Empirical Economics (2023) 65:805–829
https://doi.org/10.1007/s00181-022-02356-9

Forecasting GDP with many predictors in a small open economy: forecast or information pooling?

Hwee Kwan Chow[1] · **Yijie Fei**[2] · **Daniel Han**[3]

Received: 28 February 2022 / Accepted: 28 December 2022 / Published online: 9 January 2023

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 consider various popular weighting schemes in the literature when conducting forecast pooling. As for factor extraction, both the conventional dynamic factor model and the three-pass regression filter approach are considered. We investigate the relative predictive performance of all methods 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. In comparison, we find information pooling tends to dominate both the quarterly autoregressive benchmark model and the forecast pooling strategy particularly during the Global Financial Crisis.

Keywords Forecast evaluation · Factor MIDAS · Three-pass regression filter · Pooling GDP forecasts · Global Financial Crisis · COVID-19 pandemic crisis

JEL Classification C22 · C53 · C55

✉ Yijie Fei
yijiefei@hnu.edu.cn

Hwee Kwan Chow
hkchow@smu.edu.sg

Daniel Han
daniel.han@busaracenter.org

1 School of Economics, Singapore Management University, Singapore, Singapore

2 College of Finance and Statistics, Hunan University, Changsha, China

3 Busara Center for Behavioral Economics, Nairobi, Kenya

1 Introduction

When projecting future economic growth, it is common for forecasters and decision makers to draw on information from a wide variety of economic and financial indicators sampled at different frequencies. 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. This is particularly useful in a fast-evolving environment such as in a crisis where the information content of higher frequency variables provides a timelier assessment of current and near-term future economic conditions. 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 inefficient and inconsistent estimation of the parameters (Andreou et al. 2010). Besides, time aggregation would also lead to a loss of valuable 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 et al. (2004) and Ghysels et al. (2007) directly relates mixed frequency variables in a highly parsimonious way. Since the seminal work of Clements and Galvão (2008, 2009), there has been a burgeoning literature on the application of MIDAS techniques to macroeconomic forecasting, such as for large economies like the US and the euro-area (Armesto et al. 2010; Foroni and Marcellino 2014). Applications to small open economies are found in more recent studies such as Rusnák (2016) for Czech; Kim and Swanson (2018) for Korea; Yau and Hueng (2019) for Taiwan; Galli et al. (2019) for Switzerland; den Reijer and Johansson (2019) for Sweden; Laine and Lindblad (2021) for Finland; and Marcellino and Sivec (2021) for Luxembourg.

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-based MIDAS model. Another promising approach to forecast with a huge mixed frequency dataset is to use penalised regressions, such as Lasso or elastic net. Recent contributions along this line include Uematsu and Tanaka (2019) and Babii et al. (2021). Since the computational costs of these techniques are higher, we focus on the forecast pooling and factor models approaches in this study.

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 Timmermann (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 produce more accurate forecasts is fuelled by mixed evidence from empirical studies. Some in the literature, including Heinisch and Scheufele (2018) and Kuck and

Schweikert (2021), provide empirical evidence that the pooled single indicator forecasts strategy outperforms the factor-based information pooling approach in the case of Germany, while Kuzin et al. (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.

For the information pooling strategy, a critical question to consider is how to extract informative factors from a large unbalanced panel of indicators. In conventional dynamic factor models, the factors represent the underlying movement in the predictors. They can be estimated using either static or dynamic principal components as proposed in Stock and Watson (2002) and Forni et al. (2005), with missing data handled by EM algorithm or vertical realignment. Another popular approach, discussed in Doz et al. (2011), is to cast the factor model into a state-space form before applying the Kalman filter and smoother. Marcellino and Schumacher (2010) provide more details on these approaches in a MIDAS context. More recently, Hepenstrick and Marcellino (2019) proposed a mixed frequency three-pass regression filter (MF-3PRF), which is built upon the work of Kelly and Pruitt (2015). This method enables us to obtain targeted factors in an intuitive and less computationally intensive manner, thereby overcoming the drawback of principal components-based methods that ignore the comovement between indicators and the variable of interest when computing factors.

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 fore- casting exercise by recursively generating GDP growth forecasts for the period 2007Q4 to 2020Q3. Within each approach, we generate the forecasts using a variant of the basic MIDAS model known as the autoregressive distributed lag unrestricted MIDAS (ADL-U-MIDAS) model. This model simply regresses current-quarter GDP growth on lagged GDP growth and contemporaneous and lagged values of the predictors instead of using a functional form restriction on the lags of the predictor as in the basic MIDAS framework. 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.

This paper makes two 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 for short-term forecasting of GDP growth under different regimes, namely crisis versus non-crisis periods. Secondly, 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. For instance, Tsui et al. (2018) employs only daily stock prices in a MIDAS forecast model while Abeysinghe (1998) uses only the external trade variable in a nonlinear dynamic regression model. 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. 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 various forecasting models used in this study, while Sect. 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. Section 5 concludes.

2 An overview of forecasting models

This section describes the models used in our study. We first offer brief descriptions of a variant of MIDAS models and then discuss the factor MIDAS approach and mixed frequency 3PRF.

2.1 Autoregressive distributed lag unrestricted MIDAS (ADL-U-MIDAS) model

To simplify the description of the model, we consider the case of a quarterly dependent variable (y_t), which is Singapore GDP growth in this study, along with one monthly indicator ($x_{i,t}$). The discussion can be extended in a straightforward fashion to incorporate more predictors. The ADL-U-MIDAS model has a forecast equation for h quarters ahead as follows:

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

where the disturbance term is assumed to be independently and identically distributed (i.i.d.) with zero-mean and constant variance. In equation (1), $L^{j/3} x_{i,t+w} = x_{i,t+w-j/3}$ is the lagged value of monthly indicator i. While the dependent variable is only available up to T, the last available observation of the regressor is at $T + w$. For example, $w = 1/3$ 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. $w = 1/3, 2/3$ and 1 correspond to the nowcasts made at the beginning of the first, second and third months of the forecast quarter, respectively.

Instead of having a weighting function depending on a low-dimensional parameter as in the standard MIDAS model proposed by Ghysels et al. (2007), we do not impose any functional constraints on the distributed lags. The unrestricted model, proposed by Foroni et al. (2015), increases flexibility in the specification and eschews estimation via nonlinear least squares, thereby incurring a lower computational cost. We utilise the unrestricted model in our empirical exercise in view of the small frequency mismatch from mixing monthly and quarterly data, resulting in the estimation of relatively few parameters. Following Andreou et al. (2013), the model in (1) is augmented by lagged dependent variables $\sum_{q=0}^{Q} \gamma_{1+q}^{(h)} L^q y_t$. In the rest of the paper, we will simply use the abbreviation 'MIDAS' to refer to the ADL-U-MIDAS model. The benchmark model we consider is an autoregressive model that involves only quarterly GDP data, and whose forecast equation is the same as equation (1) but without the indicator terms.

2.2 Forecast pooling strategies

Given N single-indicator MIDAS models, one can pool the forecasts linearly using

$$\hat{y}_{t+h} = \sum_{i=1}^{N} c_{i,t} \hat{y}_{i,t+h}$$

where $c_{i,t}$ denotes the weight given to indicator i in period t and $\hat{y}_{i,t}$ is the forecast of y_{t+h} made using only $x_{i,t}$. Though it is usually found in empirical research that combining forecasts can lead to better performance than using individual models, there is no consensus regarding how to choose the weighting schemes. Therefore, in this study, we consider various methods that are commonly used in literature.

