

Faria, Gonçalo; Verona, Fabio

Working Paper

Forecast combination in the frequency domain

Bank of Finland Research Discussion Papers, No. 1/2023

Provided in Cooperation with:

Bank of Finland, Helsinki

Suggested Citation: Faria, Gonçalo; Verona, Fabio (2023) : Forecast combination in the frequency domain, Bank of Finland Research Discussion Papers, No. 1/2023, ISBN 978-952-323-430-7, Bank of Finland, Helsinki

This Version is available at:

<https://hdl.handle.net/10419/268426>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Bank of Finland Research Discussion Papers
1 • 2023

Gonçalo Faria – Fabio Verona

Forecast combination in the frequency domain



Bank of Finland
Research

Bank of Finland Research Discussion Papers
Editor-in-Chief Esa Jokivuolle

Bank of Finland Research Discussion Papers 1/2023
19 January 2023

Gonçalo Faria – Fabio Verona:
Forecast combination in the frequency domain

ISBN 978-952-323-430-7, online
ISSN 1456-6184, online

Bank of Finland
Research Unit

PO Box 160
FIN-00101 Helsinki

Phone: +358 9 1831

Email: research@bof.fi

Website: www.suomenpankki.fi/en/research/research-unit/

The opinions expressed in this paper are those of the authors and do not necessarily reflect the views of the Bank of Finland.

Forecast combination in the frequency domain*

Gonçalo Faria[†] Fabio Verona[‡]

Abstract

Predictability is time *and* frequency dependent. We propose a new forecasting method – forecast combination in the frequency domain – that takes this fact into account. With this method we forecast the equity premium and real GDP growth rate. Combining forecasts in the frequency domain produces markedly more accurate predictions relative to the standard forecast combination in the time domain, both in terms of statistical and economic measures of out-of-sample predictability. In a real-time forecasting exercise, the flexibility of this method allows to capture remarkably well the sudden and abrupt drops associated with recessions and further improve predictability.

Keywords: forecast combination, frequency domain, equity premium, GDP growth, Haar filter, wavelets

JEL classification: C58, G11, G17

* We are grateful for the comments of participants at the Bank of Finland and Católica Porto Business School seminars. The views expressed here are those of the authors and do not necessarily reflect those of the Bank of Finland. Faria gratefully acknowledges financial support from Fundação para a Ciência e Tecnologia through project UIDB/00731/2020.

[†] Universidade Católica Portuguesa, Católica Porto Business School and CEGE, and University of Vigo – *RGEA* (gfaria@ucp.pt).

[‡] Bank of Finland – Monetary Policy and Research Department, and University of Porto – *CEFUP* (fabio.verona@bof.fi). Corresponding author.

1 Introduction

Multiple forecasts of the same variable are often available. An issue decision makers face is thus how best to exploit the information of each individual forecast. This is challenging as predictability is time dependent: no individual variable is clearly superior to other variables consistently throughout time (see e.g. Stock and Watson, 2004 and Henkel, Martin and Nardari, 2011).

The fact that the best forecasting variable changes over time renders individual variables unreliable predictors. A method proposed to overcome this problem is forecast combination. Since Bates and Granger's (1969) seminal paper, it has been known that combining forecasts across models often produces a forecast that performs better than the best individual model. Forecast combination achieves a compromise between smoothing out the excessive noise in the individual forecasts and the need to retain some of the volatility that allows to capture the time-varying behavior of the variable of interest.¹

More recent empirical literature has shown that predictability is *also* frequency dependent. Some frequencies of a variable might be good predictors for the variable of interest, others might not. For instance, Faria and Verona (2020) show that the low frequency of the term spread has good predictive power (for equity returns), while the remaining frequencies don't. Likewise, some frequencies of the targeted variable need to be forecasted well: Faria and Verona (2021) and Martins and Verona (2021) show that it is crucial to predict well the low frequencies of the equity premium and inflation, respectively, while the other frequencies mainly bring noise to the forecast exercise.

In this paper we propose a method that reduces the forecast noise simultaneously in the time and in the frequency domain. It is a forecast combination method that takes the frequency dependence between variables into account. We apply it to forecast a financial variable – the equity premium – and

¹ Timmermann (2006) provides an extensive survey on forecast combination methods. Recent contributions include Rapach, Strauss and Zhou (2010), Pettenuzzo and Ravazzolo (2016), and Pirschel and Wolters (2018). Other methods that incorporate information from a large set of potential predictors in a predictive regression framework include Bayesian model averaging (Cremers, 2002) and factor models (Stock and Watson, 2002 and Caggiano, Kapetanios and Labhard, 2011).

a macro variable – real GDP growth rate. The method goes as follows. We decompose our target variables and a given set of n ($n=15$) predictor variables into f ($f=5$) time series frequency components, each of them capturing the oscillations of the original variable within a specific frequency band. We then forecast, separately, each of the f frequency component of the target variables using the corresponding frequency component f of one predictor at a time. We obtain n forecasts for each frequency f of the target variables. Subsequently, the forecast of each frequency component f of the target variables is computed as the combination of the n forecasts of that frequency component from the n predictor variables. Finally, the overall forecast of each target variable is computed as the sum of the f forecasts of its f frequency components.

We find that combining forecasts in the frequency domain produces markedly more accurate predictions relative to existing alternatives, both in terms of statistical and economic measures of out-of-sample predictability. The advantage of this approach relies on its flexibility, as it allows to exclude some of the frequencies of the target variables when computing the forecasts. This is important: we show that, in recessions, it is crucial to ignore the forecast of the low-frequency components of the target variables to have good forecasts. In a real-time (robustness) exercise, the possibility of ignoring the low-frequency components in recessions allows to improve the forecasting results even further, as the forecasts capture remarkably well the sudden and abrupt drops associated with recessions.

The remaining of this paper is organized as follows. In section 2 we present the data and the band-pass filter used to extract the frequency components from the original variables. Section 3 outlines the econometric methodology. The out-of-sample forecasting results for the individual predictive regression models and the forecast combination methods are reported in section 4. We analyze the predictability over the business cycle in section 5. Robustness tests are briefly described in section 6, and section 7 concludes.

2 Data

We use U.S. quarterly data from 1947:Q1 until 2019:Q4. The target variables are the equity premium and real GDP growth rate (quarter-over-quarter). The equity premium in quarter t is measured by the difference between the log (total) return of the S&P500 index in quarter t and the log return on a three-month Treasury bill at the beginning of quarter t .

As predictors, we use fifteen variables from Goyal and Welch (2008). Specifically, we use the log dividend-price ratio (DP), the log dividend yield (DY), the log earnings-price ratio (EP), the log dividend-payout ratio (DE), the stock variance (SVAR), the book-to-market ratio (BM), the net equity expansion (NTIS), the Treasury bill rate (TBL), the long-term bond yield (LTY), the long-term bond return (LTR), the term spread (TMS), the default yield spread (DFY), the default return spread (DFR), the lagged inflation rate (INFL), and the lagged investment rate (IK). While these predictors are extensively used to forecast equity returns, several of them have predictive ability with respect to real GDP growth as well (see e.g. Stock and Watson, 2003). A classical example is the slope of the yield curve (proxied by the term spread), which has long been used to forecast recessions (see e.g. Estrella and Hardouvelis, 1991).

In appendix 1 these predictors are briefly explained. The time series of the target variables and of the predictors are plotted in figure 1 and 9, respectively. Table 1 reports summary statistics for all the variables. We note here that both target variables are negatively skewed, suggesting that both the real economy and the equity market have more crashes than what would happen if they were normally distributed.

To decompose the variables into their time series frequency components, we band-pass the data with the Haar filter.² We consider five frequency components: the first one (D_1) captures fluctuations of the original variable with a period between 2 and 4 quarters, while components D_2 , D_3 , and

² Besides its simplicity and wide use (see e.g. Faria and Verona, 2018, Bandi, Perron, Tamoni and Tebaldi, 2019, and Lubik, Matthes and Verona, 2019), the Haar filter makes a clean connection to temporal aggregation as the filter coefficients are simply differences of moving averages.

D_4 capture fluctuations with a period of 1-2, 2-4, and 4-8 years, respectively. Finally, component D_5 captures fluctuations with a period longer than 8 years. We note that the sum of these five time series frequency components gives exactly the original time series of the original variable.

As an example, figure 2 shows the time series of investment rate (upper left plot) and of its time series frequency components (remaining plots). Component D_1 captures the high-frequency movements of investment rate (the noisy component) and frequency component D_5 its trend, while the remaining frequencies ($D_2 - D_4$) broadly capture its business cycle frequency fluctuations.

3 Econometric methodology

The one-step ahead out-of-sample (OOS) forecasts are generated using a sequence of expanding windows. We use an initial in-sample period (1947:Q1 to 1964:Q4) to make the first one-step ahead OOS forecast. The in-sample period is then increased by one observation and a new one-step ahead OOS forecast is produced. This is the procedure until the end of the sample. The full OOS period therefore spans from 1965:Q1 to 2019:Q4 (that is, 55 years of quarterly forecasts).

3.1 Predictive regression model

Let r be the target variable (the equity premium or real GDP growth rate). For each predictor $x_i, i = 1, \dots, n$ ($n = 15$), the predictive regression is

$$r_{x_i,t} = \alpha_{x_i} + \beta_{x_i} x_{i,t-1} + \varepsilon_t, \quad (1)$$

and the corresponding forecasts are given by

$$\hat{r}_{x_i,t+1} = \hat{\alpha}_{x_i} + \hat{\beta}_{x_i} x_{i,t}, \quad (2)$$

where $\hat{\alpha}_{x_i}$ and $\hat{\beta}_{x_i}$ are the OLS estimates of α_{x_i} and β_{x_i} in equation (1), respectively, using data from the beginning of the sample until quarter t .

