to quarterly frequency to overcome the mixed frequency problem.

description of the models, we consider the case of a quarterly dependent variable (y_t), which is Singapore GDP growth in this study, along with one monthly indicator (x_t). The models can be extended in a straightforward fashion to incorporate more predictors.

2.1 Variants of the MIDAS model

The basic MIDAS model proposed by Ghysels, Sinko and Valkanov (2007) has a forecast equation for h quarters ahead as follows:

$$y_{t+h} = \beta_0^{(h)} + \beta_1^{(h)} \sum_{j=0}^{J} \omega(j, \boldsymbol{\theta}) L^{j/3} x_{t+w}^{(3)} + \varepsilon_{t+h} \tag{1}$$

where $\omega(j, \boldsymbol{\theta}) = \frac{\exp\{\theta_1(j+1)+\theta_2(j+1)^2\}}{\sum_{j=0}^{J} \exp\{\theta_1(j+1)+\theta_2(j+1)^2\}}$ is the normalised exponential Almon polynomial weighting function with hyperparameters $\boldsymbol{\theta} = \{\theta_1, \theta_2\}$, $t = 1, 2, \ldots, T$, $L^{j/3} x_{t+w}^{(3)} = x_{t+w-j/m}^{(3)}$ and the disturbance term is assumed to be independently and identically distributed with zero-mean and constant variance. While the dependent variable is only available up to T, the last available observation of the regressor is at $T + w$. For example, $w = \frac{1}{3}$ corresponds to indicator information available for the first month of the forecast quarter. Hence, the forecast at $t = T$ is conditioned on the information set at $T + w$. The superscript on the distributed lag term in equation (1) indicates skip-sampling of monthly observations across quarters. The superscript h on the coefficients indicates that they are specific to the forecast horizon due to direct forecasting. As an example, $h = \frac{1}{3}, \frac{2}{3}$ and 1 correspond to the nowcasts made at the beginning of the third, second and first months of the forecast quarter respectively.

We select the lag length for all MIDAS variants by applying the Bayesian Information Criterion (BIC) with a maximum of 12 lags (1 year) for the monthly indicators. The normalised exponential Almon lag polynomial allows for weighting functions of different shapes depending on the values of the hyperparameters. As in most empirical studies, we restrict the hyperparameter vector $\boldsymbol{\theta}$ to the two-dimensional case for the sake of parsimony. In this way, the number of parameters to be estimated in the model is greatly reduced. Estimation of the models is carried out by the non-linear least-squares technique that applies the Broyden-Fletcher-Goldfarb-Shanno algorithm. Since the estimation outcome may be sensitive to initial values, we first choose models with the lowest residual sum of squares by searching over the set of hyperparameter values -100 ≤ $\theta_{\mathbf{1}}$ ≤ 5 and -100 ≤ $\theta_{\mathbf{2}}$ ≤ 0.[2] That is, we use the estimated values of the hyperparameters from the model with the lowest residual sum of squares as initial values in the estimation process. We also set the initial value of $\beta_{\mathbf{1}}$ to 1.

A commonly used variant of MIDAS (e.g., Andreou, Ghysels & Kourtellos, 2013; Tsui Xu and Zhang, 2018) is the ADL-MIDAS model, which is analogous to the standard autoregressive distributed lag model, except that the regressors and regressand are now sampled at different frequencies. The forecast equation for *h* quarters ahead is given by:

$$y_{t+h} = \beta_0^{(h)} + \beta_1^{(h)} \sum_{j=0}^{J} \omega(j, \boldsymbol{\theta}) L^{j/3} x_{t+w}^{(3)} + \sum_{q=0}^{Q} \gamma_{1+q}^{(h)} L^q y_t + \varepsilon_{t+h} \tag{2}$$

where the weighting function $\omega(j, \boldsymbol{\theta})$ is as defined in equation (1), and all coefficients are estimated via non-linear least squares as with the basic MIDAS model.

Unlike the two aforementioned models, the unrestricted MIDAS (U-MIDAS) model, due to Foroni, Marcellino and Schumacher (2015), does not impose any functional constraints on the distributed lags, increasing flexibility in the specification. This approach may be feasible in our case because of the small difference in sampling frequencies between the dependent variable and indicator, which implies that there are relatively few parameters to be estimated. As is the case of the ADL-MIDAS model, we can augment the U-MIDAS model with lagged dependent variables. This is known as the ADL-U-MIDAS model, and its forecast equation for *h* quarters ahead is as follows:

$$y_{t+h} = \beta_0^{(h)} + \sum_{j=0}^{J} \beta_{1+j}^{(h)} L^{j/3} x_{t+w}^{(3)} + \sum_{q=0}^{Q} \gamma_{1+q}^{(h)} L^q y_t + \varepsilon_{t+h} \tag{3}$$

[2] The grid search is performed in increments of one for each hyperparameter. For multiple-factor models, we restrict the initial values to be the same across indicators to reduce computational complexity.

The corresponding forecast equation for the U-MIDAS model is the same as equation (3) but without the autoregressive terms. Meanwhile, the benchmark model is an autoregressive model that involves only quarterly GDP data, and its corresponding forecast equation is the same as equation (3) but without the indicator terms.

2.2 Factor MIDAS models

Factor MIDAS, due to Marcellino and Schumacher (2010), synthesises dynamic factor models with MIDAS models. The single indicators in the MIDAS models described in the previous sub-section are simply replaced with estimated monthly factors. These few factors summarise the systematic information in our large dataset. There are various factor extraction methods, but the literature on dynamic factor models reports conflicting results on which is superior. Nonetheless, Marcellino and Schumacher (2010) show that the choice of factor extraction technique does not substantially impact the short-term forecasting performance of factor MIDAS models. Similarly, others such as Kuzin, Marcellino and Schumacher (2013) find no systematic difference in forecast accuracy across different factor extraction methods. In this study, we adopt the two-step estimator by Doz, Giannone and Reichlin (2011), which relies on a state-space framework and can handle the ragged-edge structure in the data panel.

The state-space model for monthly variables and monthly factors is given below:

$$\boldsymbol{X}_{\boldsymbol{t_m}} = \boldsymbol{\Lambda}\boldsymbol{F}_{\boldsymbol{t_m}} + \boldsymbol{\xi}_{\boldsymbol{t_m}} \quad (4)$$

$$\boldsymbol{\Psi}(L_m)\boldsymbol{F}_{\boldsymbol{t_m}} = \boldsymbol{B}\boldsymbol{\eta}_{\boldsymbol{t_m}} \quad (5)$$

where $\boldsymbol{X}_{t_m}$ is a *n x 1* vector of monthly indicators, $\boldsymbol{\Lambda}$ is a *n x r* matrix of factor loadings for the *r* static factors, and $\boldsymbol{\xi}_{\boldsymbol{t_m}}$ is the idiosyncratic component. Equation (4) is the static factor representation as in Stock and Watson (2002). Equation (5) specifies the dynamics of the factors using a vector autoregression (VAR) of order *p*. $\boldsymbol{\eta}_{\boldsymbol{t_m}}$ is a *q x 1* vector of dynamic shocks orthogonal to the factors and their lags $\boldsymbol{\Psi}(L_m) = \sum_{i=1}^{p} \boldsymbol{\psi}_i L_m^i$, and $\boldsymbol{B}$ is an *r x q* coefficient matrix. We follow Galli, Hepenstrick and Scheufele (2019) in setting *p* = 1 for the sake of parsimony and consider *r* = 1, 2 or 3.

Unlike static principal components analysis (PCA), the state-space approach explicitly specifies the dynamics of the factors. As the dimension of $\boldsymbol{X}_{t_m}$ is large for our dataset, iterative maximum likelihood is not feasible. Instead, single-step Kalman smoothing is applied outside the model. The estimation procedure to extract the factors is as follows:

1. Static PCA is used to produce factor estimates $\widehat{\boldsymbol{F}_{\boldsymbol{t_m}}}$ using a balanced dataset (truncated at the end of sample).

2. The factor loading matrix $\boldsymbol{\Lambda}$ is estimated by regressing $\boldsymbol{X}_{t_m}$ on $\widehat{\boldsymbol{F}_{\boldsymbol{t_m}}}$, hence obtaining the covariance matrix of idiosyncratic components $\boldsymbol{\xi}_{\boldsymbol{t_m}}$.

3. The VAR of order p = 1 is estimated to obtain $\widehat{\boldsymbol{\Psi}}(L_m)$ and the residual covariance matrix.

4. Kalman smoothing applied over the unbalanced dataset produces the updated estimate $\widehat{\boldsymbol{F}_{\boldsymbol{t_m}}}$.

The extracted monthly factors are used to forecast quarterly GDP growth in the four variants of the MIDAS models. We abbreviate the corresponding counterparts as factor MIDAS, factor ADL-MIDAS, factor U-MIDAS and factor ADL-U-MIDAS respectively.

3. Data and Empirical Procedure

3.1 Data

Our large scale dataset comprises 95 monthly indicators, mostly collected from the CEIC and FRED databases.[3] Our dataset is similar to that of Chow and Choy (2009b), and a complete data listing is found in Appendix A. The broad categories of data are the GDP and leading indicators of major trading partners, foreign financial data, world electronics sales and indexes, world prices, industrial production, business expectations, sectoral indicators, external trade, domestic prices, financial indicators and exchange rates, and monetary and credit aggregates. The download date is 2 January 2021.

[3] The CEIC and FRED databases compile data from various official sources such as government agencies, national statistical sources and multilateral organisations. To achieve more parsimonious model specifications, weekly indicators are time-aggregated to monthly frequency.

Data for some series are available as early as 1955, but we perform the empirical exercise using information from 1990Q1 to 2020Q3, for which we have data for the vast majority of the time series. The different publication delays of the indicators result in an unbalanced panel dataset with a ragged edge. Due to a lack of data availability, we are unable to obtain real-time vintages and can only simulate a `pseudo-real-time' forecasting scenario. However, we think this need not be a concern since past studies, including Boivin and Ng (2005) and Breitung and Schumacher (2008), have shown that data revisions do not considerably impact forecast accuracy.

The data collected have generally been seasonally adjusted. Otherwise, manual adjustment is performed using the X-13 ARIMA procedure whenever seasonality is detected in the time series. We determine at the 5% significance level the order of integration of the individual series based on the Augmented Dickey-Fuller breakpoint unit root test and the KPSS test.[4] Appropriate transformations are taken to induce stationarity in the series, and these mostly involve taking the log-difference or first difference to obtain month-on-month growth rates. Lastly, outliers are identified by the cutoff rule—six times the interquartile range from the sample median—and these extreme observations are replaced by the lower or upper quartiles depending on which tail of the distribution they lie on. The data listing in Appendix A provides details on the data source, applied transformation and publication lags of the individual indicators.