First, we consider a simple average of all forecasts (FP-EW), where the weight is given as $c_{i,t} = 1/N$. Though it does not take into account historical performance of indicators, many studies have found that this simple method performs remarkably well in out-of-sample exercise (Bec and Mogliani 2015; Galli et al. 2019; Kuzin et al. 2013).

Second, we consider giving a weight inversely proportional to the mean-square forecast error (MSFE). The weights are given by:

$$\begin{aligned} c_{i,t} &= \frac{1/MSFE_{i,t}}{\sum_{i=1}^{N} 1/MSFE_{i,t}}, \\ MSFE_{i,t} &= \sum_{s=t-3}^{t} (y_s - \hat{y}_{i,s})^2, \end{aligned} \tag{2}$$

where $\hat{y}_{i,s}$ is model i's forecast for period s. Theoretically, this approach corresponds to the optimal weighting scheme derived in Bates and Granger (1969) when all forecasts are uncorrelated. It is possible to use a more general regression-based method instead of assuming orthogonality. However, for the sake of simplicity, we follow the existing literature and only consider the weights given in equation (2). The MSFE is computed using a rolling window over previous four quarters (FP-ROLL) to better reflect the rapid changes in the economy, following Stock and Watson (2004) and Kuzin et al. (2013).

Our third scheme assigns a weight proportional to the inverse of discounted MSFE (DMSFE), as discussed in Stock and Watson (2004). The weights are given as

$$c_{i,t} = \frac{1/DMSFE_{i,t}}{\sum_{i=1}^{N} 1/DMSFE_{i,t}},$$
$$DMSFE_{i,t} = \sum_{s=T_0+1}^{t} \delta^{t-s}(y_s - \hat{y}_{i,s})^2, \tag{3}$$

where $\hat{y}_{i,s}$ is again model i's forecast for period s and T_0 is the length of in-sample period. The discount factor δ allows one to focus more on the recent performance of an indicator while still accounting for all historical errors. Following Stock and Watson (2004), we set $\delta = 0.95$ in our empirical study. This scheme will be called FP-DISC.

The last scheme we consider sets the weight proportional to the inverse of a model's performance rank (FP-RANK). This approach is often more robust than other pooling schemes, as ranks are less sensitive to the presence of outliers. With this method, the weights are given as

$$c_{i,t} = \frac{1/r_{i,t}}{\sum_{i=1}^{N} 1/r_{i,t}}, \tag{4}$$

where $r_{i,t}$ denotes the model i's rank at period t determined by $MSFE_{i,t}$ defined in equation (2) Specifically, indicators with lower $MSFE_{i,t}$ will rank higher and vice versa. An alternative rank-method is to use median forecasts. However, we find that this strategy is dominated by other methods in most cases. Hence, its performance is not reported here and is available upon request.

2.3 Factor MIDAS models

Factor MIDAS, originally proposed in Marcellino and Schumacher (2010), synthesises dynamic factor models with MIDAS models. The single indicators in the MIDAS models described in the previous subsection 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 et al. (2013) find no systematic difference in forecast accuracy across different factor extraction methods. In this study, we adopt the two-step estimator by Doz et al. (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:

$$X_t = \phi + \Lambda F_t + \xi_t \tag{5}$$
$$\Psi(L) F_t = B\eta_t \tag{6}$$

where X_t is a $n \times 1$ vector of monthly indicators, Λ is a $n \times r$ matrix of factor loadings for the r static factors, and ξ_t is the idiosyncratic component. Equation (5) is the static factor representation as in Stock and Watson (2002). Equation (6) specifies the dynamics of the factors using a vector autoregression (VAR) of order p. η_t is a $q \times 1$ vector of dynamic shocks orthogonal to the factors and their lags $\Psi(L) = \sum_{i=1}^{p} \psi_i L^i$, and B is an $r \times q$ coefficient matrix.

Banerjee et al. (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. Hence, we follow Galli et al. (2019) in setting $p = 1$ for the sake of parsimony and consider $r = 1, 2$ or 3. We will denote factor MIDAS models with one, two and three factors by DFM1, DFM2 and DFM3, respectively.

Unlike static principal components analysis (PCA), the state-space approach explicitly specifies the dynamics of the factors. As the dimension of X_t 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 $\widehat{F_t}$ using a balanced dataset (truncated at the end of sample).
2. The factor loading matrix Λ is estimated by regressing X_t on $\widehat{F_t}$, hence obtaining the estimated covariance matrix of idiosyncratic components $\widehat{\xi_t}$.
3. The VAR of order $p = 1$ is estimated to obtain $\widehat{\Psi}(L)$ and the residual covariance matrix.
4. Kalman smoothing applied over the unbalanced dataset produces the updated estimate $\widehat{F_t}$.

The extracted monthly factors are then plugged in MIDAS forecasting equation

$$y_{t+h} = \beta_0^{(h)} + \sum_{j=0}^{J} \beta_{1+j}^{(h)} L^{j/3} \hat{F}_{t+w} + \sum_{q=0}^{Q} \gamma_{1+q}^{(h)} L^q y_t + \epsilon_{t+h}, \quad t = 1, 2, \ldots, T \tag{7}$$

to forecast quarterly GDP growth.

2.4 Mixed frequency 3PRF

Factor MIDAS models introduced in the last section are built upon conventional principal component analysis and they do not take into account the targeted variable when extracting factors. There are various approaches available in the literature to compute targeted factors. For instance, Bai and Ng (2008) propose hard and soft thresholds, while Fuentes et al. (2015) consider partial least squares.

In this study, we consider mixed frequency three-pass regression filter recently proposed in Hepenstrick and Marcellino (2019), who extend the 3PRF method in Kelly and Pruitt (2015) to a mixed frequency context. MF-3PRF enables us to obtain targeted factors for forecasting a specific variable of interest, rather than only summarising the information in a large number of predictors. Moreover, it has several asymptotic optimality properties and is quite straightforward to implement, being primarily based on ordinary least squares. Based on a dataset drawn from many different countries, Hepenstrick and Marcellino (2019) and Marcellino and Sivec (2021) provide evidence that MF-3PRF is a promising tool for nowcasting GDP growth.

Kelly and Pruitt (2015) consider the following model:

$$\begin{aligned} y_{t+h} &= \beta_0 + \beta' F_t + \epsilon_{t+h} \\ z_t &= \lambda_0 + \Lambda F_t + \omega_t \\ X_t &= \phi_0 + \Phi F_t + \eta_t \end{aligned} \tag{8}$$

where $F_t = (f_t', g_t')'$ are the $K = K_f + K_g$ common factors driving all indicators; $\beta = (\beta_f', 0')'$, such that y_{t+h} only depends on f_t and not on g_t; z_t is a small set of l proxies of target variable y_t such that $\Lambda = (\Lambda_f, 0)$ with Λ_f non-singular. Compared with the conventional factor-based forecasting model, here the large dataset X_t is possibly driven by more factors than the target variable y_t. It is well-known that, in finite samples, estimating and using only relevant factors f_t would improve forecast performance. Kelly and Pruitt (2015) propose a general and simple method in model (8), based on three steps of OLS regressions. They show that their factor estimator is consistent and the 3PRF forecast converges to the unfeasible forecast, generated if factors were observable, in large samples.

