

Frequency-domain information for active portfolio management*

Gonçalo Faria[†] Fabio Verona[‡]

Abstract

We assess the benefits of using frequency-domain information for active portfolio management. To do so, we forecast the bond risk premium and equity risk premium using a methodology that isolates frequencies (of the predictors) with the highest predictive power. The resulting forecasts are more accurate than those of traditional forecasting methods for both asset classes. When used in the context of active portfolio management, the forecasts based on frequency-domain information lead to better portfolio performances than when using the original time series of the predictors. It produces higher information ratio (0.57 vs 0.45), higher CER gains (1.12% vs 0.81%), and lower maximum drawdown (19.1% vs 19.6%).

Keywords: equity risk premium, bond risk premium, predictability, multiresolution analysis, active portfolio management

JEL classification: C58, G11, G17

* We are grateful for the excellent comments of participants at the Católica Porto Business School seminar and the 2019 SAET conference. 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 Ciéncia e Tecnologia through project UID/GES/00731/2019.

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

‡ Bank of Finland – Monetary Policy and Research Department, and University of Porto – CEF.UP (fabio.verona@bof.fi)

1 Introduction

Active portfolio management relies on good return forecasts of asset classes under management. It is then of key interest for active asset managers to identify reliable predictors and good forecasting methods.

There is an extensive literature on the out-of-sample predictability of the equity risk premium (see the reviews of Rapach and Zhou, 2013 and Harvey, Liu, and Zhu, 2016), but the literature on the predictability of the bond risk premium is limited (notable contributions include Ludvigson and Ng, 2009; Thornton and Valente, 2012; Sarno, Schneider, and Wagner, 2016; and Gargano, Pettenuzzo, and Timmermann, 2019). The literature is dominated by time series analysis. Frequency domain techniques, like Fourier transformations, are rather new tools in finance applications (e.g. Dew-Becker and Giglio, 2016). In the context of forecasting equity returns, Faria and Verona (2018) and Bandi, Perron, Tamoni, and Tebaldi (2019) introduce models where equity returns and predictors are linear aggregates of components operating over different frequencies and predictability is frequency-specific.

The first contribution of this paper is to compare the performance of alternative predictive models of the bond risk premium (*BRP*) and the equity risk premium (*ERP*). We first use frequency-domain filtering techniques to expand an initial dataset of predictors to obtain more predictors for *BRP* and *ERP* forecasting. In particular, from each original variable we extract several time series, each corresponding to a particular frequency of the original variable and each representing a new predictor. The enlarged dataset has the same amount of information as the original dataset (we start from the same number of variables), but allows forecasting the *BRP* and *ERP* with more granular information. This allows us i) to tease out those predictor frequencies with the highest predictive power from others that bring noise to the exercise, and ii) to infer the relevance of using the frequency-domain information of the original predictors.

For both the *BRP* and the *ERP*, we find that the use of frequency-domain information significantly improves the statistical performance of forecasts over forecasts that only use original variables. While this result is not new with regards to *ERP* forecasting (see Faria and Verona 2017, 2019), it is, to the best of our knowledge, the first time frequency-domain information has been used in *BRP* forecasting. Furthermore, we find that the forecasting gains from using frequency-domain information are significantly higher when combining different frequencies from different original predictors than when combining different frequencies from the same original predictor. This finding suggests that different frequencies of different variables are useful predictors of equity and bond returns as they track different frequency components of the equity and bond risk premium. This result is in line with Fama and French (1989), who find that different financial variables track different frequency components of the equity premium.

The second contribution of this paper is an evaluation of the economic significance of frequency-domain information for active portfolio management. We adopt the perspective of a power-utility maximizing investor, whereby the *BRP* and *ERP* forecasts from the first step are treated as the investor's active views on stock and bond markets. We consider a mean-variance optimization framework and, as benchmark, a conventional allocation of 60% to stocks and 40% to bonds. We find that using frequency-domain information leads to better portfolio performances than when using the original time series of the predictors. It produces higher information ratio (0.57 vs 0.45), higher CER gains (1.12% vs 0.81%), and lower maximum drawdown (19.1% vs 19.6%). This finding is robust towards the consideration of an alternative portfolio optimization setting (Black-Litterman-type model), alternative benchmarks, and various portfolio constraint settings.

The rest of the paper is organized as follows. Section 2 sets out the data and methodology. Section 3 presents the out-of-sample results and performance of the proposed active portfolio management strategy. Section 4 documents the robustness test results. Section 5 concludes.

2 Data and methodology

Our focus is on out-of-sample (OOS) predictability of bond and equity risk premiums. The OOS exercise is relevant in evaluating effective return predictability in real time, while avoiding in-sample over-fitting, distortions from small sample size, and look-ahead bias.