3.2 Forecast combination in the time domain

The forecast combination of r in the time domain (TD) made at time t for $t+1$, denoted $FC - TD_{t+1}$, is the mean of the n ($n=15$) individual forecasts based on equation (2):³

$$FC - TD_{t+1} = \frac{1}{n} \sum_{i=1}^n \hat{r}_{x_i, t+1} . \quad (3)$$

3.3 Forecast combination in the frequency domain

The first step of our method consists in decomposing all variables into their time series frequency components ($D_2 - D_5$). As in Faria and Verona (2021), we then estimate, for each predictor x_i , a model like (1) for each frequency f . That is, we estimate – separately – each frequency component D_f of r using the frequency component of the predictor x_i at the same frequency f :⁴

$$r_t^{D_f, x_i} = \alpha_{t, f}^{x_i} + \beta_{t, f}^{x_i} x_{i, t-1}^{D_f} + \varepsilon_t . \quad (4)$$

We use these estimation results to produce the one-step ahead OOS forecast of the corresponding frequency component of r :

$$\hat{r}_{t+1}^{D_f, x_i} = \hat{\alpha}_{t, f}^{x_i} + \hat{\beta}_{t, f}^{x_i} x_{i, t}^{D_f} ,$$

where $\hat{\alpha}_{t, f}^{x_i}$ and $\hat{\beta}_{t, f}^{x_i}$ are the OLS estimates of $\alpha_{t, f}^{x_i}$ and $\beta_{t, f}^{x_i}$ in equation (4), respectively, using data from the beginning of the sample until quarter t .⁵

³ We also considered other combination methods (median, trimmed mean, as well as discounted mean square prediction error). Results were usually not better than the mean average.

⁴ This setup is akin to the band spectrum regression proposed by Engle (1974).

⁵ As we use a two-sided filter in the OOS exercise, we recompute the time series frequency components of the variables recursively at each iteration of the OOS forecasting process using data from the start of the sample through

We then use the combination forecast method to compute the forecasts of each frequency components D_f of r , which is computed as the mean forecast combination for that frequency f :

$$\hat{r}_{c,t+1}^{D_f} = \frac{1}{n} \sum_{i=1}^n \hat{r}_{t+1}^{D_f, x_i} .$$

Finally, the overall forecast of r made at time t for $t+1$ in the frequency domain (FD), denoted $FC - FD_{t+1}$, is obtained by summing the forecasts of the f individual frequencies of r :

$$FC - FD_{t+1} = \sum_{f=1}^5 \hat{r}_{c,t+1}^{D_f} = \sum_{f=1}^5 \left(\frac{1}{n} \sum_{i=1}^n \hat{r}_{t+1}^{D_f, x_i} \right) . \quad (5)$$

Table 2 provides a sketch of the forecast combination models (FC-TD and FC-FD) and helps visualizing the differences between them.

3.4 Forecast evaluation

3.4.1 Statistical performance

The forecasting performances of the forecast combination models are evaluated using the Campbell and Thompson (2008) R_{OS}^2 statistic. The R_{OS}^2 statistic measures the proportional reduction in the mean squared forecast error (MSFE) for the predictive model ($MSFE_{PRED}$) relative to the benchmark model ($MSFE_{BENCHMARK}$) and is given by

$$R_{OS}^2 = 100 \left(1 - \frac{MSFE_{PRED}}{MSFE_{BENCHMARK}} \right) = 100 \left[1 - \frac{\sum_{t=t_0}^{T-1} (r_{t+1} - \hat{r}_{t+1})^2}{\sum_{t=t_0}^{T-1} (r_{t+1} - \hat{r}_{t+1}^{BENCHMARK})^2} \right] ,$$

the quarter at which the forecasts are made. This step ensures that our method does not have a look-ahead bias, as the forecasts are made with current and past information only. When using a two-sided filter some assumptions as regards how to deal with the observations at the beginning and at the end of the sample has to be made. The literature suggests several types of boundary treatment rules to deal with boundary effects (e.g. periodic rule, reflection rule, zero padding rule, and polynomial extension). Here, we use a reflection rule, whereby the original time series are reflected symmetrically at the boundaries before filtering them.

where \hat{r}_{t+1} is the forecast for $t+1$ from the FC-TD or the FC-FD model (equation (3) and (5), respectively) and r_{t+1} is the realized equity premium / GDP growth from t to $t+1$. A positive (negative) R_{OS}^2 indicates that the predictive model outperforms (underperforms) the benchmark model in terms of MSFE. The benchmark model to forecast the equity premium is the average equity premium up to time t , and to forecast the GDP growth is an AR(p) model, where p is chosen recursively according to the Akaike information criterion.

The statistical significance of the R_{OS}^2 is evaluated using the Clark and West (2007) MSFE-adjusted statistic. This statistic tests the null hypothesis that the MSFE of the benchmark model is less than or equal to the MSFE of the FC-TD or FC-FD model against the alternative hypothesis that the MSFE of the benchmark model is greater than the MSFE of the FC-TD or FC-FD model ($H_0 : R_{OS}^2 \leq 0$ against $H_A : R_{OS}^2 > 0$).

3.4.2 Economic performance

Stock return forecasts should also be analyzed with utility-based metrics, which provide a more direct measure of the value of forecasts to decision makers. In these exercises, stock return forecasts are used as inputs for asset allocation decisions derived from expected utility maximization problems. A leading utility-based metric is the average utility gain for a mean-variance investor, who allocates her portfolio between equities and risk-free bills. At the end of quarter t , the investor optimally allocates a share $w_t = \hat{R}_{t+1} / (\gamma \hat{\sigma}_{t+1}^2)$ of the portfolio to equity for period $t+1$, where γ is the investor's relative risk aversion coefficient, \hat{R}_{t+1} is the time t (FC-TD or FC-FD) model forecast of equity premium,⁶ and $\hat{\sigma}_{t+1}^2$ is the forecast of the variance of the equity premium. We assume a relative risk aversion coefficient of three, use a five-year moving window of past equity premium to estimate the variance forecast and constrain the weights w_t to lie between -0.5 and 1.5. These constraints limit the possibilities of short selling and leveraging the portfolio to realistic levels.

⁶ As in Rapach, Ringgenberg and Zhou (2016), in the asset allocation exercise, we use simple (instead of log) returns. Therefore the change in notation from r to R .

The realized portfolio return at time $t+1$, RP_{t+1} , is given by $RP_{t+1} = w_t R_{t+1} + RF_{t+1}$, where RF_{t+1} denotes the risk-free return from time t to $t+1$ (*i.e.* the market rate, which is known at time t). The average utility (or certainty equivalent return, CER) is computed as $CER = \overline{RP} - 0.5\gamma\sigma_{RP}^2$, where \overline{RP} and σ_{RP}^2 are the sample mean and variance of the portfolio return, respectively. We report the annualized utility gain, which is computed as the difference between the CER for an investor that uses the FC-TD or FC-FD model to forecast the equity premium and the CER for an investor who uses the historical mean benchmark for forecasting. The difference is multiplied by 4 to annualize quarterly performance, which allows to interpret it as the annual portfolio management fee that an investor would accept to pay to have access to the alternative forecasting model versus the benchmark model forecast. Following Gargano, Pettenuzzo and Timmermann (2019) and Bianchi, Buchner, Tamoni and Nieuwerburgh (2021), we use a Diebold and Mariano (1995) test to assess if the annualized CER gains are statistically greater than zero.

4 Results

4.1 Equity premium

The second through seventh columns of table 3 reports the results for equity premium predictions for the full OOS 1965:Q1-2019:Q4 forecast evaluation period (second and third columns), and separately for NBER-dated business-cycle expansions (fourth and fifth columns) and recessions (sixth and seventh columns). Panel A show the results for individual predictive regression, while panel B for different forecast combination models.⁷

The R_{OS}^2 statistics in the second column clearly show that individual predictive regression forecasts of the equity premium (panel A) frequently fail to beat the benchmark model in terms of MSFE.

⁷ In panel B we do not report the results of the so-called kitchen sink model, which corresponds to a multiple predictive regression model that includes all fifteen economic variables as regressors, as it performs very poorly.

Indeed, 13 of the 15 R_{OS}^2 statistics are negative, and only one (IK) of the two predictors with a positive R_{OS}^2 (INFL and IK) is statistically significant. That is, only one of these 15 economic variables displays statistically significant OOS predictive ability at conventional levels. The third column reports the average utility gains. Relative to the R_{OS}^2 statistics, the individual predictive regression forecasts appear more valuable from an economic point of view, as 9 of the 15 economic variables offer positive gains. However, only one of them (*TMS*) has statistically significant positive annualized gains (369 basis points).

When analyzing stock return predictability over the business cycle, the R_{OS}^2 statistics in the fourth and sixth columns confirm the well-known fact that predictability is higher in recessions than in expansions (see e.g. Cujean and Hasler, 2017). Only one variable (LTR) delivers a positive and statistically significant R_{OS}^2 in expansions, while four variables (DP, DY, TMS, and IK) are remarkably good predictors of equity returns in recessions. Differences in predictability over business cycle phases are even more pronounced when looking at the utility gains (in the fifth and seventh columns).

Overall, the results in panel A between columns two and seven show that no single variable is clearly and constantly better (from a statistical or economic point of view) in all sub-samples than the others as equity premium predictor. In the same columns but in panel B are reported the R_{OS}^2 statistics and average utility gains for the combining methods for equity premium predictions. The first row in Panel B demonstrates the usefulness of the forecast combination in the time domain. The FC-TD model delivers positive and statistically significant (at conventional levels) R_{OS}^2 and positive CER gains (but not statistically significant) for the full forecast evaluation period, as well as in both expansion and recession periods.