To nowcast quarterly GDP growth using the 3PRF methodology, Hepenstrick and Marcellino (2019) consider the case in which the target variable y_t (or the proxies z_t) are of lower frequency than predictors X_t. Their MF-3PRF method can be summarised in three steps as follows.

- Pass 1: Run a (time-series) regression, in quarterly frequency, of each element of X_t, $\tilde{x}_{i,t}$ on the proxy variables z_t.

$$\tilde{x}_{i,t} = \alpha_{0,i} + z_t'\alpha_i + u_{i,t}, \quad t = 1, \ldots, T, \tag{9}$$

  for each $i = 1, \ldots, N$ and store the OLS estimate $\hat{\alpha}_i$. Here, $\tilde{x}_{i,t} = x_{i,3t} + x_{i,3t-1} + x_{i,3t-2}$ for $t = 1, \ldots, T$ is the quarterly aggregation of monthly indicators. Monthly indicators is cumulated in such a way as this transformation replicates the temporal aggregation applied to (unobservable) month-on-month GDP growth. As pointed out in Marcellino and Sivec (2021), this step is a key element to ensure the consistency of MF-3PRF estimation. Since Hepenstrick and Marcellino (2019) find that adding more factors tends to decrease forecasting performance of MF-3PRF method, we only consider the case of a single factor f_t. As for the number of proxies, results not reported here show that allowing for more proxies in our case will not improve the forecast performance. Therefore, we following Kelly and Pruitt (2015) and simply use the target variable y_t as the proxy z_t.
- Pass 2: Run a (cross-sectional) regression of $x_{i,t}$ on $\hat{\alpha}_i$:

$$x_{i,t} = \alpha_{0,t} + \hat{\alpha}_i' F_t + \epsilon_{i,t}, \quad i = 1, \ldots, N, \tag{10}$$

for each month $t = 1, \ldots, 3T$ and retain the OLS estimate $\hat{F}_t$.

- Pass 3: Split the estimated monthly factors $\hat{F}_t$ obtained in Pass 2 into three quarterly factors and forecast the target variable using a MIDAS regression similar to (7).

Hepenstrick and Marcellino (2019) argue that MF-3PRF inherits the consistency properties of 3PRF and they also discuss how to deal with other data irregularities such as ragged edges. In our empirical study, we consider two methods in this regard. The first method simply fits an autoregressive AR(2) model to each individual time series with a ragged edge, and fills in the missing observations at the end of the series using predicted values. This approach is quite intuitive and computationally straightforward. We call the corresponding model as 3PRF-AR2. The second version we consider relies on Kalman filter algorithm to handle the ragged edges, which leads to the best linear estimates conditional on a correct specification. Besides having to meet the assumption of correct specification, this method has the disadvantage of computational complexity. We denote this model as 3PRF-KF.

3 Data and empirical procedure

3.1 Data

Our large-scale dataset comprises 95 monthly indicators, mostly collected from the CEIC and FRED databases. CEIC and FRED collect 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. 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.

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 Schumacher and Breitung (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. The breakpoint unit root tests are conducted for the period January 1990 to December 2019, i.e. just before the onset of the COVID-19 pandemic crisis. The results of the unit root tests are available from the authors upon request.

Appropriate transformations are taken to induce stationarity in the series, and these mostly involve taking the first difference or log-difference to obtain month-on-month growth rates. A caveat to our main findings would be that these data alterations require full sample information. We do not remove outliers from the data because such a procedure would also rely on full sample information and may bias the results towards finding better forecast performance. The data listing in Appendix A provides details on the data sources, applied transformations and publication lags of the individual indicators.

3.2 Short-term forecasting procedure

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), with a maximum of 12 lags (1 year) for the monthly indicators. The first set of forecasts are thus generated for 2007Q4. We choose the BIC instead of other information criteria like the Akaike information criterion (AIC) or Hannan–Quinn information criterion (HQ) because it tends to select more parsimonious models, which is typically favourable for forecasting.

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. Another advantage of direct forecasting is that we do not have to specify the change in unobserved factors over time for factor-based models (see Marcellino et al. 2006).

We denote $h = i/3$ as h_i for $i = 1, 2, \ldots, 6$ here and for the rest of the paper. We consider monthly forecast horizons from h_1 to h_6 and generate a sequence of six forecasts for each evaluation quarter. For instance, for the evaluation quarter 2019Q1, forecasts are computed on the 2nd of January, February and March 2019, corresponding to h_3, h_2 and h_1, respectively, and these current quarter forecasts are also known as 'nowcasts.' Horizons h_6, h_5 and h_4 refer to the one-quarter-ahead forecasts produced on 2nd of October, November and December 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 as at the end of the sample. In doing so, we implicitly assume stability in the publication lag structure.

For all single-indicator MIDAS models, 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. 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 2 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. 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 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.

As a measure of forecast accuracy, we compute the root mean square forecast error (RMSE) for the entire evaluation period and three subperiods. 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. 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 autoregressive model, 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 single-indicator 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 (henceforth 'DM-CI test') modifies the Diebold and Mariano (2002) test statistics to overcome the prevalence of negative variance estimates that typically arise in smaller samples and with forecasts at longer horizons. 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 $j = 0, 1, \ldots, m$ by

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

where $i = \sqrt{-1}$. Then, the DM-CI test statistics is given by

$$DM_{CI} = \sqrt{T}\left(\frac{\bar{d}}{\hat{V}(\bar{d})}\right) \xrightarrow{d} t_{2m}$$

where $\hat{V}(\bar{d}) = \frac{2\pi}{m}\sum_{j=1}^{m} I(\lambda_j)$. Considering the size-power trade-off when determining the bandwidth, we follow Harvey et al. (2017) in setting $m = T^{\frac{1}{3}}$. We implement the test by adapting codes courtesy of Coroneo and Iacone.

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's GDP growth. To provide an overall description of Singapore's economic performance during our sample period, we first plot the time series of its quarterly GDP growth in Fig. 1. Three significant recessions, including the Asian Financial Crisis, GFC and COVID-19 pandemic, are highlighted. Also plotted are the normalised estimated factors extracted using the first principal component and the three-pass regression filter. Both methods generate factors which trend similarly and closely track GDP growth fluctuations.

Figures 2, 3, 4 and 5 present the relative RMSEs of various pooling strategies over the autoregressive benchmark for the whole sample period and three subperiods, respectively. The benchmark model is an autoregression of the quarterly GDP series whose optimal lag length is determined by the BIC with a maximum lag length of four. Like all the other models in this paper, it is estimated with an expanding window. The length of each bar indicates the relative RMSE of a model, with its value displayed at the end of each bar and the threshold of an RMSE equal to one indicated by a red vertical line. In this way, a number less than one, or equivalently a bar to the left of the red line, indicates that the forecast pooled from the single indicator models or based on pooled information is more accurate than the benchmark forecast. '***,' '**' and '*' indicate the significance at the 1%, 5% and 10% levels, respectively, according to the DM-CI two-sided test and the numbers in parentheses are the *p*-values. The best performing model for a given horizon is highlighted in yellow, while bars with black edges indicate statistical significance at the 10% level.