3.2 Short-term Forecasting Procedure

In this study, all MIDAS estimations and forecasts are performed in R using the "midasr" package courtesy of the authors of Ghysels, Kvedaras and Zemlys (2016). Our full sample period is divided into the estimation and evaluation periods. In the first instance, we estimate each model over the initial sample from 1990Q1 to 2007Q3, selecting the lag lengths for the predictors based on the Bayesian Information Criterion (BIC). The first set of forecasts are thus generated for 2007Q4.

---

[4] The breakpoint unit root tests are conducted for the period 1990M1 to 2019M12, i.e. just before the onset of the pandemic crisis. The results of the unit root tests are available from the authors upon request.

Following this, we expand the estimation window forward by one quarter, re-select the lag lengths and re-estimate the model coefficients. The forecast is then computed for 2008Q1. This procedure continues recursively until forecasts for the entire evaluation period from 2007Q4 to 2020Q3 are generated. We use direct multistep forecasting, whereby a different forecast model is estimated for each horizon. An advantage of the direct method over the iterative approach is that misspecification in the one-step-ahead case is not carried over to the multistep-ahead forecasts. Further, direct forecasts avoid having to specify the change in unobserved factors over time for factor-based models (see Marcellino, Stock and Watson, 2006).

We consider monthly forecast horizons from h_1 *to* h_6 to generate a sequence of six forecasts for each evaluation quarter.[5] For instance, for the evaluation quarter 2019Q1, forecasts are computed on the 2nd of March, February and January 2019, corresponding to h_1, h_2 and h_3 respectively, and these current quarter forecasts are also known as "nowcasts". Horizons h_4, h_5 and h_6 refer to the one-quarter-ahead forecasts produced on 2nd of December, November and October in 2018 respectively. Each forecast is generated based on the information set available up to that point. The ragged edge structure in each recursion of our estimation procedure is replicated by imposing the same number of missing values observed for each indicator at the end of the sample. In doing so, we implicitly assume stability in the publication lag structure.

For all models with autoregressive terms, the maximum lag length of lagged dependent variables is four quarters or one year. We consider the fact that the GDP of a particular quarter is released towards the end of the second month of the next quarter.[6] Specifically, the maximum number of lagged dependent variables considered in the models will be four when the forecast horizon is h_1, three for horizons h_2, h_3 and h_4, and two for horizons h_5 and h_6. For instance, if the

[5] We denote $h = i/3$ as h_i for $i = 1,2,\ldots,6$ here and for the rest of the paper.

[6] The effective forecast horizons needed for computing the forecasts are longer when we account for the publication lag of GDP. For instance, when we are predicting GDP growth at 2019Q1 from 2nd February 2019, the end of sample data vintage would then comprise GDP data up to 2018Q3. This means we effectively require a two-quarter ahead forecast from the end of the GDP sample.

evaluation quarter is 2019Q1, last quarter's GDP growth at 2018Q4 would have been published and can therefore be used for forecasting in March 2019. However, the 2018Q4 GDP growth figure would not be available when forecasting from December 2018 to February 2019, so that a maximum of three lagged dependent variables is considered. When forecasting in October 2018 and November 2018, both 2018Q4 and 2018Q3 GDP growth figures would not be observed, so that the maximum number of lagged dependent variables is two.

We found that, in general, no individual specification of the MIDAS model in terms of the four model variants dominates in forecast performance for both factor and single indicator forecasts. Hence, for each approach, we take the average of the forecasts across all four variants of MIDAS models. Similarly, when pooling forecasts across single indicator MIDAS models, we use the simple arithmetic mean to combine forecasts throughout this study. Kuzin, Marcellino and Schumacher (2013) shows that the forecast performance of pooled MIDAS models and pooled factor MIDAS models are robust towards different weighting schemes. As a robustness check, we replace the mean forecast from single indicator MIDAS models by the median forecast. However, the results as recorded in Table C in Appendix C indicate the qualitative conclusions are the same.

On the factor-based information pooling approach, we generate forecasts from the four variants of factor MIDAS models, extracting up to three factors (*r = 1, 2 and 3*) in each case. These factors are extracted using Kalman filter estimates in the state-space factor models. Banerjee, Marcellino and Masten (2005) reveals a considerable decline in forecast performance in models when many factors are used. Using up to a maximum of six factors, Marcellino and Schumacher (2010) find that only the factor MIDAS models based on one or two factors have predictive content for German GDP. The paper also documents that differing factor estimation techniques result in only small differences in forecast performance of the factor MIDAS models.

As a measure of forecast accuracy, we compute the root mean square forecast error (RMSE) for the entire evaluation period and three subperiods.[7] Since the full evaluation period includes the occurrence of two crises, it is split into three subperiods: 2007Q4 to 2010Q2; 2010Q3 to 2019Q4; and 2020Q1 to 2020Q3 for the cycles that include the GFC, non-crisis and COVID-19 pandemic subperiods respectively. We evaluate the two approaches, namely forecast pooling versus information pooling, by comparing the RMSE obtained from the individual approaches to the RMSE from a quarterly autoregressive model. For the estimation of the autoregressive model, we use the auto-ARIMA procedure from the "forecast" package in R whereby the conditional sum-of-squares are used to find initial values, followed by maximum likelihood estimation of the model parameters. As with all other models, we replicate the publication lags of GDP, select the number of lags by BIC, and use an expanding window strategy. The maximum lag length is four, considering the release date of GDP growth in the same way as for the ADL-MIDAS and ADL-U-MIDAS models.

Since differences in forecast accuracy between two competing models may be attributed to chance, we employ a test for equal predictive accuracy proposed by Coroneo and Iacone (2020). This test modifies the Diebold and Mariano (1995) test statistics to overcome the prevalence of negative variance estimates that typically arises in smaller samples and with forecasts at longer horizons (henceforth "DM-CI test"). This is applicable to our study as we found in the computation of Diebold Mariano test statistics that our variance estimates tend to be negative for longer horizons such as at h_5 and h_6, especially when we perform the tests separately for the short crisis subperiods.

In the DM-CI test statistic, the standard rectangular kernel estimator used in the Diebold-Mariano statistic is replaced with a Daniell kernel to form the weighted periodogram estimator as follows. For a given bandwidth m, we denote the periodogram of the loss differential at time t (d_t) for the Fourier frequency $\lambda_j = \frac{2\pi j}{T}$ for $j = 0, \pm 1, \ldots, \pm m$ by

[7] For robustness checks, we use the mean absolute forecast errors instead of the root mean square forecast error but find that the qualitative conclusions are the same. These results are available from the authors upon request.

$$I(\lambda_j) = |\frac{1}{\sqrt{2\pi T}}\sum_{t=1}^{T} d_t e^{-i\lambda_j t}|^2 \text{ where } i = \sqrt{-1}$$

Then the DM-CI test statistics is given by

$$DM_{CI} = \sqrt{T}\left(\frac{\bar{d}}{\hat{V}(\bar{d})}\right) \xrightarrow{d} t_{2m} \qquad \text{where } \hat{V}(\bar{d}) = \frac{2\pi}{m}\sum_{j=1}^{m} I(\lambda_j)$$

Considering the size-power tradeoff when determining the bandwidth, we follow Harvey, Leybourne and Whitehouse (2017) to set $m = \sqrt[3]{T}$. We implement the test in R by adapting codes courtesy of Coroneo and Iacone[8] and obtained the critical values from Kiefer and Vogelsang (2005).

4. Empirical Results and Discussion

In this section, we report the forecast accuracy of the models by generating "nowcasts" and one-quarter-ahead forecasts of Singapore GDP growth. We examine the forecast accuracy of the four variants of MIDAS models, namely MIDAS, ADL-MIDAS, U-MIDAS and ADL-U-MIDAS, for every single indicator. As there is a vast number of single indicator models, we do not report their RMSEs to save space but follow Stock and Watson (2004) to present the relative RMSEs as regime graphs[9] in Appendix B. The graphical display is arranged to facilitate comparisons across the four variants of MIDAS models. It turns out that the relative performance across the four MIDAS model variants varies with indicators and forecast horizons. Following Kuzin, Marcellino and Schumacher (2013), the forecasts are pooled across the four different MIDAS model specifications for the individual indicators. These single indicator forecasts are then averaged to obtain forecasts for the forecast pooling approach.

[8] The R codes can be obtained from Laura Coroneo on her personal webpage, https://www-users.york.ac.uk/~lc1081/dm_fsa.R. The authors wish to thank Coroneo and Iacone for providing the codes.

[9] A regime graph is a scatter diagram where the horizontal and vertical axes represent the GFC and non-crisis subperiods respectively. Each point in the scatter diagram is the pair of relative RMSEs from one model and the position of the point indicates the performance of the model relative to the benchmark in the different subperiods. We do not display the regime graphs associated with the Covid-19 subperiod due to the shortness of the time period but they are available from the authors upon request.

Table 1A records the RMSE of the forecasts from the forecast pooling approach as a ratio to the RMSE of the corresponding forecasts from the benchmark autoregressive model. In this way, a number less than one indicates that the forecast pooled from the single indicator models is more accurate than the benchmark forecast. These relative RMSE figures are recorded in bold in the following tables. At the same time, the figure is marked by asterisk(s) whenever the improvement in relative RMSE is statistically significant as determined by the results of the one-sided DM-CI tests of equal predictive accuracy. A further distinction in relative forecast accuracy is made in terms of the predictions made in the non-crisis and two crises subperiods.

Table 1. Relative RMSE for Pooling Strategies
(Insert around here)

We see from Table 1A that the forecast pooling approach systematically outperforms the benchmark, though the magnitudes of the gains are not large in any subperiod. The pooled forecast strategy outperforms the benchmark model over the whole evaluation period except at the longest horizon h_6. We obtain statistically significant improvements in forecast accuracy for horizons h_2 to h_4 for the overall period as well as the GFC subperiod. For the non-crisis period, an estimated 13% gain in forecast accuracy in terms of the relative RMSE is recorded at the one-month horizon. The percentage differential declines as the horizon increases, recording 4% at the six-month horizon. We note that the gains in forecast accuracy are statistically significant for all horizons except for h_1. Meanwhile, none of the CI statistics are statistically significant in row 4 of Table 1. However, this is expected as the very small number of forecasts in the Covid-19 crisis period would lead to a large variance of the squared forecast error difference. Like the other subperiods, we note that the gains in forecast accuracy are mostly less than 10% during the pandemic crisis.