The second row in Panel B shows the results of the forecast combination in the frequency domain (FC-FD). From a statistical point of view (R_{OS}^2), the FC-FD model performs slightly better than the FC-TD model over the full OOS period, while its performance is much better (slightly worse) than that of the FC-TD model in recessions (expansions). However, from an economic point of

view (CER gains), the FC-FD model delivers clearly larger utility gains regardless of the sample period considered. Over the full OOS period and during recessions, the CER gains generated by the FC-FD model are sizable, statistically significant (at conventional levels), and twice as large as than those of the FC-TD model (309 and 1510 basis points versus 142 and 650 basis points, respectively).

Figure 3 provides information on the behavior of the portfolios based on the forecasts from the benchmark model and from the forecast combination models. Panel A and B depicts equity weights and the log cumulative wealth, respectively, over the 1965:Q1-2019:Q4 forecast evaluation period.

The equity weight for the portfolio based on the benchmark model (black line) is relatively stable throughout the OOS period, which reflects the fact that the historical mean benchmark forecast is very smooth. The equity weight for the portfolio based on the FC-TD model (red line) exhibits substantial fluctuations around the weight of the benchmark portfolio, especially until 1990. After that, the weight closely follows the one of the benchmark portfolio.

The equity weight for the portfolio based on the FC-FD model (blue line) exhibits even more fluctuations, especially around recessions. The enhanced portfolio performance of the FC-FD model, quite evident from the log cumulative wealth in panel B, is due to its better market timing, as it allows to quickly reduce the exposure to the equity market around recessions. The terminal wealth in December 2019 is 128€ if the investor uses the historical mean forecast, 211€ if the investor uses the FC-TD predictive model, and 448€ if the FC-FD predictive model is used.

4.2 GDP growth

The eighth through tenth columns of table 3 reports the results for real GDP growth rate predictions for the full OOS forecast evaluation period (eighth column), and separately for NBER-dated business-cycle expansions (ninth column) and recessions (tenth column). Panel A show the results for individual predictive regression, while panel B for different forecast combination models.

Over the full OOS period, results are similar to the equity premium ones: only one of the individual predictor (NTIS) displays statistically significant out-of-sample predictive ability at conventional levels. Differently from the equity premium case, predictability of GDP growth rates is higher in expansions than in recessions: seven variables are good predictors of GDP growth in expansions, while only two in recessions.

Looking at the forecast combination methods in Panel B, the FC-FD model perform much better than the FC-TD model over the full OOS period. The R_{OS}^2 for the FC-FD model is 9.05%, which is about three times larger than the 3.18% R_{OS}^2 for the FC-TD model (both statistically significant at the 1% level). Both forecast combination methods perform very well in expansions (with the FC-FD model being better than the FC-TD model) but rather poorly in recessions (negative R_{OS}^2 statistics).

The reason for this poor performance of the forecast combination methods in recessions is clear when looking at the actual forecasts, reported in figure 4. The black solid line is real GDP growth rate over the full 1965:Q1-2019:Q4 forecast evaluation period, the red (blue) line is the forecast of the FC-TD (FC-FD) model, and grey bars denote NBER-dated recessions. Indeed, none of the forecast combination methods is indeed able to capture the sudden and abrupt drops associated with recessions. From figure 4 it is also possible to note that the enhanced performance of the FC-FD model is due on its ability to capture well both the trend of GDP growth and some of its higher frequency fluctuations.

4.3 Placebo test

To demonstrate that our procedure does not mechanically generate predictability, we run the following placebo test. We generate 1000 datasets, each of them containing 15 variables, and each variable having the same persistence and standard deviation as the respective variable in the real dataset. Innovations in the simulated datasets are produced by a random number generator so they

are independent from the true data. Then we run the forecast with our FC-FD model to each of these simulated datasets, and record the R_{OS}^2 and CER gains.

Figure 5 shows the distribution of the R_{OS}^2 and CER gains for equity premium predictions (left and middle graph, respectively) and the R_{OS}^2 for GDP growth rate predictions. For all measures and variables, the medians (red lines within each box) as well as the entire distributions (vertical dashed lines) are below the results with the original data (black dots). This placebo analysis thus shows that the predictability power of the FC-FD model is unlikely to be driven by a mechanical bias.

4.4 Why is it important to take into account the frequency domain (in forecast combination)?

The benefits of using combination forecast methods are well known in the literature. As stressed in the seminal paper by Goyal and Welch (2008), the inconsistent out-of-sample performance of individual predictive regression models is due to of structural instability. The graphs in the top row in Figure 6 give a visual impression of the changing nature of the relationships between the target variables (equity premium on the left and GDP growth on the right) and three individual economic variables (DP, black lines; TMS, blue lines; IK, red lines). The figure depicts the OLS estimates from expanding windows that start with the sample 1947:Q1-1964:Q4 and recursively add one quarter through 2019:Q4. The regression coefficients, which are ultimately used to produce the OOS forecasts with the FC-TD model, fluctuate substantially over the period, and there are even instances where the sign switches from positive to negative (or the other way around). Given this instability over time, averaging across individual forecasts gives more stable and, ultimately, better forecasting results.

The remaining rows in Figure 6 report the OLS estimates in each frequency bands, which are used to produce the OOS forecasts with the FC-FD model. We emphasize three features about time- and frequency-varying changing relationships. First, as it happens with the original variables, there is a

clear time variation of each coefficient within a specific frequency, and switching sign is also quite common. Second, there are cases where the sign of the relationship between aggregate variables is different from the sign between the same variables at different frequencies. For instance, the estimated OLS coefficients are negative (positive) between the equity premium and IK (GDP growth and TMS), but the sign flips at frequency D_2 (D_5). Third, for a given predictor, the magnitude of the estimated coefficients also significantly varies across frequencies and are quite often different from the magnitude using the original series.⁸

Overall, figure 6 suggests important structural instabilities in the relationships between the target variables and these predictors not only across time, but also across frequencies. These findings support the relevance of taking the frequency domain into account, as the magnitude and sign of the estimated coefficients (that are used to make the forecasts) are time- and frequency-specific. As Faria and Verona (2021) show, taking these facts into account is of key importance in an OOS exercise.

Another way to understand why forecast combination performs better than individual predictive regression models is to look at the Theil (1971) MSFE decomposition into the squared forecast bias and a remainder term. The latter term depends, among other things, on the forecast volatility, and limiting forecast volatility helps to reduce the remainder term. A model's forecasting performance ultimately depends on the tradeoff between the reduction in bias and variance. To get a sense of this bias-efficiency trade-offs in the forecasts, figure 7 is a scatterplot depicting the MSFE decomposition into the squared forecast bias and the remainder term for the individual predictive regression models, the benchmark models, and the combination models for the full OOS period.

Looking at the equity premium forecast (left graph), several forecasting strategies produce relatively unbiased return predictions, many of them even better than the historical mean benchmark.

⁸ It is beyond the scope of this paper to analyze why, for some of these variables, the sign and magnitude at some frequencies are different than those with the original series. However, this fact has already been emphasized in other applications. For instance, in the context of the Q theory of investment, Gallegati and Ramsey (2013) and Verona (2020) show that the investment-Q sensitivity is not always positive at all points in time and for all frequencies.

However, their performance relative to the historical mean is negatively affected by their higher remainder term.

The FC-TD forecast has a lower forecast variance than all of the individual predictive regression models and a relatively small squared forecast bias (close to the smallest squared biases of the individual predictive regression model, IK). When compared with the historical mean benchmark, both the squared bias and the remainder term are substantially below. Hence, the FC-TD model achieves a higher R_{OS}^2 (that is, a smaller MSFE) than the historical mean benchmark and any of the individual predictive regression models.

The FC-FD model delivers more accurate forecasts (higher R_{OS}^2 and smaller MSFE) than the FC-TD model due to its ability to further reduce both the forecast bias and the remainder term.

Similar conclusions can be drawn from the analysis of the scatterplot for real GDP growth forecast in the right graph of figure 7. Thus, forecasts based on the FC-FD model are generally both less biased and more efficient than all the other forecasts analyzed here, including the forecast combination in the time domain.

5 Predictability over the business cycle

So far we sum all (five) frequencies when making the forecast in the frequency domain, *i.e.* $FC - FD_{t+1} = \sum_{f=1}^5 \hat{r}_{c,t+1}^{D_f}$. However, as shown by Faria and Verona (2021) and Martins and Verona (2021), ignoring some frequencies of the target variable leads to better forecasts. For instance, forecasting with a model like $FC - FD_{t+1} = \sum_{f=2}^4 \hat{r}_{c,t+1}^{D_f}$, which ignores both the highest and the lowest frequency forecasts of the target variable and only uses the business-cycle frequencies to make the forecasts, might produce more accurate forecast than an identical model that sums all (five) frequencies.

We now exploit the flexibility of our method and check if and when it is profitable to ignore some

frequencies of the target variables. We start from an ex-post exercise to gain some intuition about the predictability over business cycle phases. We then move to a real-time exercise where the status of the business cycle is assessed in real time and the forecaster switches between two forecasts according to the state of the economy.

5.1 Ex-post exercise

To gain some idea about how our method performs in recessions and in expansions, we analyse which frequencies of the target variables are important to include (or exclude) to have good forecasts, and whether there are differences between expansions and recessions.