4.1 Comparison of forecast pooling schemes

We see from Fig. 2 that, for the whole sample, the forecast pooling approaches generally outperform the benchmark at the shorter horizons from h_1 to h_3 but under-perform relative to the benchmark at horizons h_4 to h_6. However, none of the differences in RMSE between the forecast pooling models and the benchmark are statistically significant at the conventional levels. There are interesting differences across the different subperiods. The poor performance at the longer horizons reflects the low forecast accuracy of the forecast pooling methods at these horizons during the pandemic crisis, as shown in Fig. 5. In contrast, we observe from Fig. 3 that the forecast pooling

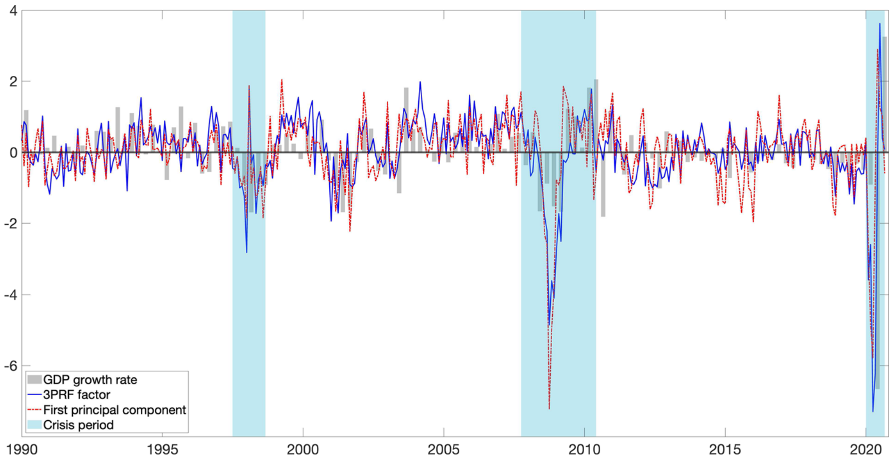

Fig. 1 GDP growth rate and estimated factors: the figure presents the monthly static 3PRF factor (blue solid line) and the first principal component (red dotted line), both estimated using the full sample. Also plotted are observed quarterly GDP growth rates (grey bars). To extract each factor, only cross-sectional information is used. Shaded areas are three recession periods, namely Asian Financial Crisis, Great Financial Crisis and COVID-19 pandemic. Both factors are normalised to have zero mean and unit variance. (Color figure online)

methods systematically outperform the benchmark across all six horizons during the GFC period. There are more considerable gains of up to 15% at horizons h_1 and h_2, and some of the improvements are statistically significant at the 10% level at horizons h_1 and h_4. As for the non-crisis period, the gains in forecast accuracy are all less than 10% and statistically insignificant at conventional levels (see Fig. 4).

Table A.1 (available in Online Appendix) records the RMSEs of alternative weighting schemes relative to the simple average method to facilitate comparisons across the different forecast combination methods. The relative RMSE figures smaller than one are recorded in bold. Each figure is marked by an asterisk(s) whenever the difference in relative RMSE is statistically significant as determined by the results of the two-sided DM-CI tests of equal predictive accuracy. The figures in red indicate the best strategy for a particular horizon (when two models have the same relative RMSE, we choose the one with a smaller p-value). We can conclude from Table A.1 that all three methods are, in general, no worse than the simple average approach, and they could significantly outperform it at the short horizons. Specifically, the FP-ROLL and FP-DISC weighting schemes show statistically significant gains of about 3% to 5% over the simple average method in the GFC period at horizons h_1 and h_2. By incorporating the past performance of different predictors, these pooling strategies can generate more accurate nowcasts.

The combination strategies involve time-varying weights based on the past performance of all single indicator models. This allows us to investigate which variables are more informative for predicting Singapore GDP growth. To this end, Fig. 6 plots the horizon-specific heat maps of forecast pooling weights, where darker blue indicates larger weights. For better visualisation, the average pooling weights are computed for

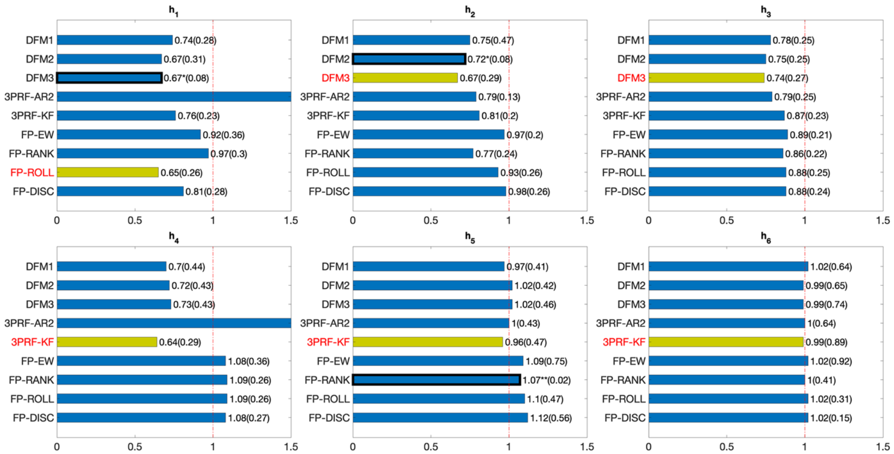

Fig. 2 Relative RMSE over AR (whole period): this figure presents the relative RMSE of various pooling strategies over the AR benchmark for the evaluation period 2007Q4 to 2020Q3. The height of each bar indicates the relative RMSE and its value is displayed at the tips of each bar. '***,' '**' and '*' indicate statistical significance at the 1%, 5% and 10% levels, respectively, according to the Coroneo–Iacone two-sided test. The numbers in parentheses are the *p*-values. The red vertical line indicates RMSEs equal to one. h_1 to h_6 denote the forecast horizons. Yellow bars indicate the best performer for a horizon. Bars with black edges indicate that the model is significant at least at the 10% level. (Color figure online)

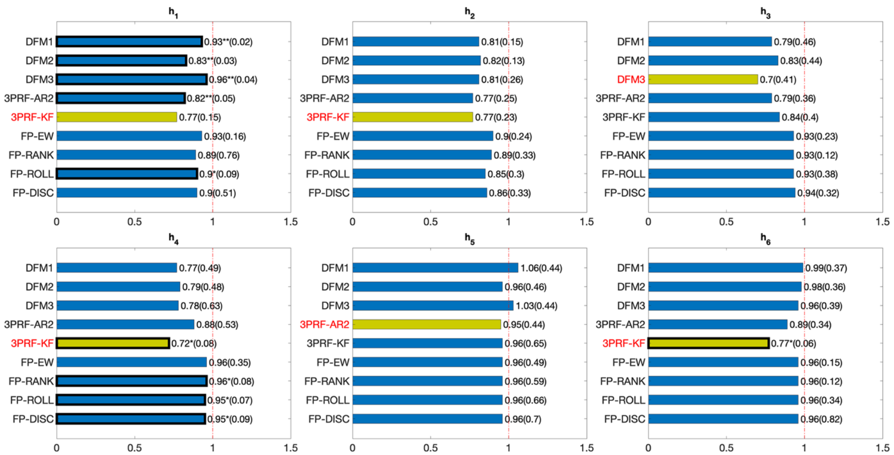

Fig. 3 Relative RMSE over AR (GFC period): this figure presents the relative RMSE of various pooling strategies over the AR benchmark for the evaluation period 2007Q4 to 2010Q2. The height of each bar indicates the relative RMSE and its value is displayed at the tips of each bar. '***,' '**' and '*' indicate statistical significance at the 1%, 5% and 10% levels, respectively, according to the Coroneo–Iacone two-sided test. The numbers in parentheses are the *p*-values. The red vertical line indicates RMSEs equal to one. h_1 to h_6 denote the forecast horizons. Yellow bars indicate the best performer for a horizon. Bars with black edges indicate that the model is significant at least at the 10% level. (Color figure online)