Turning to the factor-based information pooling approach, we also find that the relative performance across the four variants of factor MIDAS models vary with forecast horizons. We use regime graphs to view the relative RMSEs of the individual factor MIDAS models in Appendix B and arrange the graphical display to facilitate comparisons across MIDAS models with different number of factors. It turns out that factor MIDAS models with three factors dominate the corresponding models with one or two factors in forecast performance. This finding concurs with Chow and Choy (2009b), which identified three dynamic factors to capture the underlying movements of a similar panel dataset by applying the Bai and Ng information criteria (2002). Hence, we pool the forecasts from only those factor MIDAS models that use three factors.[10]

Table 1B records the RMSEs of the forecasts from the factor-based information pooling approach relative to those from the benchmark model. We see from Table 1B that the information pooling strategy outperforms the benchmark model, particularly during the two crisis subperiods and at the shorter horizons. Notably, the gains in forecast accuracy are substantially greater in the GFC subperiod than in the non-crisis period over all horizons, and they are statistically significant at the 5% level at h_1, h_3 and h_4. Unsurprisingly, none of the DM-CI statistics are statistically significant during the Covid-19 crisis subperiod due to the very short time period. However, we observe a huge gain in forecast accuracy of greater than 30% up to four months ahead in the pandemic subperiod. These gains in the crisis subperiods largely contributed to the outperformance of factor MIDAS over the benchmark model in the whole evaluation period, offering statistically significant differentials for all horizons except at h_4 and h_6. For the non-crisis period, we obtain improvements in forecast accuracy in the shorter horizons of up to three months ahead even though none of the gains turn out to be statistically significant.

Factor MIDAS models are extensions of single frequency dynamic factor models to the mixed frequency case. Instead of time-aggregating the monthly data to quarterly frequency, the factor MIDAS models exploits the newly-published high frequency information within the quarter into the forecasts, which could explain the improvement in forecast accuracy over the benchmark models at short horizons. Our results are broadly consistent with those of Marcellino and Schumacher

[10] Robustness checks using forecasts pooled across the three factors yield qualitatively similar results. The results are available from the authors upon request

(2010) and Kim and Swanson (2017), who also find factor MIDAS models are better than single frequency time series models for short-term forecasting of GDP.

We next perform a direct comparison of the forecast accuracy between the two pooling strategies. Table 1C records the RMSEs of the forecasts from the information pooling approach as a ratio to the RMSE of the corresponding forecasts from the forecast pooling approach. We see from Table 1C that for the GFC subperiod, differences in predictive accuracy between the two pooling strategies are statistically significant at all horizons, with the sole exception of h_2. The statistically significant gains range from 12% to 25%, with the larger ones realised at the shorter horizons. As for the Covid-19 crisis subperiod, factor MIDAS forecasts are more accurate than those from pooled single indicator models with large gains ranging from 22% to 31% at horizons of up to four months ahead and not significantly worse at other horizons. For the non-crisis period, the factor MIDAS forecasts have lower (higher) RMSEs compared to pooled single indicator forecasts at the shorter (longer) horizons. However, none of the differentials are statistically significant at conventional significance levels for this stable growth period.

To aid the discussion of results, forecast plots generated from the pooled single indicator approach (orange line) and pooled factor MIDAS approach (blue line) for the six individual horizons are displayed in Figure 1. We observe that real output growth, as represented by the grey line in the figure, plunged during the global financial crisis, and there was some volatility in the immediate aftermath of the crisis. However, real GDP growth soon steadied and remained fairly stable till the onset of the pandemic crisis. It appears the two pooling strategies have similar performance in this stable non-crisis subperiod, apart from the greater volatility in the forecasts from factor MIDAS models compared to those from the pooled single indicator models.[11] We agree with Tay (2007) that any reasonable forecasting model, including the quarterly autoregressive benchmark model, would record decent forecast performance during a stable period. Hence, the absence of substantial gains

[11] The greater volatility in the factor MIDAS forecasts leads to less systematic gains over the benchmark model which is a plausible explanation for the lack of statistical significance despite larger differentials in the relative RMSEs compared to pooled single indicator models in the non-crisis subperiod, compare Table 1A and 1B.

in forecast accuracy between the individual pooling approaches versus the benchmark model in the non-crisis subperiod (see Tables 1A and 1B).

Figure 1. Real GDP Growth Rate Forecasts vs. Actual GDP by Forecast Horizon (Insert around here)

It follows that the choice between the two approaches lies more in their predictive ability in the crisis subperiods when the economic environment is fast-evolving. We see from the forecast plots in Figure 1 that the pooled indicator forecasts are less volatile than the actual data at the individual horizons. While pooling forecasts would aid in the cancelling out of misspecification errors of single indicator models, averaging over such a large number of models could have resulted in overly-smooth forecasts that fail to sufficiently capture the dive in GDP growth during the GFC and Covid-19 crises. Similarly, Galli, Hepenstrick and Scheufele (2019) finds that forecast pooling produces less volatile forecasts that are not flexible enough to capture sudden swings in output growth in Switzerland.

In comparison, the charts show the factor MIDAS forecasts can better track rapid changes in output growth during the crises, especially for horizons h_1 to h_4. This observation is consistent with the higher forecast precision for the information pooling approach vis-a-vis the forecast pooling approach during crises, particularly at shorter horizons, as reported in Tables 1C. In other words, the visual comparisons are in agreement with the numbers. A natural question that follows is whether the factor MIDAS models would continue to do better in capturing the myriad of shocks hitting the small open economy when we have accounted for the over-smoothing of forecasts seen in Figure 1. The next section attempts to answer this question.

5. Forecast pooling from smaller datasets

Against the backdrop of overly-smooth forecasts from the forecast pooling strategy, we investigate whether using a smaller dataset could improve the predictive ability of this approach. In place of the full sample, we consider pooling forecasts from single indicator MIDAS models based on a smaller subgroup of indicators. We select from our dataset the categories of indicators that are highly relevant for driving GDP growth in Singapore based on evidence from our regime graphs for single indicator models. This approach is similar to that adopted by Banerjee, Marcellino and Masten (2005), whereby the simple average of a small number of the best performing single indicator forecasts is taken. As the Singapore economy has huge exposure to international economic developments with a trade to GDP ratio exceeding three over the past decades and it has free capital flows due to its status as an international financial centre (Wilson, 2015), external factors play a vital role in determining output growth in Singapore. Hence, we consider broad categories of external sector-related indicators when forming subgroups of predictors.

We consider a medium-sized dataset and a small dataset. The former comprises the external trade category, which is the subgroup of domestic indicators most directly affected by external developments, as well as the foreign indicators that cover multiple dimensions of the external economy. Our full sample has five categories of foreign indicators: GDP and leading indicators of major trading partners, foreign stock price data, foreign real interest rates, world electronics sales and indexes, and world prices. We observe from the regime graphs in Figure B1 that, as a group, the indicators in these six categories visually outperform the rest of the indicators in terms of their forecast accuracy even at the longer horizons.[12] In total, the medium-sized dataset has 40 indicators, which is approximately 42% of the entire dataset.

As for the small dataset, we first compare the RMSEs across the six categories of indicators in the medium-sized dataset. To this end, we obtain individual regime graphs for each separate category of indicators, as shown in Figure B3 in Appendix B. A visual comparison reveals that the two categories, external trade as well as GDP and leading indicators of major trading partners, are discernibly the better-performing ones, especially during the GFC and at shorter horizons. The high predictive content of these indicators is not surprising given the trade-dependent nature of the

[12] The RMSEs of the indicators in the six categories relative to the benchmark are represented by red dots, which appear to be lower and further left than the blue dots representing the corresponding relative RMSEs of the rest of the indicators.

Singapore economy. Hence, we select the 18 indicators in these two categories to constitute the small dataset.[13] The forecasts obtained earlier on from the single indicator MIDAS models are first averaged across the four variants of MIDAS models and then pooled separately over the medium size and the small size datasets. Table 2 records the RMSE of the pooled forecast for the two subgroups relative to the benchmark model. We include the previous results based on the full sample for ease of comparison.

Table 2. Relative RMSE for Forecast Pooling based on Subgroups
(Insert around here)

It is clear from Table 2 that the predictive ability of the forecast pooling strategy improves in both crisis periods when pooling is restricted to subgroups. Moreover, the differentials in the crisis period tend to be larger when the pooled forecasts are based on the small dataset vis-à-vis the medium-sized dataset. Pooling forecasts from the small dataset offers statistically significant gains over the benchmark model that range from 13% to 24% for horizons of up to four months ahead during the GFC. In the Covid-19 crisis, the corresponding gains in forecast precision are at least 10% at h_1, h_3 and h_5. Notably, these gains do not come at the expense of a deterioration in forecast accuracy in the non-crisis period. In general, the predictive ability of the forecast pooling strategy in the stable growth period does not appear to be only marginally affected by the sample size. We infer that the forecast pooling strategy performs best when based on the small dataset.

We did not attempt to extract factors from the small dataset as there would be too few series to do so precisely, as suggested by Boivin and Ng (2005). However, we repeat the factor MIDAS forecasting procedure using the medium-sized dataset, which produces forecasts that are less accurate than those based on the full sample.[14] Similarly, Galli, Hepenstrick and Scheufele (2019)

[13] As a robustness check, we pooled separately over the individual categories, but this produces less precise forecasts.

[14] The results are available from the authors upon request.

find dynamic factor models based on a subgroup of indicators yield inferior forecasting results in the case of Switzerland. Hence, we compare the forecasts pooled from single indicator models based on the small dataset with the forecasts from factor MIDAS models based on the full sample.

Table 1D records the RMSEs of the forecasts from the information pooling approach as a ratio to the RMSEs of the corresponding forecasts from the small subgroup forecast pooling approach. We observe from the table that the forecast performance of factor MIDAS still dominates the small subgroup pooled single indicator models at the shorter horizons during crises, albeit by a reduced margin of around 15%. The gains in predictive accuracy in the GFC subperiod are statistically significant at h_1 and h_4. Meanwhile, large differentials of 20% to 30% are recorded when predicting up to four months ahead in the Covid-19 pandemic crisis. Overall, the factor MIDAS model remains superior to forecast pooling during crisis periods.[15]

To better understand the results, we generate the forecast plots from the small subgroup pooled single indicator approach (orange line) and the factor MIDAS approach (blue line) for the six individual horizons, as displayed in Figure 2. A visual comparison of Figures 1 and 2 reveals that the small subgroup pooled indicator strategy delivers more flexible forecasts vis-à-vis those based on the full sample at the shorter horizons during crises, especially in the GFC subperiod. It appears the over-smoothing of forecasts based on the full sample of indicators is partly ameliorated by averaging over a smaller number of forecasts which enables the pooled forecast to track rapid changes in output growth more closely. This finding is consistent with the improvements we obtained from forecast pooling based on the small subgroup (see Table 2).