The second to last row in table 3 reports the forecasts of the FC-FD model when the forecast is given by $FC - FD_{t+1} = \sum_{f=1}^4 \hat{r}_{c,t+1}^{D_f}$. This is the case when the low-frequency component (D_5) of the target variable is ignored when making the forecast (we denote this method as FC-FD (no LF)). The gains over the full OOS period are not very impressive (for the equity premium) or really bad (for GDP growth). However, for both target variables, there are huge forecasting gains in ignoring the low-frequency forecast of the target variables in recessions. The intuition is that ignoring the forecast of the trend allows to better track the quick and sudden drop associated with recessions (recall that both variables are negatively skewed). However, it is crucial to forecast well the trend (as well as some high-frequency fluctuations) in expansions, and ignoring it leads to bad forecasts.⁹

However, this finding of enhanced return predictability during recessions is ex post, since the NBER dates of business-cycle peaks and troughs are known retrospectively. The question is now whether this information is useful in real time. In particular, how large are the statistical / economic gains if we were able to switch between the forecasts of two different frequency combinations – $FC - FD_{t+1} = \sum_{f=1}^5 \hat{r}_{c,t+1}^{D_f}$ for expansions and $FC - FD_{t+1} = \sum_{f=1}^4 \hat{r}_{c,t+1}^{D_f}$ for recessions – in real time

⁹ Regardless of the forecasts of the other frequency components $D_1 - D_4$, we find similar results for almost all possible frequency combinations in the spirit of Faria and Verona (2021): whenever we exclude (include) the low-frequency forecasts (D_5) of the target variable, the forecasts in recessions are much better (worse) and in expansions are much worse (better).

according to the perception of the state of the business cycle? We address this question in the next subsection.

5.2 Real-time exercise

To guide switching between forecasts over the business cycle, we rely on a well-known leading indicator of the business cycle – the stock market. In particular, we use the information from some stock market technical indicators (TIs), which are widely employed by practitioners, to compute a real-time indicator of the state of the business cycle. Technical indicators rely on past stock market price and volume patterns to identify trends believed to persist into the future, so they provide useful forward looking information about the business cycle.¹⁰

Following Neely, Rapach, Tu and Zhou (2014), we use two moving average indicators, two momentum indicators, and two volume indicators.¹¹ A value of 0 (1) for each of these indicators implies a sell (buy) signal at the end of quarter t , hence quarter $t+1$ is considered to be a recession (expansion) according to this specific technical indicator. Relying on a single TI might however generate too many false recession signals. Hence, we introduce a novel business cycle leading indicator, that we name as coincident index, that summarizes the information from the six TIs. In particular, for quarter $t+1$ to be considered a recession, all six TIs have to be 0 at the end of quarter t . In this case, we use the forecasts for $t+1$ (made at the end of quarter t) that exclude the low-frequency forecast of the target variable (i.e. we use $FC - FD_{t+1} = \sum_{f=1}^4 \hat{r}_{c,t+1}^{Df}$). The coincident index, plotted in figure 8, captures most of the actual NBER-dated recession quarters, albeit triggering some false recession signals.

¹⁰ This method is similar in spirit but much simpler than the one proposed by Aruoba, Diebold and Scotti (2009), who use high-frequency data to compute a real-time indicator of economic activity. Other variables commonly used a real-time / leading indicator of the business cycle are the Chicago Fed National Activity Index, the Business Conditions Index, the term spread, and indicators based on survey data (Survey of Professional Forecasters, Livingston Survey, and Purchasing Managers' index). Markov-switching models (see e.g. Guidolin and Timmermann, 2007) provide a different framework for switching between forecasting models according to estimated probabilities of the state of the economy.

¹¹ These technical indicators are described and plotted in appendix 2.

The results for the equity premium and GDP growth forecasts in real time, denoted FC-FD real time, are reported in the last row in table 3. Being able to switch forecasts in real time according to the state of business cycle allows to improve forecast even further when compared to the FC-FD model. For the equity premium (second and third column), the R_{OS}^2 statistic and CER gain both are statistically significant and are 4.73% and 427 basis points, respectively. The blue dashed line in figure 3 reports the log cumulative wealth for an investor who trades using the FC-FD real time model. The terminal wealth in December 2019 is 655€: the possibility of being able to switch between forecasts significantly increases overall profits throughout the entire OOS period.

Looking at the real GDP growth rate forecasts (eighth column), the R_{OS}^2 statistic of the FC-FD real time model is statistically significant and markedly higher than those of the other combination models (15.9% against 9.05% and 3.18% of the FC-FD and FC-TD model, respectively). The blue dashed line in figure 4 shows the forecast of the FC-FD real time model. This method produces more accurate forecast as it allows to capture remarkably well the drops associated with recessions. Likewise, from the bias-variance scatterplot (figure 7), the improved performance of the FC-FD real time model is due to its ability to cut the forecast bias almost to zero and to reduce even further the remainder term.

6 Robustness

We run the following robustness tests.

The choice of the number of frequencies and of the band-pass filter to use to filter the data affect both the equity premium and real GDP growth forecast. We run the analysis with 4 or 6 frequencies (instead of 5),¹² and use the Daubechies filter of length two and four.

¹² In this case, frequencies $D_1 - D_3$ are the same. When using 4 frequencies, frequency component D_4 captures fluctuations with a period longer than 4 years. When using 6 frequencies, frequency component D_5 captures fluctuations with a period of 8-16 years, while frequency component D_6 captures fluctuations with a period longer than 16 years.

The choice of the parameters related with the asset allocation exercise only affects the CER gain results. We consider a risk aversion coefficient of 5 (instead of 3), different set of portfolio constraints for the equity weights w_t (no leverage and/or no short selling instead of 50% leverage and short selling), a ten-year (instead of five-year) moving window of past equity premium to estimate the variance forecast, and CER gains net of transactions costs of 50 basis points.

For the real-time exercise, we use different technical indicators to compute the coincident index.

Results turn out to be robust to all these changes, so we do not report them here (but they available upon request).

Finally, we add year 2020 to our sample. The COVID-19 recession does not (significantly) change the results as regards equity premium predictability. The drop and recover of the equity market during the COVID-19 recession were, in fact, comparable with (or even less abrupt than) those in previous recessions. Furthermore, as we are using quarterly data, all high frequency stock market fluctuations are, by construction, smoothed. Quarter-to-quarter GDP growth, on the other hand, experienced fluctuations ranging from -9.4% in 2020:Q2 to 7.3% in 2020:Q3. In comparison, in the global financial crises, it fluctuated between -2.2% in 2008:Q4 and 1.1% in 2009:Q4. These huge fluctuations in 2020 render forecast (evaluation) extremely difficult. The forecast combination models perform overall well but their R_{OS}^2 statistics are not statistically different from those of the benchmark models. However, the R_{OS}^2 statistic for the FC-FD real time model is positive (27.1%) and statistically significant (at the 5% level). So, despite these large GDP fluctuations, our real-time model keeps providing reliable GDP growth rate forecasts.

7 Conclusion

In this paper we propose a new forecasting method – forecast combination in the frequency domain – that takes into account that predictability is time *and* frequency dependent. We apply this method

to forecast the equity premium and real GDP growth rate. Combining forecasts in the frequency domain produces markedly more accurate predictions relative to the traditional forecast combination in the time domain, both in terms of statistical and economic measures of out-of-sample predictability. This method is flexible enough that it allows to exclude some of the frequencies of the target variables when making the forecasts. In particular, we show that, in recessions, it is of major relevance to ignore the forecast of the low-frequency components of the target variables. This flexibility turns out to be crucial in a real-time forecasting exercise, as it allows to capture remarkably well the sudden and abrupt drops associated with recessions and further improve predictability of the equity premium and real GDP growth rate.

References

- Aruoba, S. Boragan, Francis X. Diebold, and Chiara Scotti**, “Real-Time Measurement of Business Conditions,” *Journal of Business & Economic Statistics*, 2009, 27 (4), 417–427.
- Bandi, Federico, Bernard Perron, Andrea Tamoni, and Claudio Tebaldi**, “The Scale of Predictability,” *Journal of Econometrics*, 2019, 208 (1), 120–140.
- Bates, J. M. and C. W. J. Granger**, “The Combination of Forecasts,” *Operational Research Quarterly*, 1969, 20 (4), 451–468.
- Bianchi, Daniele, Matthias Buchner, Andrea Tamoni, and Stijn Van Nieuwerburgh**, “Bond Risk Premiums with Machine Learning,” *Review of Financial Studies*, 2021, 34 (2), 1046–1089.
- Caggiano, Giovanni, George Kapetanios, and Vincent Labhard**, “Are more data always better for factor analysis? Results for the euro area, the six largest euro area countries and the UK,” *Journal of Forecasting*, December 2011, 30 (8), 736–752.
- Campbell, John Y. and Samuel B. Thompson**, “Predicting Excess Stock Returns Out of Sample: Can Anything Beat the Historical Average?,” *Review of Financial Studies*, 2008, 21 (4), 1509–1531.
- Clark, T.E. and K.D. West**, “Approximately normal tests for equal predictive accuracy in nested models,” *Journal of Econometrics*, 2007, 138 (1), 291 – 311.
- Cremers, K. J. Martijn**, “Stock Return Predictability: A Bayesian Model Selection Perspective,” *Review of Financial Studies*, 2002, 15 (4), 1223–1249.
- Cujean, Julien and Michael Hasler**, “Why Does Return Predictability Concentrate in Bad Times?,” *Journal of Finance*, 2017, 72 (6), 2717–2758.