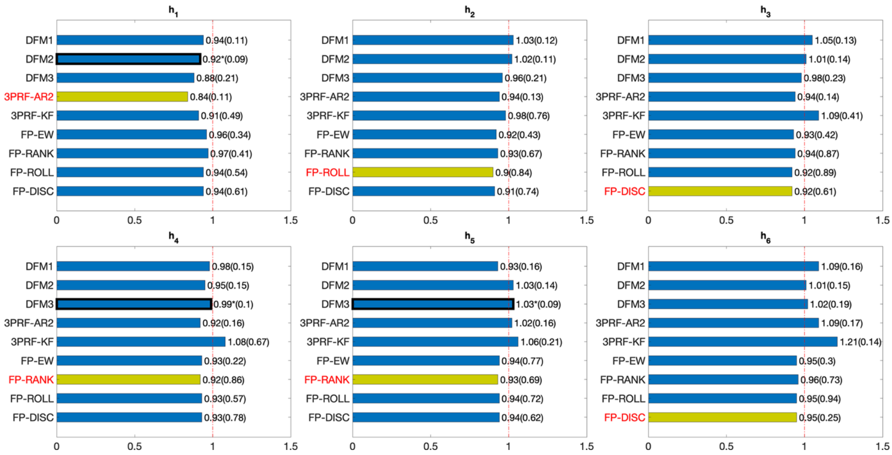

Fig. 4 Relative RMSE over AR (non-crisis period): this figure presents the relative RMSE of various pooling strategies over the AR benchmark for the evaluation period 2010Q3 to 2019Q4. The height of each bar indicates the relative RMSE and its value is displayed at the tips of each bar. '***,' '**' and '*' indicate statistical significance at the 1%, 5% and 10% levels, respectively, according to the Coroneo–Iacone two-sided test. The numbers in parentheses are the p-values. The red vertical line indicates RMSEs equal to one. h_1 to h_6 denote the forecast horizons. Yellow bars indicate the best performer for a horizon. Bars with black edges indicate that the model is significant at least at the 10% level. (Color figure online)

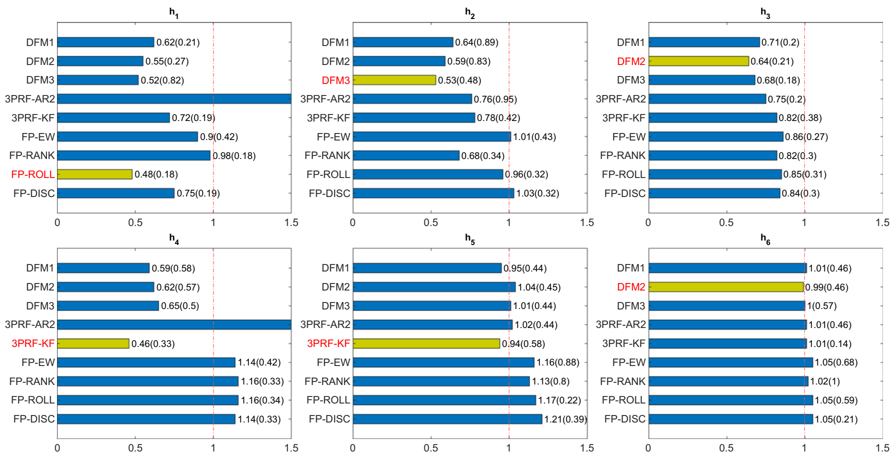

Fig. 5 Relative RMSE over AR (COVID-19 period): this figure presents the relative RMSE of various pooling strategies over the AR benchmark for the evaluation period 2020Q1 to 2020Q3. The height of bars indicates the relative RMSE and its value is displayed at the tips of each bar. '***,' '**' and '*' indicate statistical significance at the 1%, 5% and 10% levels, respectively, according to the Coroneo–Iacone two-sided test. The numbers in parentheses are the p-values. The red vertical line indicates RMSEs equal to one. h_1 to h_6 denote the forecast horizons. Bars with yellow face colour indicate the best performer for a horizon. Bars with black edges indicate that the model is significant at least at the 10% level. (Color figure online)

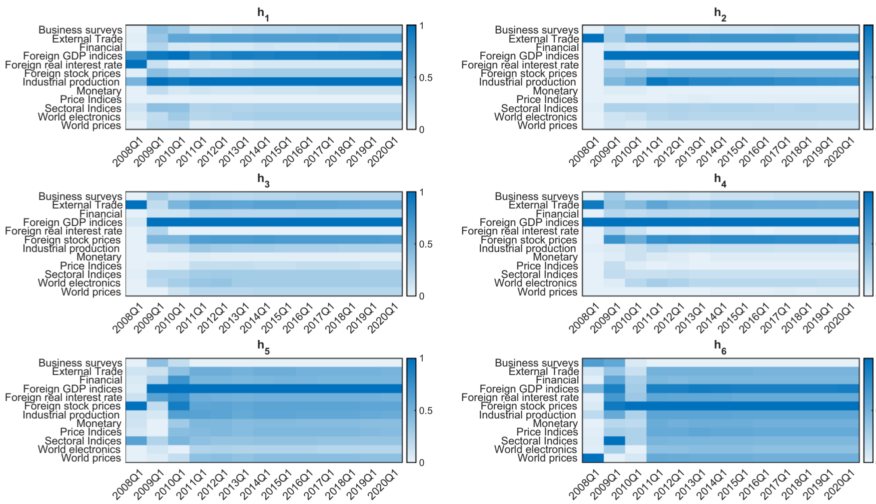

Fig. 6 Heatmaps of average forecast pooling weights: this figure displays horizon-specific heatmaps for the forecasting weights over the evaluation period spanning from 2008Q1 to 2020Q3. For each category, we compute the average forecasting weights, where the weight for an indicator is inversely proportional to its discounted MSFE (discount rate = 0.95). Publication lags of all indicators are taken into account when weights are computed. Higher average weights are indicated by darker areas. h_1 to h_6 denote the forecast horizons

the different categories of indicators described in Sect. 3.1. We present the results of the FP-DISC strategy since it leads to more stable weights and hence more recognisable patterns. It is clear from Fig. 6 that, for all evaluation periods and all forecast horizons, Foreign GDP indices which include the GDP and composite leading indexes of its major trading partners are crucial for forecasting Singapore GDP growth. This is consistent with the small-open-economy nature of Singapore. Other significant short-horizon predictors are those in the Industrial Production and External Trade categories, while Foreign Stock Prices indices become increasingly important as the horizon lengthens. These findings suggest that real macro variables are more relevant for nowcasting current-quarter GDP growth while the financial indicators, perhaps due to their forward-looking nature, provide additional information for short-term forecasting. This is in line with Tsui et al. (2018), who also find financial indicators are better predictors of GDP when they are lagged by about a month or more.