Figure 2. Real GDP Growth Rate Forecasts vs. Actual GDP by Forecast Horizon
(Insert around here)

[15] As a robustness check, the relative RMSEs for forecast pooling based on the median forecast against the benchmark model and the factor MIDAS model are recorded in Table C in Appendix C. We obtain qualitatively similar results as in the case of the mean forecast.

In an exceptionally open economy like Singapore, valuable predictive content is concentrated in relatively few foreign-related indicators. Consequently, forecast pooling over a small number of trade-related indicators as well as foreign GDP and composite leading indexes could outperform the quarterly benchmark model in all subperiods as recorded in Table 2. We infer that, in the context of a very small open economy, the subgroup forecast pooling strategy is a quick method to use monthly data for generating short-term quarterly GDP growth predictions, and is particularly useful when practitioners have difficulty maintaining and regularly updating a large dataset.

However, a visual inspection of the charts in Figure 2 that compares the forecast performance across the two pooling strategies suggests the factor MIDAS forecasts are still better able to track the swings in output growth during both crises for horizons of up to four months ahead. This observation is consistent with the numbers in Table 1D, which indicate the information pooling strategy retains its superior short-term predictive ability in crisis subperiods. Our findings suggest that extracting the underlying movements from many monthly indicators into a few factors is a more effective way to capture volatile short-term fluctuations in output growth. Incorporating higher frequency and most recently published information in factor MIDAS models seems to produce a better assessment of current quarter economic growth during crises.

6. Conclusion

It is well recognised that the publication lag of quarterly GDP growth hampers the early assessment of the current and near-term economic environment. Hence, the use of higher frequency such as monthly indicators provides more timely information on economic fluctuations, particularly in fast-evolving conditions such as in a crisis. In this paper, we evaluate forecast pooling across a large set of single indicator MIDAS models versus pooling information from indicators into factor MIDAS models to predict Singapore GDP growth. Both pooling approaches aim at extracting the predictive content from our large scale ragged-edge dataset that spans the GFC, the Covid-19 pandemic crisis and the non-crisis period in between. Since Singapore’s output growth was fairly stable in the non-crisis period, any reasonable predictive model would record a decent forecast performance in this period. Indeed, we found that neither pooling strategy records any substantive improvements to the forecast performance of the quarterly autoregressive benchmark model in the non-crisis period. It follows that the choice between the two pooling strategies lies more in their predictive ability during crises.

Our results point to the information pooling strategy dominating both the benchmark model and the forecast pooling strategy by wide margins in the GFC and Covid-19 pandemic crisis. As pooling forecasts over the entire set of single indicator models leads to over-smoothing and thus the inability of pooled forecasts to capture sudden movements in the GDP growth series during crises, we find that pooling forecasts over a small subgroup of indicators increases the predictive accuracy of the forecast pooling approach. Given the small open nature of the Singapore economy, forecast pooling based on only a small subgroup of indicators related to the external sector beats the benchmark model and thus offers an alternative method to incorporate timely predictive information without the need to maintain and update a large dataset. However, even after accounting for over-smoothing by pooling forecasts across a smaller number of single indicator models, the factor MIDAS models retain their superior short-term predictive ability compared to the forecast pooling strategy. These findings suggest the information pooling strategy is better suited to capture the myriad of shocks hitting the small open economy in periods of wide economic fluctuations.

Acknowledgements

The authors would like to thank participants of the China Meeting of the Econometrics Society held in July 2021 for helpful comments and suggestions, and Tze Yong Yew for excellent research assistance.

References

Abeysinghe, T. (1998). Forecasting Singapore’s quarterly GDP with monthly external trade. International Journal of Forecasting, 14, 505–513.

Armesto, M. T., Engemann, K. M., & Owyang, M. T. (2010). Forecasting with mixed frequencies. Federal Reserve Bank of St. Louis Review, 92 (6), 521-536.

Andreou, E., Ghysels, E. and Kourtellos, A. (2010). Regression models with mixed sampling frequencies, Journal of Econometrics, 158(2), 31-53.

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

Bai, J. and Ng, S. (2002). Determining the number of factors in approximate factor models. *Econometrica* 70, 191–221.

Banbura & Runstler (2011). A look into the factor model black box: Publication lags and the role of hard and soft data in forecasting GDP. International Journal of Forecasting, 27, 333-346.

Banerjee, A., Marcellino, M., & Masten, I. (2005). Leading indicators for euro-area inflation and GDP growth. Oxford Bulletin of Economics and Statistics, 67 (s1), 785-813.

Bates, J. M., & Granger, C. W. J. (1969). The combination of forecasts. Operational Research Quarterly, 20(4), 451–468.

Boivin, J., & Ng, S. (2005). Understanding and comparing factor-based forecasts. International Journal of Central Banking, 1 (3), 117-151.

Breitung, J., & Schumacher, C. (2008). Real-time forecasting of German GDP based on a large factor model with monthly and quarterly data. International Journal of Forecasting, 24 (3), 386-398.

Chow, H. K., & Choy, K. M. (2009a). Analysing and forecasting business cycles in a small open economy: A dynamic factor model for Singapore. OECD Journal: Journal of Business Cycle Measurement and Analysis, 2009 (1), 19-41.

Chow, H. K., & Choy, K. M. (2009b). Monetary policy and asset prices in a small open economy: A factor-augmented VAR analysis for Singapore. Annals of Financial Economics, 2009 (5), 51-78.

Clements, M. P., & Galvao, A. B. (2008). Macroeconomic forecasting with mixed-frequency data. Journal of Business & Economic Statistics, 26 (4), 546-554.

Coroneo, L., & Iacone, F. (2020). Comparing predictive accuracy in small samples using fixed-smoothing asymptotics. Journal of Applied Econometrics, 35, 391–409.

Diebold, F. X., & Mariano, R. S. (1995). Comparing predictive accuracy. Journal of Business and Economic Statistics, 13 (3), 253-263.

Doz, C., Giannone, D., & Reichlin, L. (2011). A two-step estimator for large approximate dynamic factor models based on Kalman filtering. Journal of Econometrics, 164 (1),188-205.

Forni, M., Hallin, M., Lippi, M., & Reichlin, L. (2003). Do financial variables help forecasting inflation and real activity in the euro area? Journal of Monetary Economics, 50 (6), 1243-1255.

Foroni, C., & Marcellino, M. (2014). A comparison of mixed frequency approaches for nowcasting euro area macroeconomic aggregates. International Journal of Forecasting, 30 (3), 554-568.

Foroni, C., Marcellino, M., & Schumacher, C. (2015). Unrestricted MIDAS: MIDAS regressions with unrestricted lag polynomials. Journal of the Royal Statistical Society: Series A (Statistics in Society), 178 (1), 57-82.

Galli, A., Hepenstrick, C., & Scheufele, R. (2019). Mixed-frequency models for tracking short-term economic developments in Switzerland. International Journal of Central Banking, 15 (2), 151-178.

Ghysels, E., Kvedaras, V., & Zemlys, V. (2016). Mixed frequency data sampling regression models: The R package midasr. Journal of Statistical Software, 72 (4), 1-35.

Ghysels, E., Sinko, A., & Valkanov, R. (2007). MIDAS regressions: Further results and new directions. Econometric Reviews, 26 (1), 53-90.

Harvey, D., Leybourne, S., & Whitehouse, E. (2017). Forecast evaluation tests and negative long-run variance estimates in small samples. International Journal of Forecasting, 33, 833-847.

Heinisch, K, & Scheufele, R. (2018). Bottom up or direct? Forecasting German GDP in a data-rich environment, Empirical Economics, 54, 705–745

Hueng, C. J., & Yau, R. (2019). Nowcasting GDP growth for small open economies with a mixed frequency structural model. Computational Economics, 54 (1), 177-198.

Kiefer,N. M., & Vogelsang, T. J. (2005). A new asymptotic theory for heteroskedasticity–autocorrelation robust tests. *Econometric Theory*, *21*(6), 1130–1164.

Kim, H. H., & Swanson, N. R. (2017). Methods for backcasting, nowcasting and forecasting using factor MIDAS: With an application to Korean GDP. Journal of Forecasting, 37 (3), 281-302.

Kuzin, V., Marcellino, M., & Schumacher, C. (2013). Pooling versus model selection for nowcasting GDP with many predictors: Empirical evidence for six industrialised countries. Journal of Applied Econometrics, 28 (3), 392-411.

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

Marcellino, M., Stock, J. H., & Watson, M. W. (2006). A comparison of direct and iterated multistep AR methods for forecasting macroeconomic time series. Journal of Econometrics, 135 (1), 499-526.

Stock, J. H., & Watson, M. W. (2002). Macroeconomic forecasting using dffusion indexes. Journal of Business and Economic Statistics, 20 (2), 147-162.

Stock, J. H., & Watson, M. W. (2004). Combination forecasts of output growth in a seven country data set. Journal of Forecasting, 23 (6), 405-430.

Tay, A.S. (2007). Financial variables as predictors of real output growth. Working paper 7-2007, Singapore Management University.

Timmermann, A. (2006). Forecast combinations. In G. Elliott, C. W. J. Granger, & A. Timmermann (Eds.), Handbook of economic forecasting, Vol. 1 (pp. 135–196). Amsterdam, The Netherlands: North-Holland.

Tsui, A. K., Xu, C. Y., & Zhang, Z. (2018). Macroeconomic forecasting with mixed data sampling frequencies: Evidence from a small open economy. Journal of Forecasting, 37 (6), 666-675.

Wilson, P. (2015). Monetary policy and financial sector development. Singapore Economic Review, 60 (3), 1550031 (25 pages).