- Diebold, F.X. and R.S. Mariano**, “Comparing Predictive Accuracy,” *Journal of Business & Economic Statistics*, 1995, 13 (3), 253–263.
- Engle, Robert F.**, “Band Spectrum Regression,” *International Economic Review*, 1974, 15 (1), 1–11.
- Estrella, Arturo and Gikas A Hardouvelis**, “The Term Structure as a Predictor of Real Economic Activity,” *Journal of Finance*, June 1991, 46 (2), 555–576.
- Faria, Gonçalo and Fabio Verona**, “Forecasting stock market returns by summing the frequency-decomposed parts,” *Journal of Empirical Finance*, 2018, 45, 228 – 242.
- **and** — , “The yield curve and the stock market: Mind the long run,” *Journal of Financial Markets*, 2020, 50 (C).
- **and** — , “Out-of-sample time-frequency predictability of the equity risk premium,” *Quantitative Finance*, 2021, 21 (12), 2119–2135.
- Gallegati, Marco and James B. Ramsey**, “Bond vs stock market’s Q: Testing for stability across frequencies and over time,” *Journal of Empirical Finance*, 2013, 24 (C), 138–150.
- Gargano, A., D. Pettenuzzo, and A. Timmermann**, “Bond Return Predictability: Economic Value and Links to the Macroeconomy,” *Management Science*, 2019, 65 (2), 508–540.
- Goyal, A. and I. Welch**, “A Comprehensive Look at The Empirical Performance of Equity Premium Prediction,” *Review of Financial Studies*, 2008, 21 (4), 1455–1508.
- Guidolin, Massimo and Allan Timmermann**, “Asset allocation under multivariate regime switching,” *Journal of Economic Dynamics and Control*, November 2007, 31 (11), 3503–3544.
- Henkel, Sam James, J. Spencer Martin, and Federico Nardari**, “Time-varying short-horizon predictability,” *Journal of Financial Economics*, 2011, 99 (3), 560–580.

- Lubik, Thomas A., Christian Matthes, and Fabio Verona**, “Assessing U.S. aggregate fluctuations across time and frequencies,” Research Discussion Papers 5/2019, Bank of Finland February 2019.
- Martins, Manuel M. F. and Fabio Verona**, “Inflation dynamics and forecast : frequency matters,” Research Discussion Papers 8/2021, Bank of Finland June 2021.
- Neely, C., D. Rapach, J. Tu, and G. Zhou**, “Forecasting the Equity Risk Premium: The Role of Technical Indicators,” *Management Science*, 2014, 60 (7), 1772–1791.
- Pettenuzzo, Davide and Francesco Ravazzolo**, “Optimal Portfolio Choice Under Decision-Based Model Combinations,” *Journal of Applied Econometrics*, 2016, 31 (7), 1312–1332.
- Pirschel, Inske and Maik H. Wolters**, “Forecasting with large datasets: compressing information before, during or after the estimation?,” *Empirical Economics*, September 2018, 55 (2), 573–596.
- Rapach, David E., Jack K. Strauss, and Guofu Zhou**, “Out-of-Sample Equity Premium Prediction: Combination Forecasts and Links to the Real Economy,” *Review of Financial Studies*, February 2010, 23 (2), 821–862.
- , **Matthew C. Ringgenberg, and Guofu Zhou**, “Short interest and aggregate stock returns,” *Journal of Financial Economics*, 2016, 121 (1), 46 – 65.
- Stock, James H. and Mark W. Watson**, “Macroeconomic Forecasting Using Diffusion Indexes,” *Journal of Business & Economic Statistics*, 2002, 20 (2), 147–162.
- **and** —, “Forecasting Output and Inflation: The Role of Asset Prices,” *Journal of Economic Literature*, September 2003, 41 (3), 788–829.
- **and** —, “Combination forecasts of output growth in a seven-country data set,” *Journal of Forecasting*, 2004, 23 (6), 405–430.

Theil, Henri, *Applied Economic Forecasting*, North-Holland, Amsterdam, 1971.

Timmermann, Allan, “Forecast Combinations,” in G. Elliott, C. Granger, and A. Timmermann, eds., *Handbook of Economic Forecasting*, Vol. 1 of *Handbook of Economic Forecasting*, Elsevier, 2006, chapter 4, pp. 135–196.

Verona, Fabio, “Investment, Tobin’s Q, and Cash Flow Across Time and Frequencies,” *Oxford Bulletin of Economics and Statistics*, 2020, 82 (2), 331–346.

	mean	median	1st perc.	99th perc.	std. dev.	skew.	kurt.	AR(1)
Equity premium	0.02	0.03	-0.24	0.17	0.08	-0.97	4.94	0.09
Real GDP growth	0.77	0.75	-1.87	3.56	0.93	-0.07	4.58	0.36
DP	-3.50	-3.47	-4.47	-2.64	0.44	-0.10	2.29	0.98
DY	-3.48	-3.46	-4.48	-2.60	0.44	-0.10	2.36	0.98
EP	-2.77	-2.83	-4.30	-1.88	0.45	-0.59	5.42	0.95
DE	-0.73	-0.73	-1.23	0.65	0.29	2.86	20.9	0.90
SVAR	0.01	0.00	0.00	0.04	0.01	7.81	80.0	0.42
BM	0.53	0.51	0.14	1.13	0.25	0.52	2.48	0.98
NTIS	0.01	0.02	-0.05	0.04	0.02	-0.90	3.48	0.94
TBL	0.04	0.04	0.00	0.15	0.03	0.99	4.26	0.96
LTY	0.06	0.05	0.02	0.14	0.03	0.85	3.25	0.98
LTR	0.02	0.01	-0.09	0.20	0.05	0.97	6.04	-0.03
TMS	0.02	0.02	-0.02	0.04	0.01	-0.09	3.14	0.84
DFY	0.01	0.01	0.00	0.03	0.00	1.93	8.60	0.88
DFR	0.00	0.00	-0.07	0.05	0.02	0.24	14.5	-0.09
INFL	0.01	0.01	-0.01	0.04	0.01	0.31	5.93	0.41
IK	0.04	0.03	0.03	0.04	0.00	0.42	2.70	0.97

Table 1: Summary statistics, U.S. data, 1947:Q1-2019:Q4

This table reports summary statistics for the (log) equity premium, real GDP growth, and for the 15 predictive variables. See appendix A for a description of the predictors.

Predictor	Variable to forecast r	Predictor	Variable to forecast			
			r	r^{D_1}	...	r^{D_5}
DP	$\hat{r}_{DP,t+1} = \hat{\alpha}_{DP} + \hat{\beta}_{DP} DP_t$	DP^{D_j}	$\hat{r}_{t+1}^{D_1,DP} = \hat{\alpha}_{t,1}^{DP} + \hat{\beta}_{t,1}^{DP} DP_t^{D_1}$...	$\hat{r}_{t+1}^{D_5,DP} = \hat{\alpha}_{t,5}^{DP} + \hat{\beta}_{t,5}^{DP} DP_t^{D_5}$	
DY	$\hat{r}_{DY,t+1} = \hat{\alpha}_{DY} + \hat{\beta}_{DY} DY_t$	DY^{D_j}	$\hat{r}_{t+1}^{D_1,DY} = \hat{\alpha}_{t,1}^{DY} + \hat{\beta}_{t,1}^{DY} DY_t^{D_1}$...	$\hat{r}_{t+1}^{D_5,DY} = \hat{\alpha}_{t,5}^{DY} + \hat{\beta}_{t,5}^{DY} DY_t^{D_5}$	
\vdots	\vdots	\vdots	\vdots	...	\vdots	
IK	$\hat{r}_{IK,t+1} = \hat{\alpha}_{IK} + \hat{\beta}_{IK} IK_t$	IK^{D_j}	$\hat{r}_{t+1}^{D_1,IK} = \hat{\alpha}_{t,1}^{IK} + \hat{\beta}_{t,1}^{IK} IK_t^{D_1}$...	$\hat{r}_{t+1}^{D_5,IK} = \hat{\alpha}_{t,5}^{IK} + \hat{\beta}_{t,5}^{IK} IK_t^{D_5}$	
	$FC - TD_{t+1} =$ $\frac{1}{n} \sum_{i=1}^n \hat{r}_{x_i,t+1}$		$\hat{r}_{c,t+1}^{D_1} =$ $\frac{1}{n} \sum_{i=1}^n \hat{r}_{t+1}^{D_1,x_i}$...	$\hat{r}_{c,t+1}^{D_5} =$ $\frac{1}{n} \sum_{i=1}^n \hat{r}_{t+1}^{D_5,x_i}$	
			$FC - FD_{t+1} =$ $\sum_{j=1}^5 \hat{r}_{c,t+1}^{D_j}$			

Table 2: Forecast combination methods

This table summarizes how the forecasts with the forecast combination model in the time domain (FC-TD) and in the frequency domain (FC-FD) are computed.