4.2 Comparison of information pooling approaches

Turning to the information pooling approach, we see from Fig. 2 that, with a few exceptions, both the factor MIDAS and MF-3PRF strategies clearly outperform the benchmark model for the whole sample particularly at the first four horizons. The vast majority of the gains exceed 20%, and two of the improvements are statistically significant at horizons h_1 and h_2. The exceptions are the 3PRF-AR2 method at horizons h_1 and h_4, which are due to forecast failure at the onset of the pandemic (see Fig. 5). Nonetheless, during the COVID-19 crisis, we observe substantive gains ranging from 29% to 48% in forecast accuracy for the factor-MIDAS models and a dramatic improvement of 54% at the four-month horizon for the 3PRF-KF model. However, none of the DM-CI statistics are statistically significant due to the small number of forecasts in the COVID-19 crisis subperiod that would lead to a large variance of the squared forecast error difference.

Referring to Fig. 3, the information pooling strategies outperform the benchmark model consistently during the GFC subperiod. Factor MIDAS produces statistically significant gains of 4% to 17% at h_1, while MF-3PRF yields statistically significant improvements of 18%, 28% and 23% at h_1, h_4 and h_6. In particular, 3PRF-KF is the best performing model at most horizons. Notably, the gains in forecast accuracy are greater during the GFC period compared with the non-crisis period for all horizons. During the non-crisis period, improvements in forecast accuracy up are obtained mostly at horizon h_1, with only one of these gains exhibiting statistical significance (see Fig. 4).

Factor MIDAS models and MF-3PRF are extensions of single frequency DFM and 3PRF to the mixed frequency case. Instead of time-aggregating the monthly data to quarterly frequency, these techniques exploit the newly-published high frequency information within the quarter into 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 (2010) and Hepenstrick and Marcellino (2019), who also find factor MIDAS and MF-3PRF approaches are better than single frequency time series models for short-term forecasting of GDP.

4.3 Comparison between information and forecast pooling

We next perform a direct comparison of forecast accuracy between the information pooling and forecast pooling strategies, focusing on the better performing models. Table 1 records the RMSEs of the forecasts from two information pooling approaches, namely DFM2 and 3PRF-KF, as a ratio to the RMSEs of the corresponding forecasts from the forecast pooling approach based on inverse rolling MSFE weights. These information pooling approaches are selected to represent the dynamic factor MIDAS model and mixed frequency three pass regression filter methods, respectively. The format is similar to Table A.1.

We see from Table 1 that information pooling tends to outperform forecast pooling in the GFC subperiod. Nonetheless, we obtain statistical significance in the differences in predictive accuracy only at the four-month horizon. These gains for the DFM2 and 3PRF-KF models are large at 17% and 25%, respectively. It appears that the 3PRF-KF method outperforms the DFM2 model during the GFC subperiod at all horizons except h_3. As for the COVID-19 crisis subperiod, the forecasts from both information pooling approaches are more accurate than those from pooled single indicator models at all but the shortest horizon. The gains from factor MIDAS forecasts are considerable ranging from 24% to 46% at horizons up to four months ahead, while the 3PRF-KF forecasts offer an impressive 60% gain at h_4 during the pandemic. Again, we do not observe statistical significance in the differentials due to the small number of observations in this subperiod. In contrast, forecast pooling tends to yield lower RMSEs compared to the information pooling methods at all but the shortest horizon in the non-crisis

Table 1 Relative RMSE of DFM2 and 3PRF-KF over FP-ROLL

	h_1	h_2	h_3	h_4	h_5	h_6
Horizon						
2007Q4–2020Q3						
DFM2	*1.03 (0.71)*	***0.72 (0.30)***	***0.84 (0.20)***	**0.67 (0.31)**	**0.92 (0.49)**	***0.98 (0.51)***
3PRF-KF	1.16 (0.54)	**0.87 (0.29)**	**0.99 (0.80)**	***0.59 (0.32)***	***0.87 (0.43)***	**0.98 (0.66)**
2007Q4–2010Q2						
DFM2	**0.92 (0.34)**	**0.97 (0.74)**	***0.89 (0.34)***	**0.83*** (0.01)**	1.00 (0.97)	1.02 (0.72)
3PRF-KF	***0.86 (0.31)***	***0.91 (0.44)***	**0.90 (0.40)**	***0.75** (0.04)***	***0.99 (0.95)***	***0.80 (0.15)***
2010Q3–2019Q4						
DFM2	**0.97 (0.82)**	1.12 (0.21)	*1.10 (0.31)*	*1.02 (0.80)*	*1.10** (0.05)*	*1.06 (0.52)*
3PRF-KF	***0.96 (0.75)***	*1.08 (0.34)*	1.18 (0.11)	1.16 (0.39)	1.14* (0.06)	1.27 (0.12)
2020Q1–2020Q3						
DFM2	*1.13 (0.89)*	***0.61 (0.20)***	***0.76 (0.55)***	**0.54 (0.20)**	**0.88 (0.48)**	***0.94 (0.29)***
3PRF-KF	1.48 (0.61)	**0.81 (0.19)**	**0.97 (0.90)**	***0.40 (0.19)***	***0.80 (0.26)***	**0.96 (0.53)**

'***,' '**' and '*' indicate that the differences are statistically significant at the 1%, 5% and 10% levels, respectively, according to the Coroneo–Iacone two-sided test. Figures in the parenthesis are corresponding p-values of the test. Figures in bold indicate information pooling strategies that outperform the forecast pooling benchmark (relative RMSE less than one). Figures in italic, bolditalic indicate the better strategy for a particular horizon (when two models have same relative RMSE, we choose the one with a smaller p-value). h_1 to h_6 denote the forecast horizons period. However, the differentials are statistically significant only at horizons h_5, at 10% and 14% for DFM2 and 3PRF-KF, respectively, for this stable growth period.

To aid the discussion of results, forecast plots for each subperiod are generated from FP-ROLL, DFM2 and 3PRF-KF approaches for the six individual horizons are displayed in Figs. A.1 to A.3 (available in Online Appendix), respectively. We observe that real output growth, as plotted in Fig. A.1, plunged during the GFC, 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 (see Fig. A.2). It appears that all pooling strategies have similar performance in this stable non-crisis subperiod, apart from the greater volatility in the forecasts from information pooling models compared to those from the pooled single indicator models. 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 in forecast accuracy between the individual pooling approaches versus the benchmark model in the non-crisis subperiod, as seen in Fig. 4, is not surprising.

It follows that the choice between the forecasting 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 Figs. A.1 to A.3 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 (see the blue lines in Figs. A.1 and A.3, respectively). Similarly, Galli et al. (2019) find that forecast pooling produces less volatile forecasts that are not flexible enough to capture sudden swings in output growth in Switzerland. By contrast, the charts show that information pooling forecasts can better track rapid changes in output growth during crises (see the red lines in Figs. A.1 and A.3). This observation is consistent with the higher forecast precision reported in Table 1 for the information pooling approach vis-a-vis the forecast pooling approach during crises. In other words, the visual comparisons are in agreement with the numerical results.

5 Conclusion

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. 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, using higher frequencies such as monthly indicators provides more timely information on economic fluctuations, particularly in fast-evolving conditions such as in a crisis. Both pooling approaches under study aim to extract the predictive content from our large-scale ragged-edge dataset that spans the GFC, the COVID-19 pandemic crisis and the non-crisis period.