Our monthly data extend from January 1973 to December 2018. *BRP* and *ERP* of month t are measured as the difference between the return on the 10-year US Treasury bond and the return on the S&P500 index in month t , respectively, and the one-month T-bill known at the beginning of month t (lagged-risk free rate). We use twelve variables taken from Goyal and Welch (2008) as the predictors: log dividend-price ratio (DP), log dividend yield (DY), log earnings-price ratio (EP), excess stock return volatility (RVOL), book-to-market ratio (BM), net equity expansion (NTIS), long-term bond yield (LTY), long-term bond return (LTR), term spread (TMS), default yield spread (DFY), default return spread (DFR), and lagged inflation rate (INFL). These predictors are briefly described in Appendix 1. Table 1 reports the summary statistics for *BRP*, *ERP* and the predictors. Figure 1 provides their time series.

The first step of our forecasting methodology is based on a wavelet multiresolution analysis as described in sub-section 2.1. The OOS procedure is explained in sub-section 2.2. The asset allocation framework is covered in sub-section 2.3.

2.1 Wavelet multiresolution analysis

Wavelet multiresolution analysis (MRA) allows decomposition of any time series into its frequency components in a way similar to bandpass filtering (e.g. Baxter and King, 1999). Given a time series y_t , its wavelet multiresolution representation can be written as

$$y_t = \sum_{j=1}^J y_t^{D_j} + y_t^{S_J} , \quad (1)$$

where $y_t^{D_j}$ are the J wavelet detail components and $y_t^{S_J}$ is the wavelet smooth component. Equation (1) shows that the original series y_t can be decomposed in several time series components, each capturing the fluctuation of the original time series within a specific frequency band. For small j , the j wavelet detail components represent the higher frequency components of the time series (the short-term dynamics). As j increases, the j wavelet detail components represent lower frequencies fluctuations of the series. Finally, the smooth component captures the lowest frequency component (series trend).

Here, we perform our wavelet decomposition analysis using the Haar wavelet filter¹ and the maximal overlap discrete wavelet transform (MODWT) MRA. This methodology is not restricted to a particular sample size and is not sensitive to the choice of starting point for the examined time series. Moreover, it does not introduce phase shifts in the wavelet coefficients, i.e. peaks and troughs of the original time series are perfectly aligned with similar events in the MODWT MRA.²

Given the length of the data series under analysis, we consider a $J=6$ level MRA for each of the original predictors, so that the decomposition delivers seven time-frequency series: six wavelet detail components ($y_t^{D_1}$ to $y_t^{D_6}$) and a wavelet smooth component ($y_t^{S_6}$).³ As we use monthly data, the first detail component $y_t^{D_1}$ captures oscillations between 2 and 4 months, while detail components $y_t^{D_2}$, $y_t^{D_3}$, $y_t^{D_4}$, $y_t^{D_5}$ and $y_t^{D_6}$ capture oscillations with a period of 4-8, 8-16, 16-32, 32-64 and 64-128 months, respectively. Finally, the smooth component $y_t^{S_6}$

¹ Besides its simplicity and wide use, the Haar filter makes a neat connection to temporal aggregation as the wavelet coefficients are simply differences of moving averages (see Bandi, Perron, Tamoni, and Tebaldi, 2019 and Lubik, Matthes, and Verona, 2019).

² This section provides a brief description of the theory directly relevant to our empirical analysis. A more detailed analysis of wavelet methods is provided in Appendix 2 and in Percival and Walden (2000). Recent papers using the MODWT MRA decomposition are Bekiros and Marcellino (2013), Gallegati and Ramsey (2013), Barunik and Vacha (2015), Crowley and Hughes Hallett (2015), Berger (2016), and Faria and Verona (2018), among others. See Verona (2019) for a description of the advantages of wavelet filters over other band-pass filtering techniques.

³ As regards the choice of J , the number of observations dictates the maximum number of frequency bands that can be used. In particular, if t_0 is the number of observations in the in-sample period, then J has to satisfy the constraint $J \leq \log_2 t_0$.

(re-denoted as $y_t^{D_7}$ in our later discussion) captures oscillations with a period longer than 128 months (10.6 years).⁴

To illustrate the rich set of dynamics aggregated (and therefore hidden) in the original time series, Figure 2 plots the time series of one of the predictors used (term spread) and its seven time-frequency series components. As expected, the lower the frequency, the smoother the resulting filtered time series. As can be seen, the time-frequency series components exhibit different time series properties and dynamics, so one can expect that only some are good *ERP* and *BRP* predictors. As Faria and Verona (2019) show, the lowest frequency component of the term spread (TMS^{D_7}) is a strong OOS predictor of the *ERP*, whereas the other frequency components of the term spread have much worse forecasting performances.

2.2 Out-of-sample forecasts

The OOS forecasts of the *BRP* and *ERP* are generated using a sequence of expanding windows. We use an initial sample (1973:01 to 1989:12) to make our first one-step-ahead OOS forecast. The sample is then increased by one observation and a new one-step-ahead OOS forecast is produced. We proceed this way until the end of the sample, ultimately obtaining a sequence of 348 one-step-ahead OOS forecasts. The full OOS period spans the period from 1990:01 to 2018:12.