Appendix A: Data Listing

Data	Source	Transformation	Publication lag
0. Real GDP (quarterly)	CEIC	Δln	2 months
Foreign GDP indices/composite leading indicators (6)			

1.	Major 5 Asia GDP index	FRED	Δln	4 months
2.	US CLI	OECD	Deviations from trend	1 month
3.	Japan CLI	OECD	Deviations from trend	1 month
4.	UK CLI	OECD	Deviations from trend	1 month
5.	Major 5 Asian CLI	OECD	Deviations from trend	1 month
6.	4 Big European	OECD	Deviations from trend	1 month
Foreign stock prices (11)				
7.	NASDAQ Composite Index (US)	CEIC	Δln	None
8.	Nikkei 225 (Japan)	CEIC	Δln	None
9.	DAX (Germany)	CEIC	Δln	None
10.	FTSE100 (UK)	CEIC	Δln	None
11.	KOSPI (Korea)	CEIC	Δln	None
12.	FTSE Bursa (Malaysia)	CEIC	Δln	None
13.	JKSE (Indonesia)	CEIC	Δln	None
14.	SET Index (Thailand)	CEIC	Δln	None
15.	PSEI (Philippines)	CEIC	Δln	None
16.	Shanghai Composite Index (China)	CEIC	Δln	None
17.	TSWE (Taiwan)	CEIC	Δln	None
Foreign real interest rates (3)				
18.	US (3-month treasury bill rate–CPI inflation)	CEIC	Δ	1 month
19.	Japan (3-month TIBOR–CPI inflation)	CEIC	Δ	1 month
20.	UK (3-month LIBOR–CPI inflation)	CEIC	Δ	1 month
World electronics (5)				
21.	Global semiconductor sales*	CEIC	Δln	2 months
22.	US new orders for electronics (excl. semiconductors)	US Census Bureau	Δln	2 months
23.	US electronics shipments-to-inventories ratio for electronics	US Census Bureau	Δln	2 months
24.	PPI for US electronics	FRED	None	1 month
World prices (4)				
25.	OPEC Crude oil price	CEIC	Δln	1 month
26.	Global food price index*	FRED	Δln	1 month
27.	Global non-fuel price index	FRED	Δln	1 month
28.	Global commodity price index	FRED	Δln	1 month
Industrial production (7)				
29.	Industrial production index (IPI)	Department of Statistics, Singapore (DOS)	Δln	1 month

30. IPI: Biomedicals	DOS	$\Delta \ln$	1 month
31. IPI: Transport engineering	DOS	$\Delta \ln$	1 month
32. IPI: Precision engineering	DOS	$\Delta \ln$	1 month
33. IPI: General manufacturing	DOS	$\Delta \ln$	1 month
34. IPI: Electronics	DOS	$\Delta \ln$	1 month
35. IPI: Chemicals	DOS	$\Delta \ln$	1 month
Business surveys (3)			
36. General manufacturing expectations	CEIC	None	1 month
37. Manufacturing: stocks of finished goods	CEIC	None	1 month
38. Manufacturing: new orders arrived	CEIC	None	1 month
Sectoral Indicators (11)			
39. Retail sales index	CEIC	$\Delta \ln$	2 months
40. Retail sales value*	CEIC	$\Delta \ln$	2 months
41. Car registrations Above 1600cc*	CEIC	$\Delta \ln$	1 month
42. Car registrations Below 1600cc*	CEIC	$\Delta \ln$	1 month
43. Visitor arrivals	CEIC	$\Delta \ln$	1 month
44. Air cargo loaded*	CEIC	$\Delta \ln$	1 month
45. Air cargo discharged*	CEIC	$\Delta \ln$	1 month
46. Sea cargo handled*	CEIC	$\Delta \ln$	1 month
47. Electricity generation*	CEIC	$\Delta \ln$	2 months
48. Formation of companies*	CEIC	$\Delta \ln$	1 month
49. Construction contracts awarded (BCA)	CEIC	None	2 months
External Trade (12)			
50. Total imports	CEIC	$\Delta \ln$	1 month
51. Imports: non-oil	CEIC	$\Delta \ln$	1 month
52. Imports: oil	CEIC	$\Delta \ln$	1 month
53. Total exports	CEIC	$\Delta \ln$	1 month
54. Exports: non-oil	CEIC	$\Delta \ln$	1 month
55. Exports: oil	CEIC	$\Delta \ln$	1 month
56. Domestic exports	CEIC	$\Delta \ln$	1 month
57. Domestic exports: non-oil	CEIC	$\Delta \ln$	1 month
58. Domestic exports: oil	CEIC	$\Delta \ln$	1 month
59. Re-exports	CEIC	$\Delta \ln$	1 month
60. Re-exports: non-oil	CEIC	$\Delta \ln$	1 month
61. Re-exports: oil	CEIC	$\Delta \ln$	1 month
Price Indices (7)			
62. Export price index	CEIC	$\Delta \ln$	1 month
63. Export price index: non-oil	CEIC	$\Delta \ln$	1 month

64. Import price index	CEIC	Δln	1 month
65. Import price index: non-oil	CEIC	Δln	1 month
66. Consumer price index	CEIC	Δln	1 month
67. Domestic supply price index	CEIC	Δln	1 month
68. Manufactured products price index	CEIC	Δln	1 month
Financial (18)			
69. Straits Times Index	CEIC	Δln	None
70. 1-year treasury bill yield	CEIC	Δ	None
71. 2-year treasury bill yield	CEIC	Δ	None
72. 5-year treasury bill yield	CEIC	Δ	None
73. 10-year treasury bill yield	CEIC	Δ	None
74. 3-month SIBOR	CEIC	Δ	1 month
75. Yield spread# (10-year minus 3-month SIBOR)	CEIC	None	1 month
76. Singapore dollar to Australian dollar	CEIC	Δln	None
77. Singapore dollar to Pound	CEIC	Δln	None
78. Singapore dollar to Renminbi	CEIC	Δln	None
79. Singapore dollar to Euro	CEIC	Δln	None
80. Singapore dollar to Hong Kong dollar	CEIC	Δln	None
81. Singapore dollar to Yen	CEIC	Δln	None
82. Singapore dollar to Ringgit	CEIC	Δln	None
83. Singapore dollar to US dollar	CEIC	Δln	None
84. Singapore dollar to Franc	CEIC	Δln	None
85. Nominal effective exchange rate	CEIC	Δln	1 month
86. Real effective exchange rate	CEIC	Δln	1 month
Monetary (9)			
87. M1	CEIC	Δln	1 month
88. M2	CEIC	Δln	1 month
89. M3	CEIC	Δln	1 month
90. Loans & advances	CEIC	Δln	1 month
91. Loans & advances: manufacturing	CEIC	Δln	1 month
92. Loans & advances: construction	CEIC	Δln	1 month
93. Loans & advances: commerce*	CEIC	Δln	1 month
94. Loans & advances: financial institutions	CEIC	Δln	1 month
95. Loans & advances: individuals	CEIC	Δln	1 month

Note: Series marked with asterisk `*' are seasonally adjusted via the X-13 ARIMA procedure. Yield spreads marked with '#' are calculated using 10-year treasury bill yield at the long end of the term structure instead of higher maturity treasury bills due to data availability issue.

Appendix B: Regime Graphs

We follow Stock and Watson (2004) to represent the relative RMSE results using a graphical method. For each horizon, we have one scatter diagram where the horizontal and vertical axes represent the GFC cycle and non-crisis subperiods respectively. Each point in the scatter diagram is the pair of relative RMSEs from one model. If a point lies in the third (first) quadrant so that the relative RMSEs for both periods falls below (above) one, then the model outperforms (underperforms) the benchmark in both subperiods. By contrast, a point lying in the second or fourth quadrant means the model outperforms the benchmark in one subperiod only. For all three figures B1, B2 and B3, we arrange the scatter plots by rows with each row representing one forecast horizon. For ease of comparison, the columns in Figure B1 represent the four variants of single indicator MIDAS models, the columns in Figure B2 represent factor MIDAS models with different number of factors, and the columns in Figure B3 represent the categories of foreign-related indicators for single indicator models.

Insert Figures B1, B2 and B3

Figure B1. Relative RMSEs for GFC vs Non-crisis by Horizon and Variant of Single MIDAS Model

Figure B2. Relative RMSEs for GFC vs Non-crisis by Horizon and Number of Factor in Factor MIDAS Model

Figure B3. Relative RMSEs for GFC vs Non-crisis by Horizon and Indicator Category of Single MIDAS Model