We find that for forecast pooling, various weighting schemes lead to small differences in predictive accuracy. Consistent with the small-open-economy nature of

Singapore, we show that indicators related to external sectors are more critical for improving out-of-sample performance. Moreover, forecast pooling tends to generate more accurate forecasts than the benchmark model during the GFC but not in the non-crisis period. Since Singapore's output growth was fairly stable in normal times, we expect any reasonable predictive model to record a decent forecast performance. Indeed, we found that neither pooling strategy records substantive improvements to the forecast performance of the quarterly autoregressive benchmark model in the non-crisis subperiod.

Conversely, larger differentials in the predictive ability of the two pooling approaches were recorded during crisis subperiods. Our results indicate information pooling, using either dynamic factor models or mixed frequency three pass regression filters, tends to dominate both the benchmark model and forecast pooling during the GFC crisis. Most of the improvements in forecast accuracy are evident at the shorter horizons and are statistically significant at various horizons. Overall, the findings suggest the information pooling strategy has superior short-term predictive ability during crisis periods as it is better suited to capture the myriad of shocks hitting the small open economy in periods of wide economic fluctuations.

Supplementary Information The online version contains supplementary material available at https://doi.org/10.1007/s00181-022-02356-9.

Acknowledgements The authors would like to thank two anonymous referees and participants of the 2021 China Meeting of the Econometric Society for helpful comments and suggestions. The authors also thank Tze Yong Yew for excellent research assistance. All remaining errors are our own. Yijie Fei gratefully acknowledges the financial support from the Fundamental Research Funds for the Central Universities in China.

Declarations

Conflict of interest All authors declare that they have no conflict of interest.

Ethical approval This article does not contain any studies with human participants or animals performed by any of the authors.

Appendix A: Data listing

Data	Source	Transformation	Publication lag
0. Real GDP (quarterly)	CEIC	Dln	2 months
Foreign GDP indices/composite leading indicators (6)			
1. Major 5 Asia GDP index	FRED	Dln	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 CLI	OECD	Deviations from trend	1 month
Foreign stock prices (11)			
7. NASDAQ Composite Index (US)	CEIC	Dln	None
8. Nikkei 225 (Japan)	CEIC	Dln	None
9. DAX (Germany)	CEIC	Dln	None
10. FTSE100 (UK)	CEIC	Dln	None
11. KOSPI (Korea)	CEIC	Dln	None
12. FTSE Bursa (Malaysia)	CEIC	Dln	None
13. JKSE (Indonesia)	CEIC	Dln	None
14. SET Index (Thailand)	CEIC	Dln	None
15. PSEI (Philippines)	CEIC	Dln	None
16. Shanghai Composite Index (China)	CEIC	Dln	None
17. TSWE (Taiwan)	CEIC	Dln	None
Foreign real interest rates (3)			
18. US (3-month treasury bill rate-CPI inflation)	CEIC	D	1 month
19. Japan (3-month TIBOR-CPI inflation)	CEIC	D	1 month
20. UK (3-month LIBOR-CPI inflation)	CEIC	D	1 month
World electronics (5)			
21. Global semiconductor sales*	CEIC	Dln	2 months
22. US new orders for electronics (excl. semiconductors)	US Census Bureau	Dln	2 months
23. US electronics shipments-to-inventories ratio for electronics	US Census Bureau	Dln	2 months
24. PPI for US electronics	FRED	None	1 month
World prices (4)			
25. OPEC Crude oil price	CEIC	Dln	1 month
26. Global food price index*	FRED	Dln	1 month
27. Global non-fuel price index	FRED	Dln	1 month
28. Global commodity price index	FRED	Dln	1 month
Industrial production (7)			
29. Industrial production index (IPI)	Department of Statistics, Singapore (DOS)	Dln	1 month
30. IPI: Biomedicals	DOS	Dln	1 month
31. IPI: Transport engineering	DOS	Dln	1 month
32. IPI: Precision engineering	DOS	Dln	1 month
33. IPI: General manufacturing	DOS	Dln	1 month
34. IPI: Electronics	DOS	Dln	1 month
35. IPI: Chemicals	DOS	Dln	1 month

Data	Source	Transformation	Publication lag
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	Dln	2 months
40. Retail sales value*	CEIC	Dln	2 months
41. Car registrations Above 1600cc*	CEIC	Dln	1 month
42. Car registrations Below 1600cc*	CEIC	Dln	1 month
43. Visitor arrivals	CEIC	Dln	1 month
44. Air cargo loaded*	CEIC	Dln	1 month
45. Air cargo discharged*	CEIC	Dln	1 month
46. Sea cargo handled*	CEIC	Dln	1 month
47. Electricity generation*	CEIC	Dln	2 months
48. Formation of companies*	CEIC	Dln	1 month
49. Construction contracts awarded (BCA)	CEIC	None	2 months
External Trade (12)			
50. Total imports	CEIC	Dln	1 month
51. Imports: non-oil	CEIC	Dln	1 month
52. Imports: oil	CEIC	Dln	1 month
53. Total exports	CEIC	Dln	1 month
54. Exports: non-oil	CEIC	Dln	1 month
55. Exports: oil	CEIC	Dln	1 month
56. Domestic exports	CEIC	Dln	1 month
57. Domestic exports: non-oil	CEIC	Dln	1 month
58. Domestic exports: oil	CEIC	Dln	1 month
59. Re-exports	CEIC	Dln	1 month
60. Re-exports: non-oil	CEIC	Dln	1 month
61. Re-exports: oil	CEIC	Dln	1 month
Price Indices (7)			
62. Export price index	CEIC	Dln	1 month
63. Export price index: non-oil	CEIC	Dln	1 month
64. Import price index	CEIC	Dln	1 month
65. Import price index: non-oil	CEIC	Dln	1 month
66. Consumer price index	CEIC	Dln	1 month
67. Domestic supply price index	CEIC	Dln	1 month
68. Manufactured products price index	CEIC	Dln	1 month
Financial (18)			
69. Straits Times Index	CEIC	Dln	None
70. 1-year treasury bill yield	CEIC	D	None
71. 2-year treasury bill yield	CEIC	D	None
72. 5-year treasury bill yield	CEIC	D	None
73. 10-year treasury bill yield	CEIC	D	None
74. 3-month SIBOR	CEIC	D	1 month
75. Yield spread[#] (10-year minus 3-month SIBOR)	CEIC	None	1 month
76. Singapore dollar to Australian dollar	CEIC	Dln	None
77. Singapore dollar to Pound	CEIC	Dln	None
78. Singapore dollar to Renminbi	CEIC	Dln	None
79. Singapore dollar to Euro	CEIC	Dln	None
80. Singapore dollar to Hong Kong dollar	CEIC	Dln	None
81. Singapore dollar to Yen	CEIC	Dln	None
82. Singapore dollar to Ringgit	CEIC	Dln	None