	Equity premium						Real GDP growth rate		
column	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
predictor / method	1965:Q1-2019:Q4		Expansions		Recessions		1965:Q1-2019:Q4	Expansions	Recessions
	R^2_{OS}	CER gain	R^2_{OS}	CER gain	R^2_{OS}	CER gain	R^2_{OS}	R^2_{OS}	R^2_{OS}
<i>Panel A: Individual predictive regression</i>									
DP	-0.20	-1.69	-5.21	-3.83	6.97***	8.74**	-5.02	13.3***	-34.0
DY	-0.16	-0.78	-6.52	-3.95	8.96***	15.3***	-4.48	15.0***	-35.3
EP	-1.15	0.58	-2.23	-0.39	0.39	5.19	-5.99	10.5***	-32.0
DE	-1.87	0.18	-3.52	-1.41	0.49	8.33	-10.5	-1.79	-24.1
SVAR	-9.80	-0.86	-13.9	-0.19	-3.87	-4.57	-42.6	-76.6	11.3***
BM	-2.09	-0.39	-2.48	-0.58	-1.52	0.11	-6.70	-1.23	-15.3
NTIS	-2.16	-0.85	0.02	0.29**	-5.29	-6.14	2.54***	13.2***	-14.2
TBL	-1.96	2.17	-2.84	0.08	-0.71	13.0	-7.62	-10.3	-3.41
LTY	-1.94	1.35	-2.23	-0.58	-1.51	11.4	-7.27	-5.77	-9.63
LTR	-0.51	0.14	2.01**	0.31	-4.12	-0.75	-26.1	-11.3	-49.3
TMS	-2.77	3.69***	-10.6	0.47*	8.40**	20.2**	-10.1	4.96***	-33.8
DFY	-2.41	-0.16	-2.31	0.07*	-2.54	-1.26	-3.00	-13.2	13.1*
DFR	-1.13	0.43	-5.28	-1.87	4.80	12.2	-6.61	0.64***	-18.1
INFL	0.30	0.50	-0.52	-0.24	1.46	4.68	-0.43	7.94***	-13.6
IK	2.36**	2.11	-2.01	-0.85	8.61**	17.4	-18.1	-26.3	-5.25
<i>Panel B: Forecast combination regression</i>									
FC-TD	2.67***	1.42	1.65**	0.45	4.12**	6.50	3.18***	11.4***	-9.87
FC-FD	3.04**	3.09*	1.28*	0.73	5.57**	15.1*	9.05***	15.8***	-1.65
FC-FD (no LF)	1.69***	-0.32	-5.64	-5.99	12.2***	29.8**	-73.1	-155	56.7***
FC-FD real time	4.73***	4.27*					15.9***		

Table 3: Equity premium and real GDP growth rate out-of-sample forecasting results

R^2_{OS} is the Campbell and Thompson (2008) out-of-sample R^2 statistic. CER gain is the portfolio management fee (in annualized percentage return) that an investor with mean-variance preferences and risk aversion coefficient of three would be willing to pay to have access to the forecasting model given in column (2), (4), or (6) relative to the benchmark forecasting model; the weight on stocks in the investor's portfolio is restricted to lie between -0.50 and 1.50. Statistical significance for the R^2_{OS} statistic is based on the p-value for the Clark and West (2007) out-of-sample MSFE-adjusted statistic; the statistic corresponds to a one-sided test of the null hypothesis that the competing forecasting model has equal expected square prediction error relative to the benchmark forecasting model against the alternative hypothesis that the competing forecasting model has a lower expected square prediction error than the benchmark forecasting model. Statistical significance for the CER gains is based on a one-sided Diebold and Mariano (1995) test. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

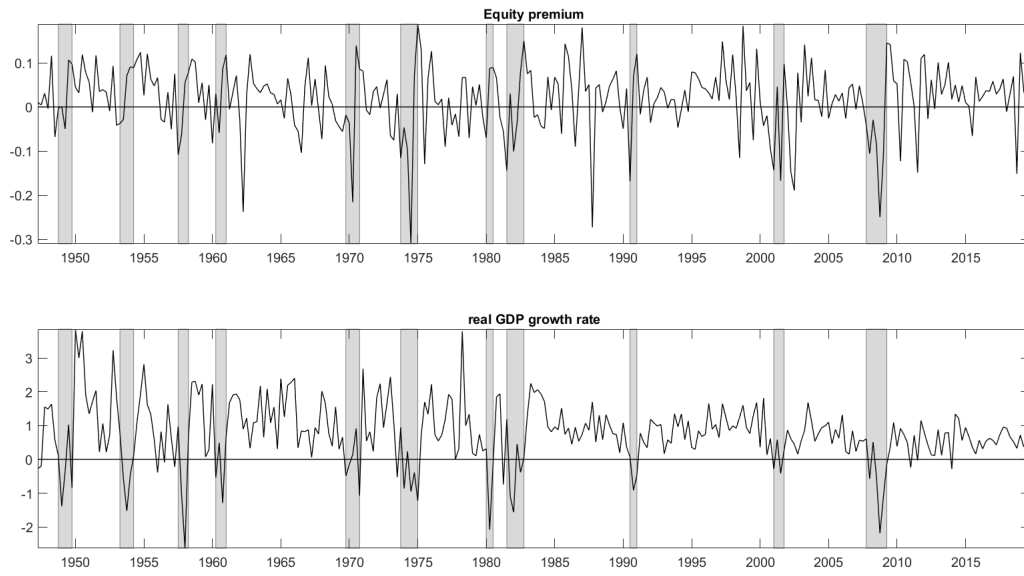


Figure 1: Equity premium and real GDP growth rate, U.S. data, 1947:Q1-2019:Q4
Time series of the quarterly log equity premium (upper graph) and quarterly real GDP growth rate (lower graph). Grey bars depict NBER-dated recessions.

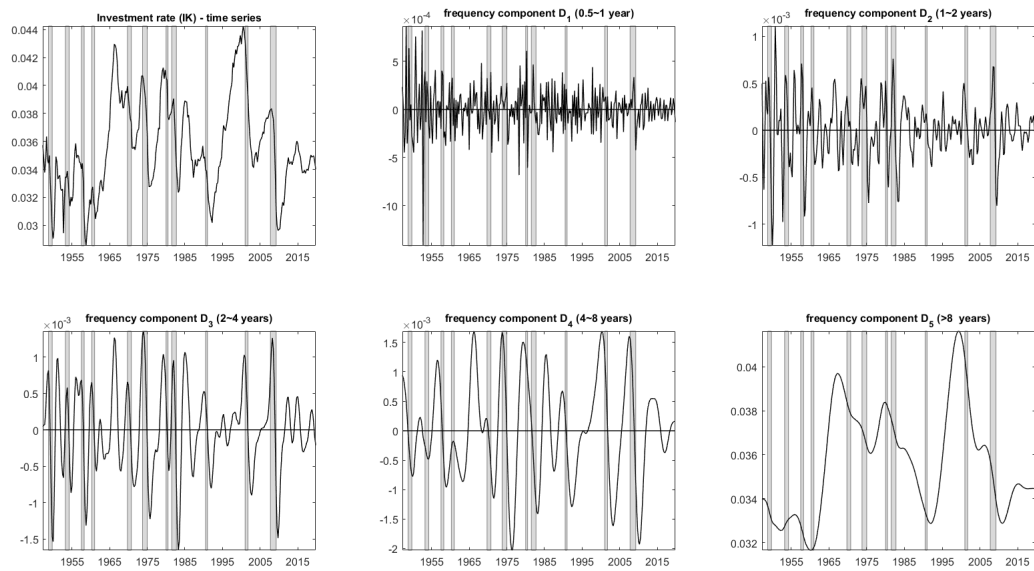
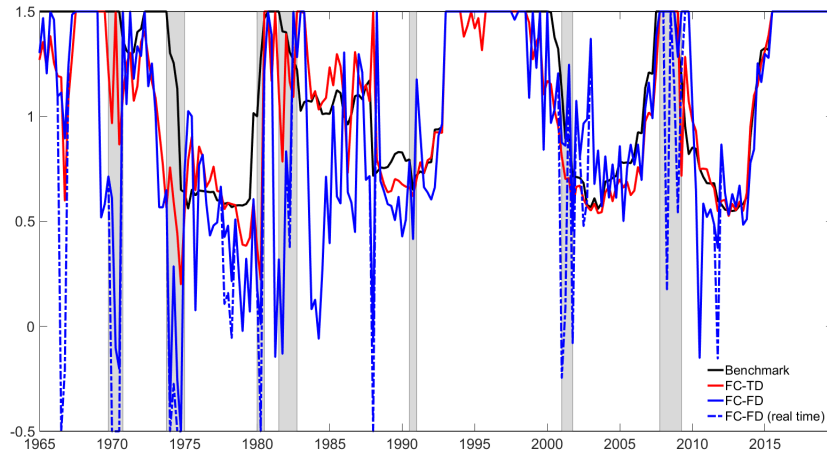


Figure 2: Investment rate (IK), time series and frequency decomposition, U.S. data, 1947:Q1-2019:Q4

The top left panel shows the time series of quarterly U.S. investment rate, while the remaining panels show the five time series frequency components into which the investment rate series is decomposed. Grey bars depict NBER-dated recessions.

A. Equity weights



B. Log cumulative wealth

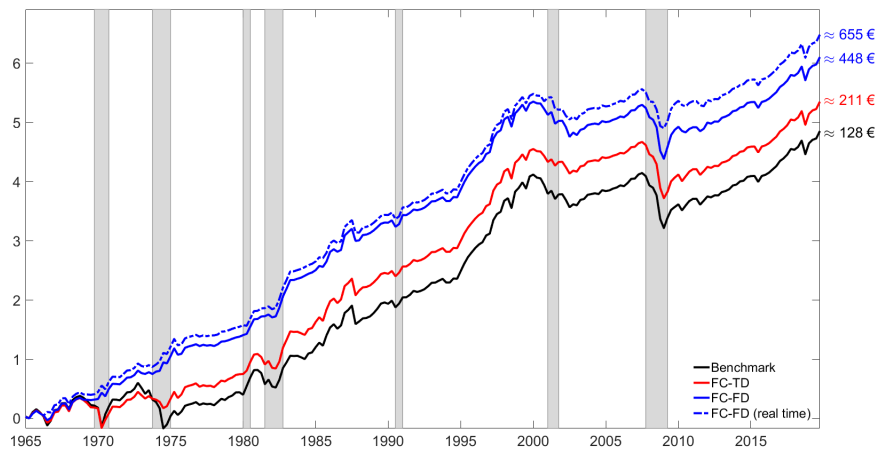


Figure 3: Equity weights and log cumulative wealth, 1965:Q1-2019:Q4

Panel A plots the dynamics of the equity weights for a mean-variance investor with relative risk aversion coefficient of three who allocates quarterly between equities and risk-free bills using a predictive regression excess return forecast based on the benchmark forecast model (black solid line), the FC-TD model (red solid line), the FC-FD model (blue solid line), or the FC-FD real time model (blue dashed line). The equity weights are constrained to lie between -0.5 and 1.5 . Panel B delineates the corresponding log cumulative wealth for the investor, assuming that the investor begins with 1€ and reinvests all proceeds. Grey bars denote NBER-dated recessions.