Appendix C: Forecast Pooling using Median Forecast

Table C. Relative RMSE for Forecast Pooling based on Median Forecast

Figure 1. Real GDP Growth Rate Forecasts vs. Actual GDP by Forecast Horizon

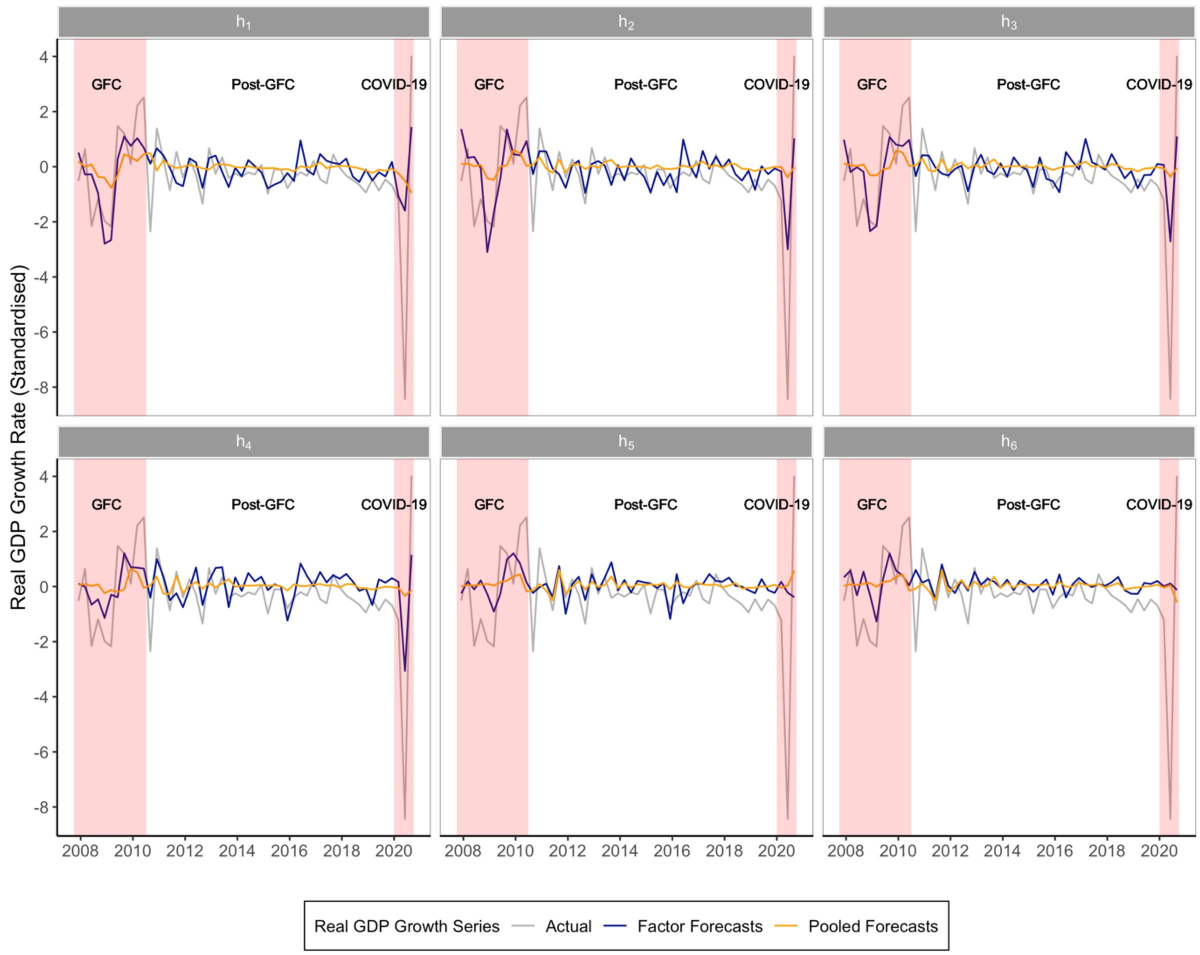

Notes: Pooled forecasts (across all single indicators) and factor MIDAS forecasts (from three-factor models and with factors estimated from all single indicators) of real GDP growth (standardised) averaged across all MIDAS methods from 2007Q4 to 2020Q3 using all information up to the forecast origin are denoted in orange and blue respectively. Actual GDP growth is denoted in grey. Crisis windows are highlighted in red. h_1, h_2, h_3, h_4, h_5 and h_6 denote the forecast horizons. All forecasts are produced using a direct forecasting with expanding window strategy with re-selection of lags and model coefficients at each forecast origin. The ragged edge structure is taken into account by imposing the same number of missing values observed at the end of sample for each recursion.

Figure 2. Real GDP Growth Rate Forecasts vs. Actual GDP by Forecast Horizon

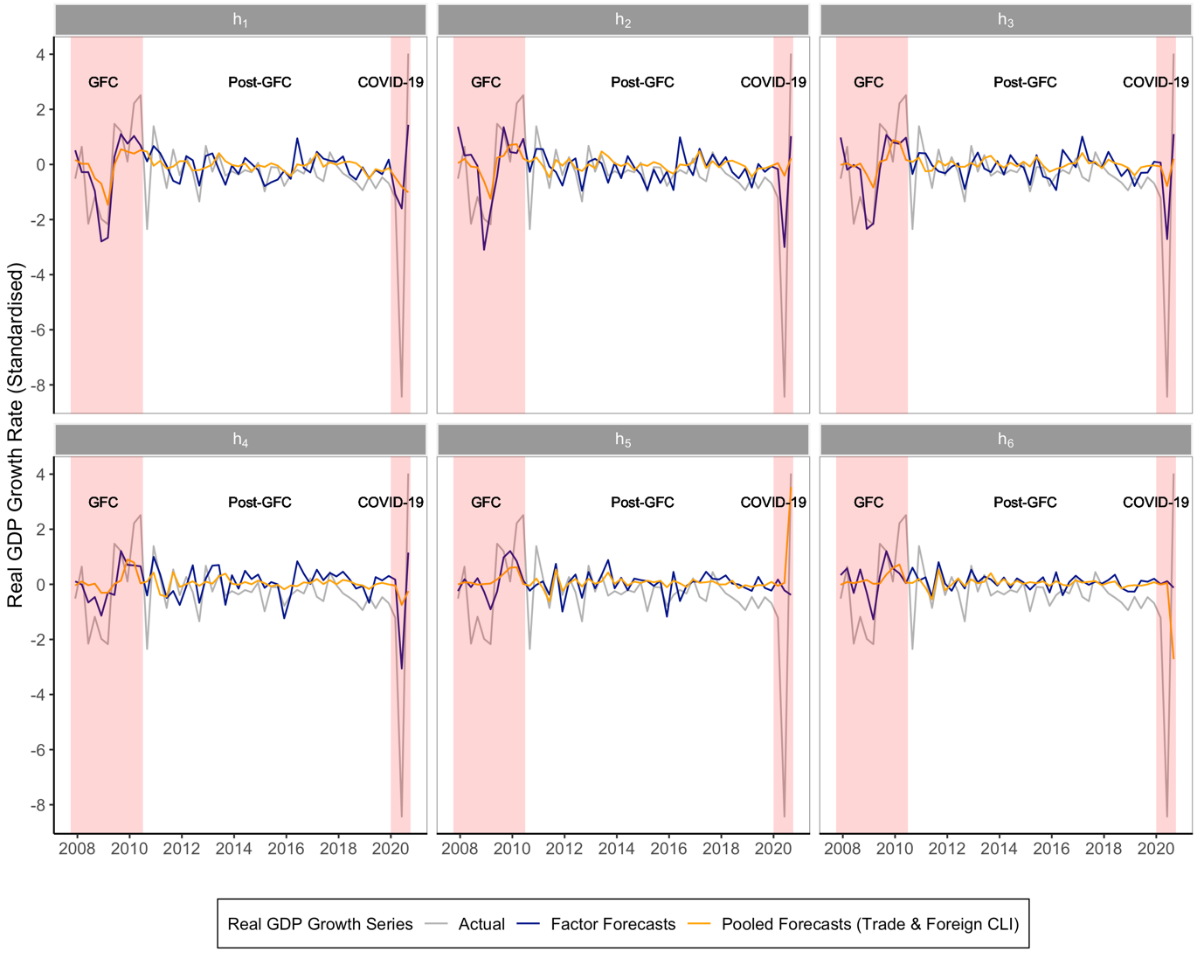

Notes: Pooled forecasts (across foreign CLI and trade indicators) and factor MIDAS forecasts (from three-factor models and with factors estimated from all single indicators) of real GDP growth (standardised) averaged across all MIDAS methods from 2007Q4 to 2020Q3 using all information up to the forecast origin are denoted in orange and blue respectively. Actual GDP growth is denoted in grey. Crisis windows are highlighted in red. h_1, h_2, h_3, h_4, h_5 and h_6 denote the forecast horizons. All forecasts are produced using a direct forecasting with expanding window strategy with re-selection of lags and model coefficients at each forecast origin. The ragged edge structure is taken into account by imposing the same number of missing values observed at the end of sample for each recursion.

Figure B1. Relative RMSEs for GFC vs Non-crisis by Horizon and Variant of Single MIDAS Model

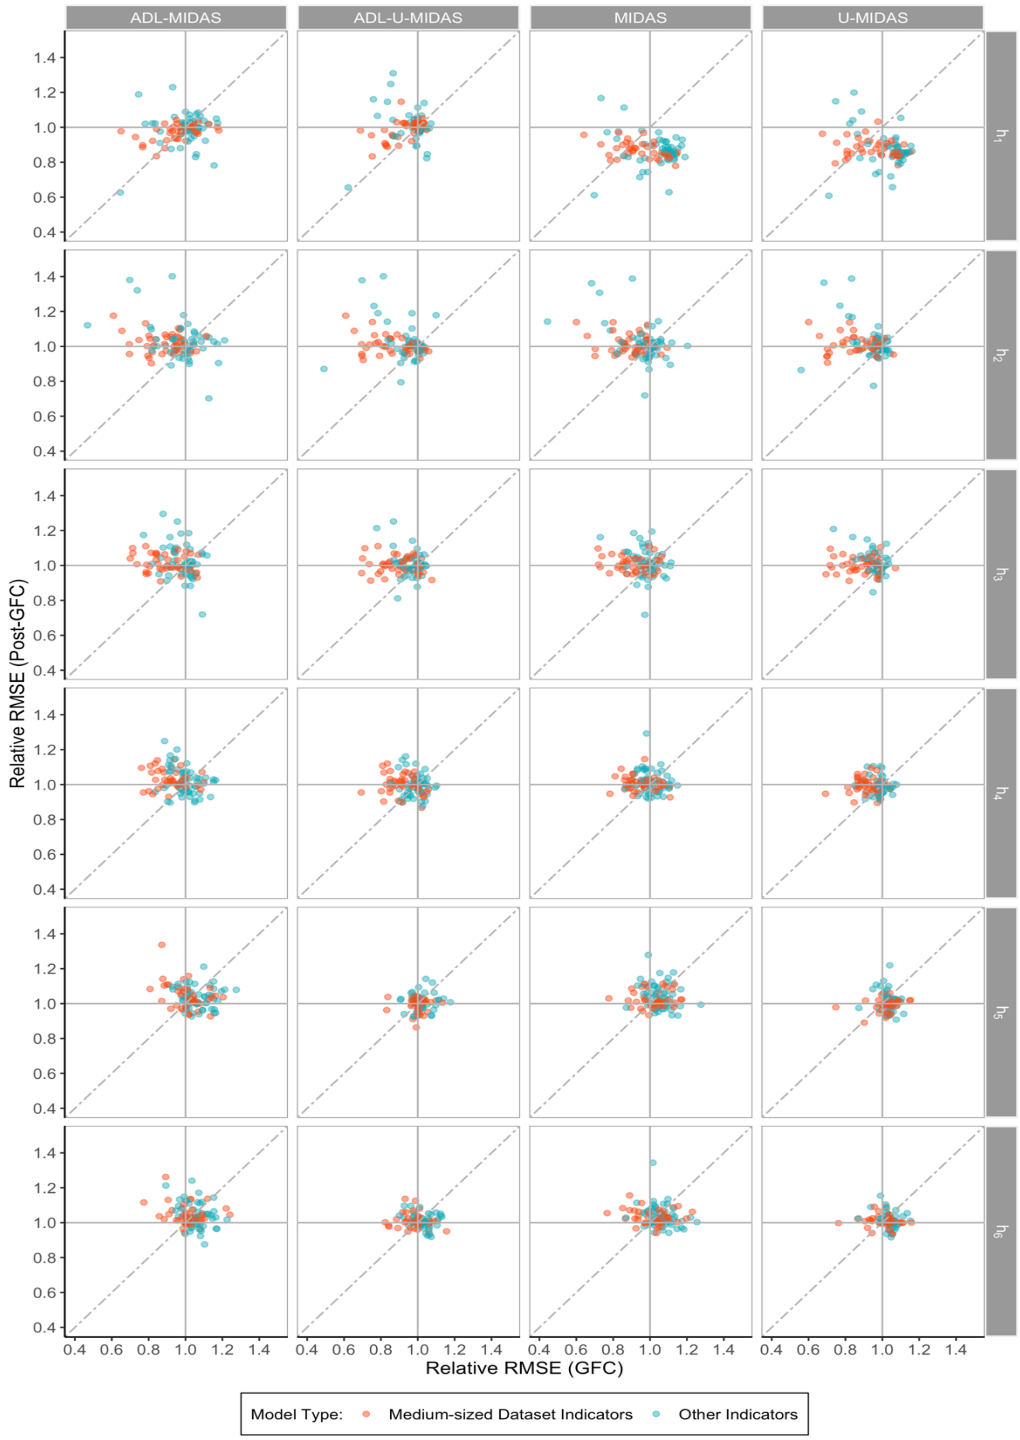

Notes: The x-coordinates represent the relative RMSEs of the forecasts to the AR benchmark for the Global Financial Crisis (2007Q4 to 2010Q2). The y-coordinates represent the relative RMSEs of the forecasts to the AR benchmark for the non-crisis period (2010Q3 to 2019Q4). h_1, h_2, h_3, h_4, h_5 and h_6 denote the forecast horizons. All forecasts are produced using a direct forecasting with expanding window strategy with re-selection of lags and model coefficients at each forecast origin. The ragged edge structure is taken into account by imposing the same number of missing values observed at the end of sample for each recursion.