Data	Source	Transformation	Publication lag
83. Singapore dollar to US dollar	CEIC	Dln	None
84. Singapore dollar to Franc	CEIC	Dln	None
85. Nominal effective exchange rate	CEIC	Dln	1 month
86. Real effective exchange rate	CEIC	Dln	1 month
Monetary (9)			
87. M1	CEIC	Dln	1 month
88. M2	CEIC	Dln	1 month
89. M3	CEIC	Dln	1 month
90. Loans advances	CEIC	Dln	1 month
91. Loans advances: manufacturing	CEIC	Dln	1 month
92. Loans advances: construction	CEIC	Dln	1 month
93. Loans advances: commerce*	CEIC	Dln	1 month
94. Loans advances: financial institutions	CEIC	Dln	1 month
95. Loans advances: individuals	CEIC	Dln	1 month

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

References

Abeysinghe T (1998) Forecasting Singapore's quarterly GDP with monthly external trade. Int J Forecast 14(4):505–513

Andreou E, Ghysels E, Kourtellos A (2010) Regression models with mixed sampling frequencies. J Econom 158(2):246–261

Andreou E, Ghysels E, Kourtellos A (2013) Should macroeconomic forecasters use daily financial data and how? J Bus Econ Stat 31(2):240–251

Armesto MT, Engemann KM, Owyang MT et al (2010) Forecasting with mixed frequencies. Federal Reserve Bank of St. Louis Review 92(6):521–536

Babii A, Ghysels E, Striaukas J (2021) Machine learning time series regressions with an application to nowcasting. *J Bus Econ Stat*, pp 1–23

Bai J, Ng S (2008) Forecasting economic time series using targeted predictors. J Econom 146(2):304–317

Banerjee A, Marcellino M, Masten I (2005) Leading indicators for euro-area inflation and GDP growth. Oxford Bull Econ Stat 67:785–813

Bates JM, Granger CW (1969) The combination of forecasts. J Oper Res Soc 20(4):451–468

Bec F, Mogliani M (2015) Nowcasting French GDP in real-time with surveys and "blocked" regressions: combining forecasts or pooling information? Int J Forecast 31(4):1021–1042

Boivin J, Ng S (2005) Understanding and comparing factor-based forecasts. Int J Central Bank 1(3)

Chow HK, Choy KM (2009) Analyzing and forecasting business cycles in a small open economy: a dynamic factor model for Singapore. OECD J: J Bus Cycle Meas Anal 2009(1):19–41

Chow HK, Choy KM (2009) Monetary policy and asset prices in a small open economy: a factor-augmented VAR analysis for Singapore. Ann Financ Econ 5(01):0950004

Clements MP, Galvão AB (2008) Macroeconomic forecasting with mixed-frequency data: forecasting output growth in the United States. J Bus Econ Stat 26(4):546–554

Clements MP, Galvão AB (2009) Forecasting US output growth using leading indicators: an appraisal using MIDAS models. J Appl Econom 24(7):1187–1206

Coroneo L, Iacone F (2020) Comparing predictive accuracy in small samples using fixed-smoothing asymptotics. J Appl Econom 35(4):391–409

den Reijer A, Johansson A (2019) Nowcasting Swedish GDP with a large and unbalanced data set. Empir Econ 57(4):1351–1373

Diebold FX, Mariano RS (2002) Comparing predictive accuracy. J Bus Econ Stat 20(1):134–144

Doz C, Giannone D, Reichlin L (2011) A two-step estimator for large approximate dynamic factor models based on Kalman filtering. J Econom 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? J Monet Econ 50(6):1243–1255
Forni M, Hallin M, Lippi M, Reichlin L (2005) The generalized dynamic factor model: one-sided estimation and forecasting. J Am Stat Assoc 100(471):830–840
Foroni C, Marcellino M (2014) A comparison of mixed frequency approaches for nowcasting euro area macroeconomic aggregates. Int J Forecast 30(3):554–568
Foroni C, Marcellino M, Schumacher C (2015) Unrestricted mixed data sampling (MIDAS): MIDAS regressions with unrestricted lag polynomials. J R Stat Soc A Stat Soc 178(1):57–82
Fuentes J, Poncela P, Rodríguez J (2015) Sparse partial least squares in time series for macroeconomic forecasting. J Appl Econom 30(4):576–595
Galli A, Hepenstrick C, Scheufele R (2019) Mixed-frequency models for tracking short-term economic developments in Switzerland. 58th issue (June 2019) of the International Journal of Central Banking
Ghysels E, Santa-Clara P, Valkanov R (2004). The MIDAS touch: mixed data sampling regression models
Ghysels E, Sinko A, Valkanov R (2007) MIDAS regressions: further results and new directions. Economet Rev 26(1):53–90
Harvey DI, Leybourne SJ, Whitehouse EJ (2017) Forecast evaluation tests and negative long-run variance estimates in small samples. Int J Forecast 33(4):833–847
Heinisch K, Scheufele R (2018) Bottom-up or direct? Forecasting German GDP in a data-rich environment. Empir Econ 54(2):705–745
Hepenstrick C, Marcellino M (2019) Forecasting gross domestic product growth with large unbalanced data sets: the mixed frequency three-pass regression filter. J R Stat Soc A Stat Soc 182(1):69–99
Kelly B, Pruitt S (2015) The three-pass regression filter: a new approach to forecasting using many predictors. J Econom 186(2):294–316
Kim HH, Swanson NR (2018) Methods for backcasting, nowcasting and forecasting using factor-MIDAS: with an application to Korean GDP. J Forecast 37(3):281–302
Kuck K, Schweikert K (2021) Forecasting Baden–Württemberg's GDP growth: MIDAS regressions versus dynamic mixed-frequency factor models. J Forecast 40(5):861–882
Kuzin V, Marcellino M, Schumacher C (2013) Pooling versus model selection for nowcasting GDP with many predictors: empirical evidence for six industrialized countries. J Appl Econom 28(3):392–411
Laine O-M, Lindblad A (2021) Nowcasting Finnish GDP growth using financial variables: a MIDAS approach. J Finnish Econ Assoc 2(1):74–108
Marcellino M, Schumacher C (2010) Factor MIDAS for nowcasting and forecasting with ragged-edge data: a model comparison for German GDP. Oxford Bull Econ Stat 72(4):518–550
Marcellino M, Sivec V (2021) Nowcasting GDP growth in a small open economy. Natl Inst Econ Rev 256:127–161
Marcellino M, Stock JH, Watson MW (2006) A comparison of direct and iterated multistep AR methods for forecasting macroeconomic time series. J Econom 135(1–2):499–526
Rusnák M (2016) Nowcasting Czech GDP in real time. Econ Model 54:26–39
Schumacher C, Breitung J (2008) Real-time forecasting of German GDP based on a large factor model with monthly and quarterly data. Int J Forecast 24(3):386–398
Stock JH, Watson MW (2002) Macroeconomic forecasting using diffusion indexes. J Bus Econ Stat 20(2):147–162
Stock JH, Watson MW (2004) Combination forecasts of output growth in a seven-country data set. J Forecast 23(6):405–430
Tay AS (2007) Financial variables as predictors of real output growth
Timmermann A (2006) Forecast combinations. Handb Econ Forecast 1:135–196
Tsui AK, Xu CY, Zhang Z (2018) Macroeconomic forecasting with mixed data sampling frequencies: evidence from a small open economy. J Forecast 37(6):666–675
Uematsu Y, Tanaka S (2019) High-dimensional macroeconomic forecasting and variable selection via penalized regression. Econom J 22(1):34–56
Yau R, Hueng CJ (2019) Nowcasting GDP growth for small open economies with a mixed-frequency structural model. Comput Econ 54(1):177–198