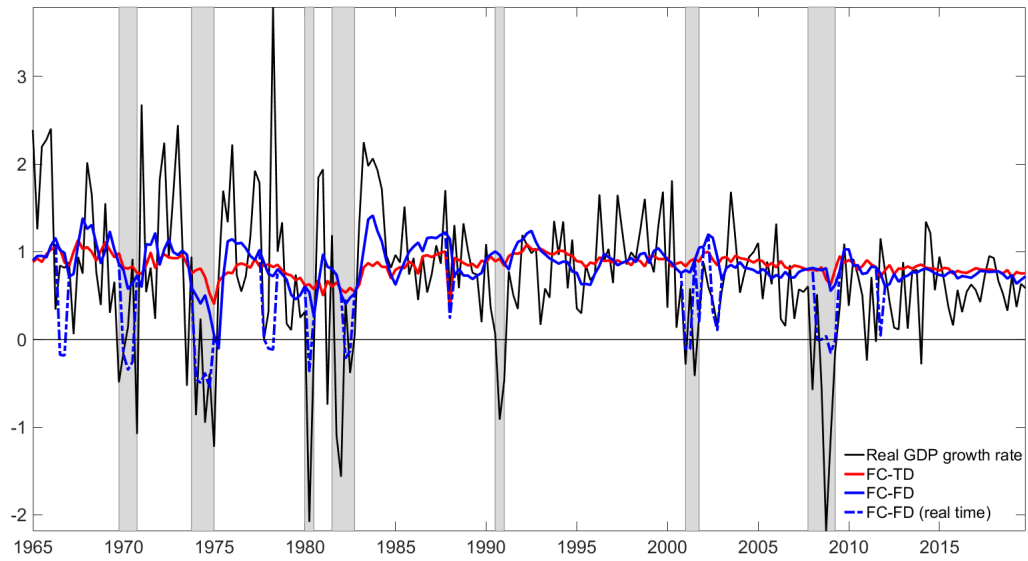


Figure 4: Real GDP growth rate, realized and out-of-sample forecasts, 1965:Q1-2019:Q4
 Quarterly U.S. real GDP growth rate (black solid line) and its out-of-sample forecasts based on the FC-TD model (red solid line), the FC-FD model (blue solid line), or the FC-FD real time model (blue dashed line). Grey bars depict NBER-dated recessions.

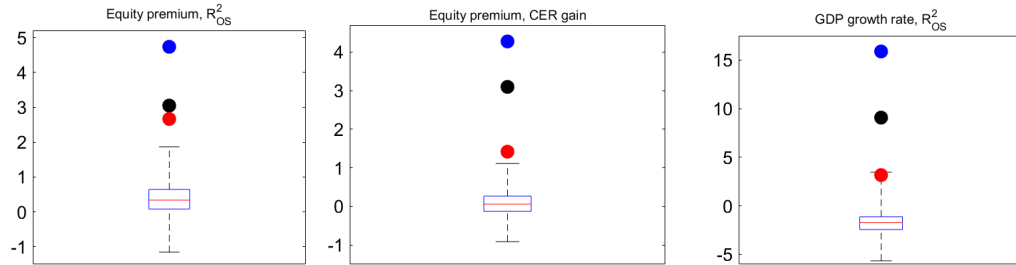


Figure 5: Placebo test R_{OS}^2 and CER gains

For each simulation, the boxes show the middle 75% of outcomes; the red line within the box is the median of the distribution; the vertical dashed lines represent the 12.5% simulated results above and below the central 75%. Red, black, and blue dots denote the results (as reported in Table 3) with the FC-TD model, the FC-FD model, and the FC-FD real time model, respectively.

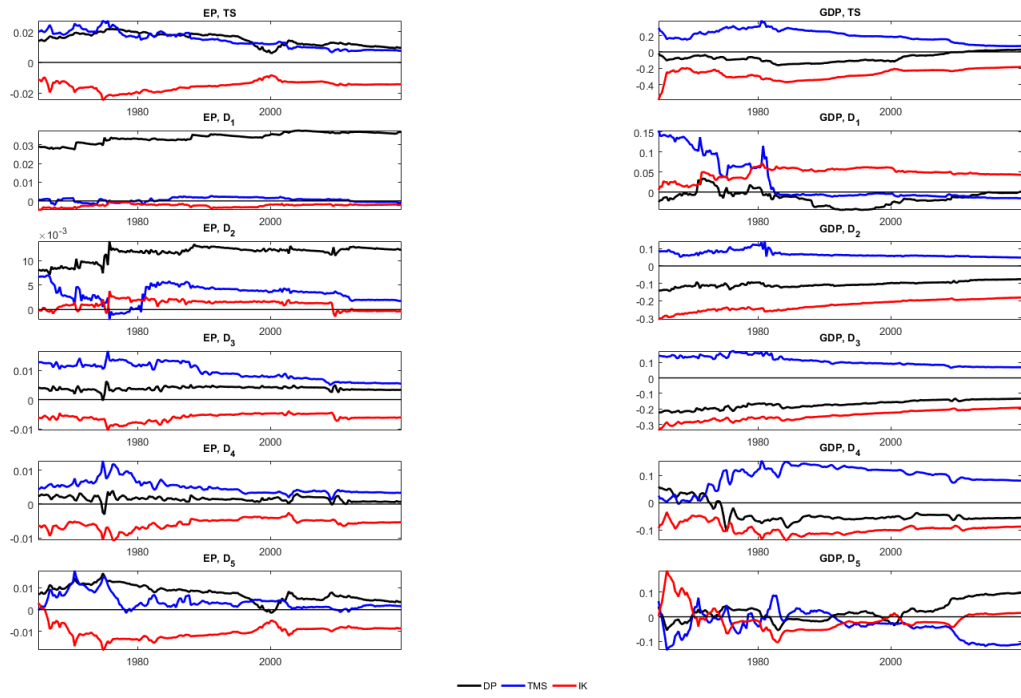


Figure 6: OLS regression coefficients between the target variables and three individual predictors based on expanding window estimates

Left (right) column: OLS regression coefficients (based on expanding window estimates starting in 1947:Q2-1964:Q4, recursively including one additional quarter through 2019:Q4) between the equity premium (real GDP growth) and three individual predictors (DP, black lines; TMS: blue lines; IK, red lines). Top row: time series. Remaining rows: coefficients in each frequency band. Each predictor variable is standardized to have a standard deviation of one before running the estimation.

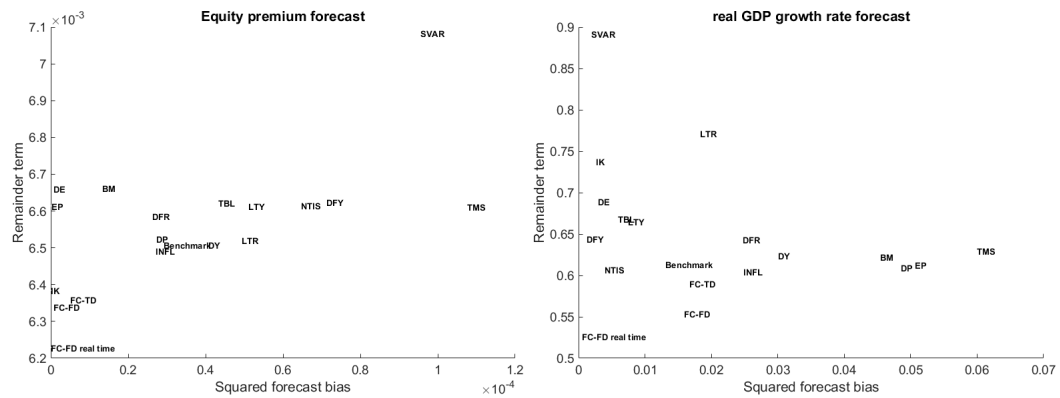


Figure 7: Scatterplot of the Theil (1971) MSFE decomposition into the squared forecast bias and a remainder term, 1965:Q1-2019:Q4

Left (right) graph: equity premium (real GDP growth rate) forecast. Benchmark corresponds to the benchmark forecast model (historical mean for equity premium, and AR(p) for GDP growth), and FC-TD, FC-FD, and FC-FD real time denote the combination forecast in the time domain, in the frequency domain, and in the frequency domain in the real time exercise in section 5.2, respectively. The other points correspond to the individual predictive regression model forecasts.