Figure B2. Relative RMSEs for GFC vs Non-crisis by Horizon and Number of Factor in Factor MIDAS Model

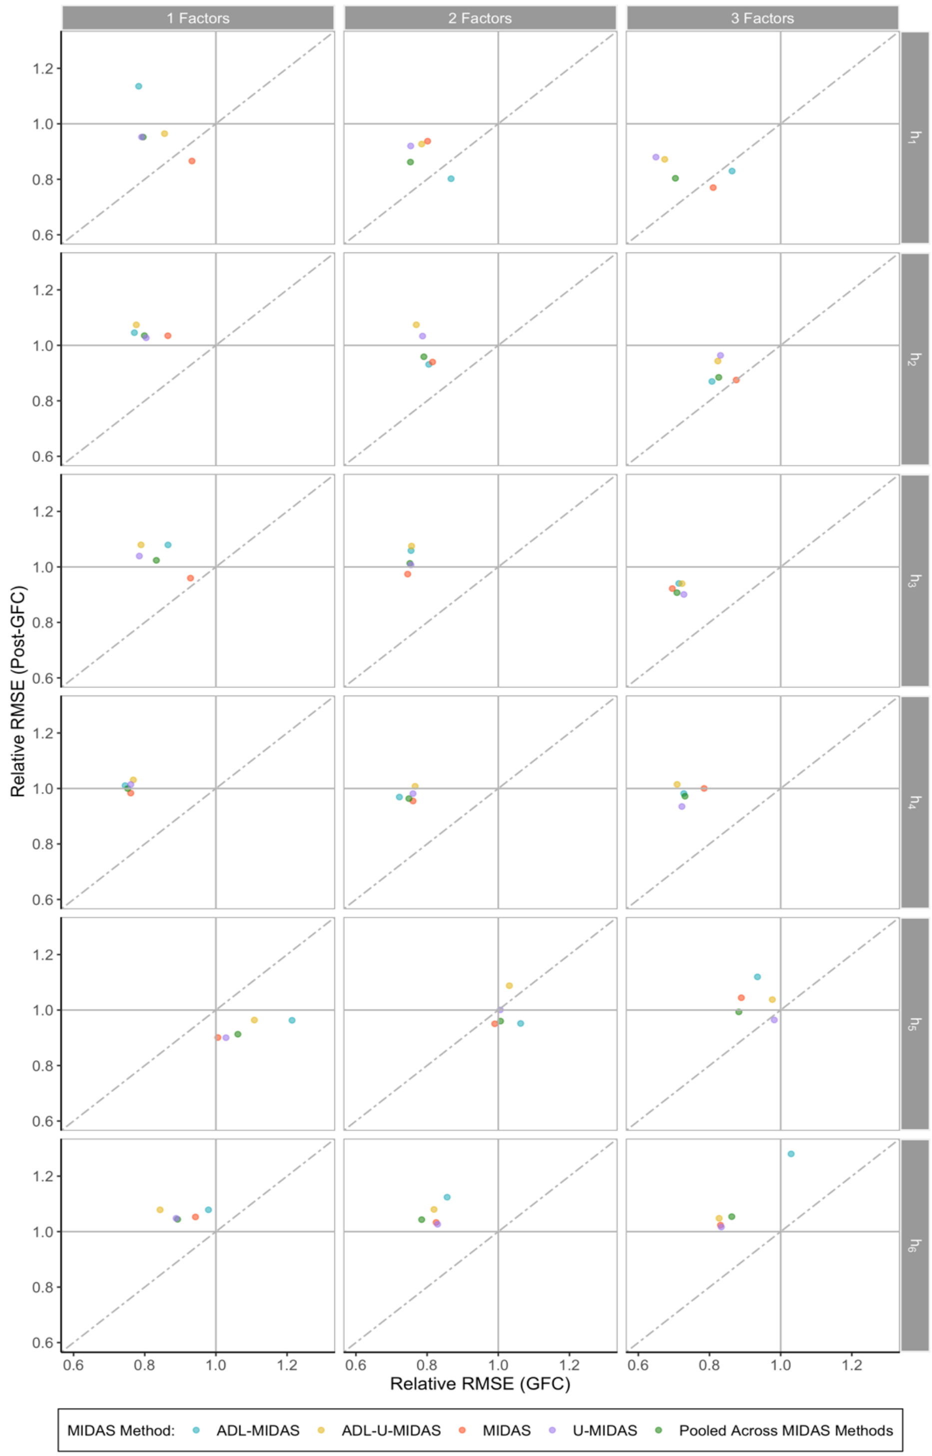

Notes: The x-coordinates represent the relative RMSEs of the forecasts to the AR benchmark for the Global Financial Crisis (2007Q4 to 2010Q2). The y-coordinates represent the relative RMSEs of the forecasts to the AR benchmark for the non-crisis period (2010Q3 to 2019Q4). h_1, h_2, h_3, h_4, h_5 and h_6 denote the forecast horizons. All forecasts are produced using a direct forecasting with expanding window strategy with re-selection of lags and model coefficients at each forecast origin. The ragged edge structure is taken into account by imposing the same number of missing values observed at the end of sample for each recursion.

Figure B3. Relative RMSEs for GFC vs Non-crisis by Horizon and Indicator Category of Single MIDAS Model

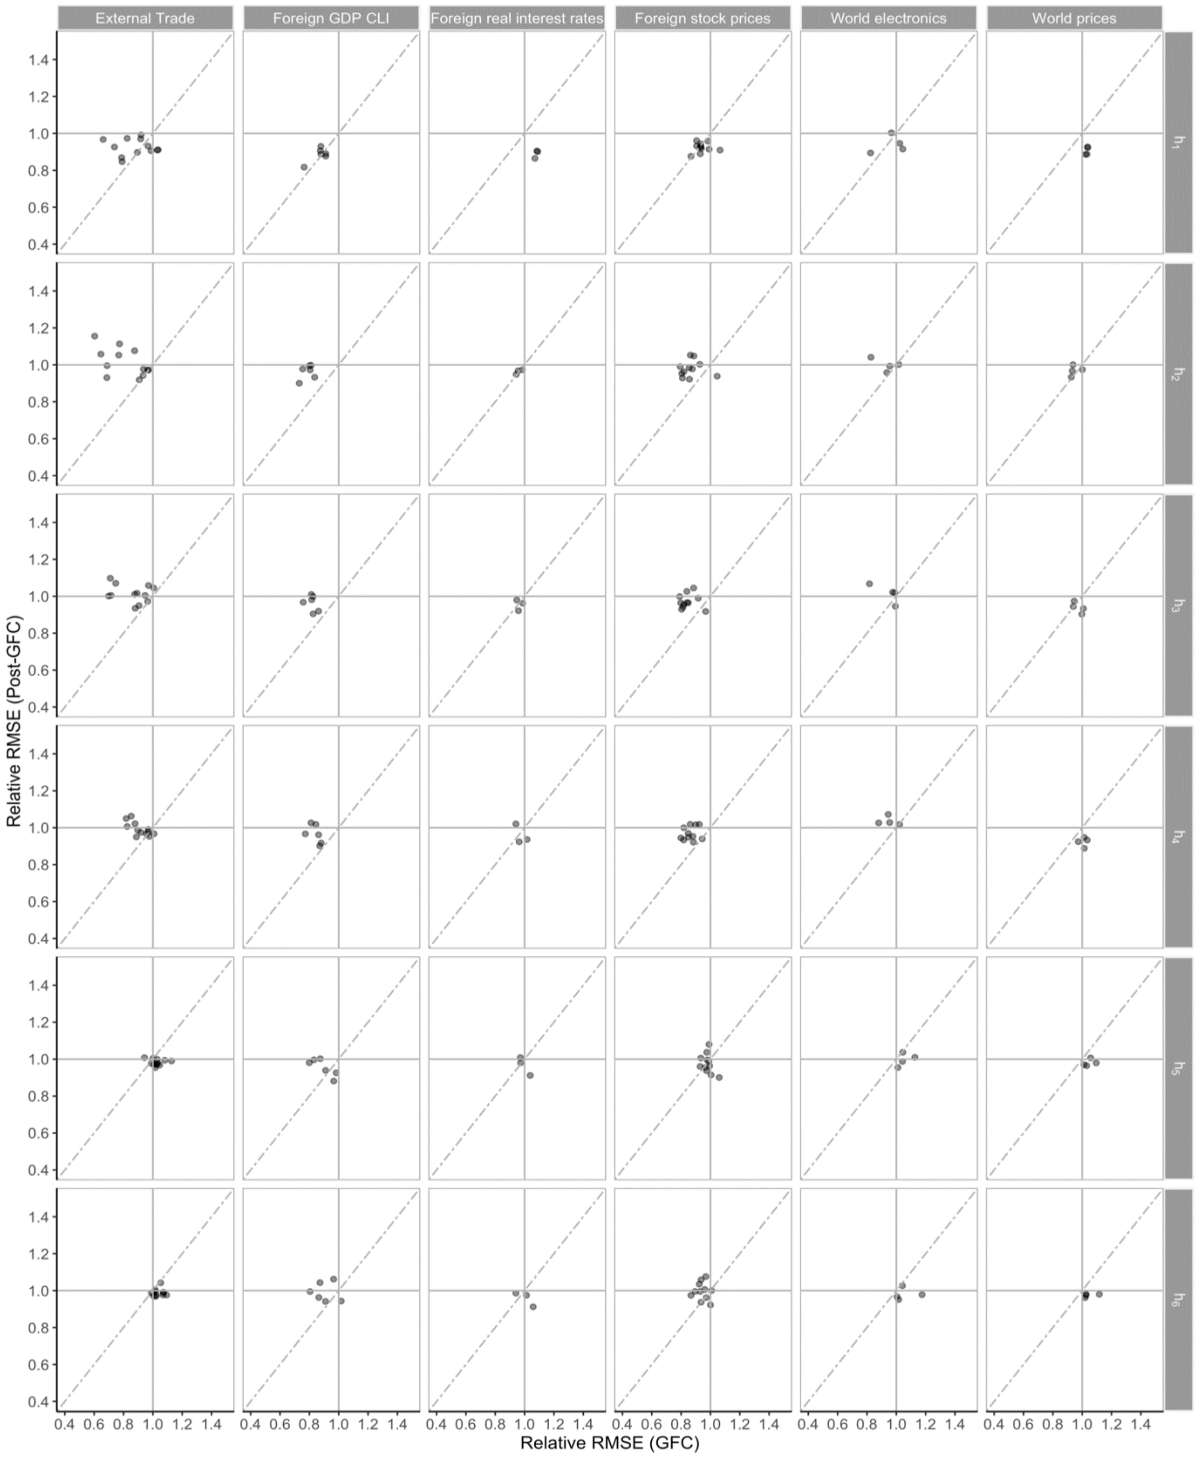

Notes: The x-coordinates represent the relative RMSEs of the forecasts to the AR benchmark for the Global Financial Crisis (2007Q4 to 2010Q2). The y-coordinates represent the relative RMSEs of the forecasts to the AR benchmark for the non-crisis period (2010Q3 to 2019Q4). h_1, h_2, h_3, h_4, h_5 and h_6 denote the forecast horizons. All forecasts are produced using a direct forecasting with expanding window strategy with re-selection of lags and model coefficients at each forecast origin. The ragged edge structure is taken into account by imposing the same number of missing values observed at the end of sample for each recursion.

Table 1. Relative RMSE for Pooling Strategies

Sample Period	Horizon (Months)					
	h$_1$	*h*$_2$	*h*$_3$	*h*$_4$	*h*$_5$	*h*$_6$
1A. Forecast Pooling vs Benchmark						
Whole: 2007Q4 to 2020Q3	**0.89**	**0.93****	**0.94****	**0.95****	**0.97**	1.01
Non-crisis: 2010Q3 to 2019Q4	**0.87**	**0.92*****	**0.93*****	**0.93*****	**0.95****	**0.96***
GFC: 2007Q4 to 2010Q2	**0.94**	**0.88****	**0.90****	**0.93****	1.00	1.00
Covid-19: 2020Q1 to 2020Q3	**0.88**	**0.95**	**0.95**	**0.96**	**0.96**	1.02
1B. Information Pooling vs Benchmark						
Whole: 2007Q4 to 2020Q3	**0.71***	**0.73***	**0.72***	**0.72**	**0.97***	**0.98**
Non-crisis: 2010Q3 to 2019Q4	**0.80**	**0.88**	**0.91**	**0.97**	**0.99**	1.05
GFC: 2007Q4 to 2010Q2	**0.70****	**0.83**	**0.71****	**0.73*****	**0.88**	**0.86**
Covid-19: 2020Q1 to 2020Q3	**0.69**	**0.66**	**0.69**	**0.65**	**0.98**	1.00
1C. Information Pooling vs Forecast Pooling (all indicators)						
Whole: 2007Q4 to 2020Q3	**0.80**	**0.79**	**0.77**	**0.76**	1.00	**0.97**
Non-crisis: 2010Q3 to 2020Q2	**0.93**	**0.96**	**0.98**	1.05	1.04	1.10
GFC: 2007Q4 to 2010Q2	**0.75****	**0.94**	**0.79***	**0.78****	**0.88***	**0.86***
Covid-19: 2020Q1 to 2020Q3	**0.78***	**0.69**	**0.72**	**0.68**	1.02	**0.98**
1D. Information Pooling vs Forecast Pooling (18 indicators)						
Whole: 2007Q4 to 2020Q3	**0.83**	**0.82**	**0.81**	**0.79**	1.05	**0.90**
Non-crisis: 2009Q2 to 2019Q4	**0.96**	**0.97**	**0.97**	1.03	1.04	1.09
GFC: 2007Q4 to 2020Q2	**0.84****	1.09	**0.87**	**0.84***	**0.92**	**0.89**
Covid-19: 2020Q1 to 2020Q3	**0.80***	**0.70**	**0.76**	**0.70**	1.10	**0.88**

Note: `***', `**' and `*' indicate that the differences are significant at the 1%, 5% and 10% levels respectively according to the Coroneo-Iacone one-sided test. Figures in bold in panels A and B indicate the pooling strategy outperform the AR benchmark (relative RMSE less than one), whilst those in panel C and D indicate the forecasts from the factor MIDAS models outperform the pooled forecasts from single indicator MIDAS models. All forecasts, for both the information and forecast pooling approach, are pooled across four MIDAS variants.

Table 2. Relative RMSE for Forecast Pooling over Subgroups

Sample Period	(Number of Indicators)	Horizon (Months)					
		h_1	h_2	h_3	h_4	h_5	h_6
Whole	(95 indicators)	**0.89**	**0.93****	**0.94****	**0.95****	**0.97**	1.01
	(40 indicators)	**0.88**	**0.91****	**0.90***	**0.92***	**0.94**	1.03
	(18 indicators)	**0.85***	**0.90****	**0.89***	**0.92***	**0.92**	1.08
Non-crisis	(95 indicators)	**0.87**	**0.92*****	**0.93*****	**0.93*****	**0.95****	**0.96***
	(40 indicators)	**0.86***	**0.91*****	**0.93*****	**0.94*****	**0.96***	**0.97**
	(18 indicators)	**0.84***	**0.91*****	**0.94****	**0.95****	**0.95****	**0.97**
GFC	(95 indicators)	**0.94**	**0.88****	**0.90****	**0.93****	1.00	1.00
	(40 indicators)	**0.90**	**0.83****	**0.85****	**0.89*****	**0.98**	**0.98**
	(18 indicators)	**0.84***	**0.76****	**0.82****	**0.87****	**0.96**	**0.97**
Covid-19	(95 indicators)	**0.88**	**0.95**	**0.95**	**0.96**	**0.96**	1.02
	(40 indicators)	**0.87**	**0.93**	**0.92**	**0.93**	**0.93**	1.05
	(18 indicators)	**0.86**	**0.94**	**0.90**	**0.93**	**0.90**	1.14

Note: `***', `**' and `*' indicate that the differences are significant at the 1%, 5% and 10% levels respectively according to the Coroneo-Iacone one-sided test. Figures in bold indicate the forecasts from the factor MIDAS models outperform the pooled forecasts from single indicator MIDAS models. All forecasts (for both the information and forecast pooling approach) are pooled across MIDAS variants. The medium-sized dataset comprising 40 of the 95 indicators contains five foreign indicator categories: GDP and leading indicators of major trading partners, foreign financial data, world electronics sales and indexes, external trade, and world prices. The small dataset comprising 18 indicators include only the external trade indicators, foreign GDP and foreign composite leading indicators.

Table C. Relative RMSE for Pooling Strategies using Median

Sample Period	Horizon (Months)					
	h_1	*h_2*	*h_3*	*h_4*	*h_5*	*h_6*
1. Forecast Pooling (median) vs Benchmark						
Whole: 2007Q4 to 2020Q3	**0.92**	**0.95****	**0.96****	**0.97****	**0.99**	**0.99**
Non-crisis: 2010Q3 to 2019Q4	**0.89**	**0.93*****	**0.93*****	**0.94*****	**0.96****	**0.96****
GFC: 2007Q4 to 2010Q2	1.00	**0.92****	**0.93****	**0.95***	1.00	1.00
Covid-19: 2020Q1 to 2020Q3	**0.91**	**0.97**	**0.97**	**0.98**	**0.99**	**0.99**
2. Information Pooling vs Forecast Pooling (median)						
Whole: 2007Q4 to 2020Q3	**0.77***	**0.77**	**0.76**	**0.75**	**0.98**	**0.99**
Non-crisis: 2010Q3 to 2020Q2	**0.90**	**0.95**	**0.97**	1.04	1.03	1.10
GFC: 2007Q4 to 2010Q2	**0.71****	**0.90**	**0.76***	**0.77****	**0.88***	**0.86***
Covid-19: 2020Q1 to 2020Q3	**0.75***	**0.68**	**0.71**	**0.67**	**0.99**	1.01
3. Subgroup Forecast Pooling (median) vs Benchmark						
Whole: 2007Q4 to 2020Q3	**0.84**	**0.88***	**0.92****	**0.95****	**0.99***	**0.99**
Non-crisis: 2010Q3 to 2019Q4	**0.86**	**0.91*****	**0.95****	**0.94*****	**0.95*****	**0.97**
GFC: 2007Q4 to 2010Q2	**0.88***	**0.79****	**0.83****	**0.89****	**0.99**	1.01
Covid-19: 2020Q1 to 2020Q3	**0.83**	**0.91**	**0.94**	**0.98**	**0.99**	**0.99**
4. Information Pooling vs Subgroup Forecast Pooling (median)						
Whole: 2007Q4 to 2020Q3	**0.85***	**0.83**	**0.79**	**0.76**	**0.98**	**0.99**
Non-crisis: 2010Q3 to 2020Q2	**0.93**	**0.97**	**0.96**	1.03	1.04	1.08
GFC: 2007Q4 to 2010Q2	**0.80****	1.05	**0.85**	**0.82****	**0.89**	**0.86**
Covid-19: 2020Q1 to 2020Q3	**0.83***	**0.72**	**0.73**	**0.67**	**0.99**	1.01

Note: `***', `**' and `*' indicate that the differences are significant at the 1%, 5% and 10% levels respectively according to the Coroneo-Iacone one-sided test. Figures in bold in panels 1 and 3 indicate the pooling strategy outperform the AR benchmark, whilst those in panels 2 and 4 indicate the forecasts from the factor MIDAS models outperform the pooled forecasts from single indicator MIDAS models. All forecasts, for both the information and forecast pooling approach, are averaged across four MIDAS variants.