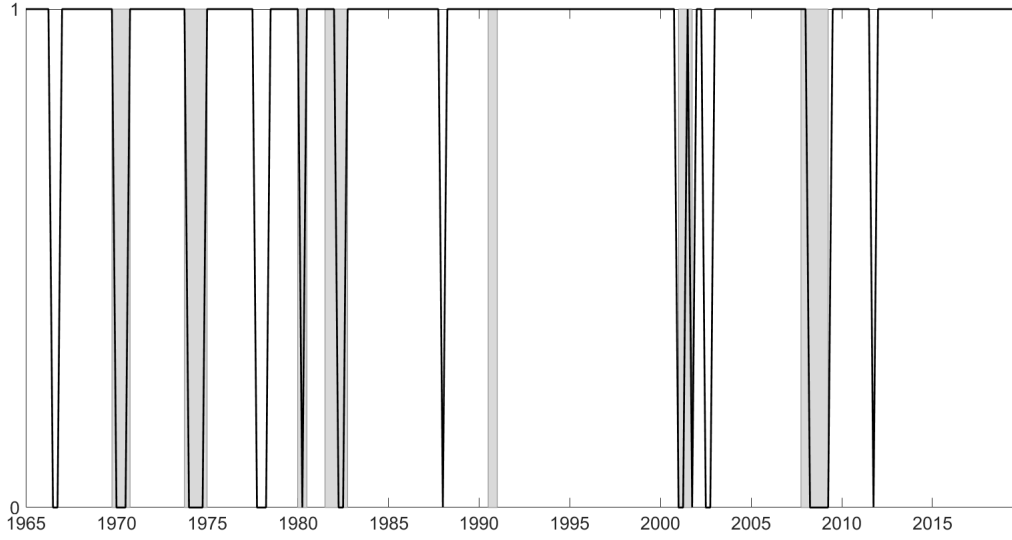


Figure 8: Coincident index and NBER-dated recessions (grey bars)

The coincident index is computed from six stock market technical indicators. A value of 0 indicates a recessions. See appendix 2 for a description of the technical indicators.

Appendix 1. List of predictors

- Log dividend-price ratio (DP): difference between the log of dividends (12-month moving sums of dividends paid on S&P 500) and the log of prices (S&P 500 index).
- Log dividend yield (DY): difference between the log of dividends (12-month moving sums of dividends paid on S&P 500) and the log of lagged prices (S&P 500 index).
- Log earnings-price ratio (EP): difference between the log of earnings (12-month moving sums of earnings on S&P 500) and the log of prices (S&P 500 index price).
- Log dividend-payout ratio (DE): difference between the log of dividends (12-month moving sums of dividends paid on S&P 500) and the log of earnings (12-month moving sums of earnings on S&P 500).
- Stock variance (SVAR): sum of squared daily returns on the S&P 500.
- Book-to-market ratio (BM): ratio of book value to market value for the Dow Jones Industrial Average.
- Net equity expansion (NTIS): ratio of 12-month moving sums of net equity issues by NYSE-listed stocks to the total end-of-year NYSE market capitalization.
- Treasury bill rate (TBL): three-month Treasury bill rate.
- Long-term yield (LTY): long-term government bond yield.
- Long-term return (LTR): long-term government bond return.
- Term spread (TMS): difference between the long-term government bond yield and the T-bill.
- Default yield spread (DFY): difference between Moody's BAA- and AAA-rated corporate bond yields.

- Default return spread (DFR): difference between long-term corporate bond and long-term government bond returns.
- Inflation rate (INFL): calculated from the Consumer Price Index (CPI) for all urban consumers.
- Investment to capital ratio (IK): ratio of aggregate (private nonresidential fixed) investment to aggregate capital for the whole economy.

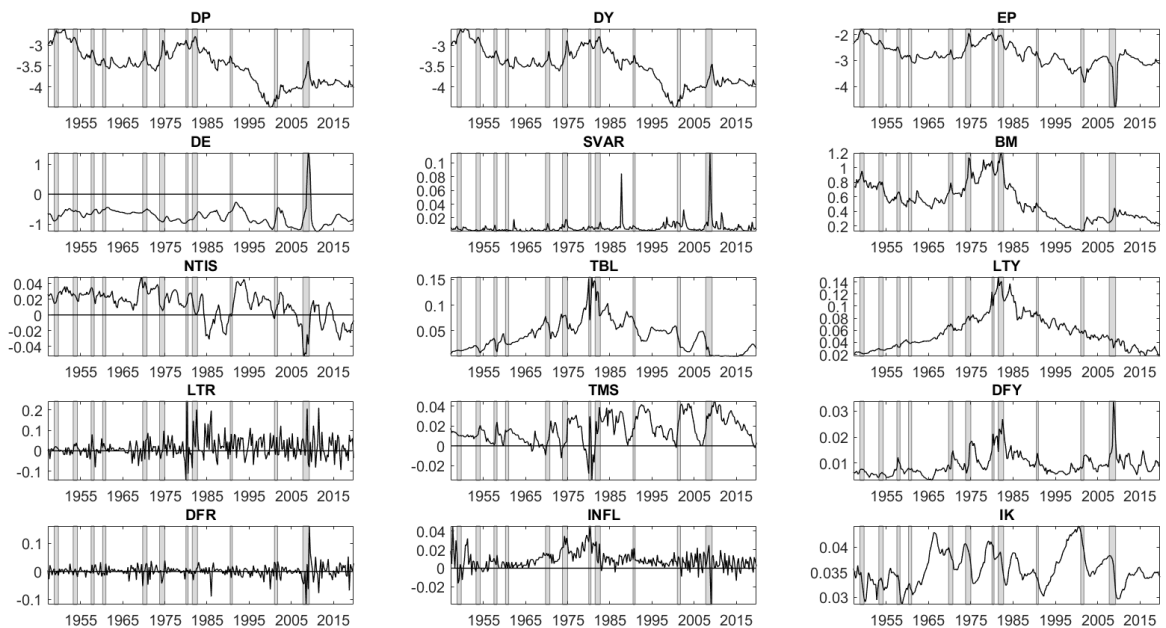


Figure 9: Time series of the predictors, U.S. data, 1947:Q1-2019:Q4

Grey bars depict NBER-dated recessions.

Appendix 2. List of technical indicators

Let P_t be the stock price index in quarter t .

Moving average indicator. The MA rule generates a buy ($S_{i,t} = 1$) or sell ($S_{i,t} = 0$) signal at the end of quarter t by comparing two moving averages:

$$S_{i,t} = \begin{cases} 1 & \text{if } MA_{short,t} \geq MA_{long,t} \\ 0 & \text{if } MA_{short,t} < MA_{long,t} \end{cases}$$

where

$$MA_{j,t} = \frac{1}{j} \sum_{i=0}^{j-1} P_{t-i} \text{ for } j = short, long$$

and *short* (*long*) is the length of the short (long) MA (*short* < *long*). The MA indicator with MA lengths *short* and *long* is denoted as MA(*short*,*long*). Intuitively, the MA rule detects changes in stock price trends because the short MA will be more sensitive to recent price movement than will the long MA. In the paper we use MA indicators with *short*=1 and *long*=3,4.¹³

Momentum indicator. The momentum rule generates the following buy ($S_{i,t} = 1$) or sell ($S_{i,t} = 0$) signal at the end of quarter t :

$$S_{i,t} = \begin{cases} 1 & \text{if } P_t \geq P_{t-m} , \\ 0 & \text{if } P_t < P_{t-m} . \end{cases}$$

Intuitively, a current stock price that is higher than its level m periods ago indicates positive momentum and relatively high expected excess returns, thereby generating a buy signal. The momentum indicator that compares P_t to P_{t-m} is denoted by MOM(m) and we compute momentum indicators for $m=3,4$.

¹³ Note that Neely et al. (2014) use monthly data, while we are using quarterly data.

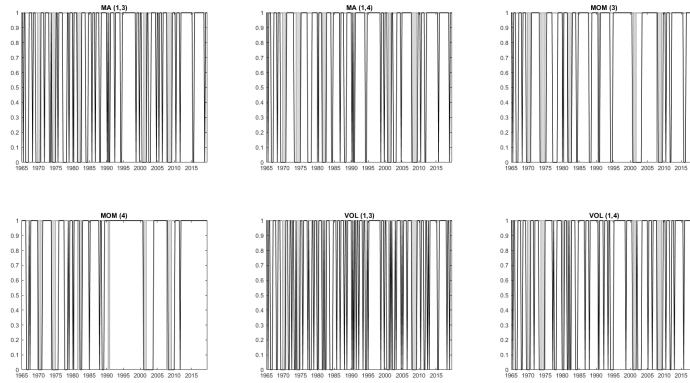


Figure 10: Time series of the technical indicators, U.S. data, 1965:Q1-2019:Q4
 Grey bars depict NBER-dated recessions.

Volume indicator. The idea of the volume indicator is to use volume data in conjunction with past prices to identify market trends. Let VOL_k be the trading volume during period k and define $OBV_t = \sum_{k=1}^t VOL_k D_k$ where $D_k = 1$ if $P_k - P_{k-1} \geq 0$ and -1 otherwise. The VOL rule generates the following buy ($S_{i,t} = 1$) or sell ($S_{i,t} = 0$) signal at the end of quarter t :

$$S_{i,t} = \begin{cases} 1 & \text{if } MA_{short,t}^{OBV} \geq MA_{long,t}^{OBV} \\ 0 & \text{if } MA_{short,t}^{OBV} < MA_{long,t}^{OBV} \end{cases}$$

where

$$MA_{j,t}^{OBV} = \frac{1}{j} \sum_{i=0}^{j-1} OBV_{t-i} \text{ for } j = short, long$$

Intuitively, relatively high recent volume together with recent price increases, say, indicate a strong positive market trend and generate a buy signal. The VOL indicator is denoted as $VOL(short, long)$ and we compute it for $short=1$ and $long=3,4$.

These TIs are plotted in figure 10.

Bank of Finland Research Discussion Papers 2023

ISSN 1456-6184, online

1/2023 Gonçalo Faria – Fabio Verona
Forecast combination in the frequency domain
ISBN 978-952-323-430-7, online