As the MODWT MRA is a two-sided filter, we recompute the frequency components of the original predictors recursively at each iteration of the OOS forecasting process using data from the start of the sample through the month of forecast formation. This important step ensures that our method does not have a “look-ahead” bias, as the forecasts are made with current and past information only. The literature suggests several types of boundary

⁴ In the MODWT, each wavelet filter at frequency j approximates an ideal high-pass filter with passband $f \in [1/2^{j+1}, 1/2^j]$, while the smooth component is associated with frequencies $f \in [0, 1/2^{j+1}]$. The level j wavelet components are therefore associated to fluctuations with periodicity $[2^j, 2^{j+1}]$ (months, in our case).

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 extended symmetrically at the right boundary to twice the time series length before computing the MODWT MRA.

2.2.1 Predictive regression model

Let \mathbf{X} be a vector of predictors. The *ERP* predictive regression model is

$$ERP_{t+1} = \alpha + \boldsymbol{\beta} \mathbf{X}_t + \varepsilon_{t+1} , \quad (2)$$

and the one-step-ahead OOS forecast of the *ERP*, \widehat{ERP}_{t+1} , is given by:

$$\widehat{ERP}_{t+1} = \hat{\alpha}_t + \hat{\boldsymbol{\beta}}_t \mathbf{X}_t , \quad (3)$$

where $\hat{\alpha}$ and $\hat{\boldsymbol{\beta}}$ are the OLS estimates of parameter α and vector of parameters $\boldsymbol{\beta}$, respectively.

The same predictive regression model is used to forecast the *BRP*.

2.2.2 Predictors used

We consider four cases when running model (2)-(3):

- \mathbf{X} includes one original predictor, i.e. we run bi-variate regressions using one original predictor at a time. We denote this model as *single_ts*.
- \mathbf{X} includes all original predictors, i.e. we run multi-variate regressions using several original predictors. We denote this model as *multi_ts*.
- \mathbf{X} includes the frequencies (obtained with the MODWT MRA) of one original predictor,

i.e. we run multi-variate regressions using different frequencies of one original predictor at a time. This model is denoted as *single_wav*.

- \mathbf{X} includes the frequencies (obtained with the MODWT MRA) of the original predictors, i.e. we run multi-variate regressions using several frequencies of different original predictors. We denote this model as *multi_wav*.

Comparison of the *ts* and *wav* models shows the value of using more granular data from frequency decomposition of the original predictors. Comparison of the *single* and *multi* models helps identify the usefulness of information from different original predictors.

2.2.3 Forecast evaluation

The forecasting performance of the predictive models are evaluated using the Campbell and Thompson (2008) R_{OS}^2 statistic. As is standard in the literature, the benchmark model is the prevailing mean forecast \bar{r}_t , i.e. the average *ERP* or *BRP* up to time t . The R_{OS}^2 statistic measures the proportional reduction in the mean squared forecast error for the predictive model ($MSFE_{PRED}$) relative to the historical mean ($MSFE_{HM}$) and is given by

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

where \hat{r}_{t+1} is the *ERP* (*BRP*) forecast for $t+1$ from the predictive model under analysis, and r_{t+1} is the realized *ERP* (*BRP*) from t to $t+1$. A positive (negative) R_{OS}^2 indicates that the predictive model outperforms (underperforms) the historical mean (HM) in terms of MSFE.

The statistical significance of the results is evaluated using the Clark and West (2007) statistic, which tests the null hypothesis that the MSFE of the HM model is less than or equal to the MSFE of the predictive model under analysis against an alternative hypothesis that the

MSFE of the HM model is greater than the MSFE of the predictive model under analysis ($H_0 : R_{OS}^2 \leq 0$ against $H_A : R_{OS}^2 > 0$).

2.3 Asset allocation

The ultimate objective of our analysis is to evaluate the economic significance of frequency-domain information for active portfolio management (*APM*). The portfolio optimization framework used in this paper is described in sub-section 2.3.1. The performance measurements of the proposed active strategy are described in section 2.3.2.

2.3.1 The portfolio optimization framework

As it is standard in the literature, we adopt the perspective of a mean-variance investor, who invests in bonds and equities. The corresponding portfolio weights are ϖ_b and ϖ_e , respectively, represented in the vector $\varpi = (\varpi_b, \varpi_e)$. Initial wealth is normalized to 1. The rebalancing decisions that underlie the *APM* are assumed to be made on a monthly basis, making use of the forecasts of bond and equity returns for the next month. The objective of the portfolio optimization framework is to optimize the trade-off between risk and return. The optimization problem is

$$\min_{\varpi} \left[\gamma \Theta_P(\varpi) - \varpi' \hat{\mathbf{R}} \right], \quad (4)$$

where γ is the relative risk aversion coefficient (which we assume to be equal to 2), $\hat{\mathbf{R}} = (\hat{R}_{b,t+1}, \hat{R}_{e,t+1})$ is the vector of one-step-ahead return forecasts of bonds ($\hat{R}_{b,t+1}$) and equities ($\hat{R}_{e,t+1}$), and $\Theta_P(\varpi)$ is the portfolio risk function.

The one-step-ahead bond return forecast ($\hat{R}_{b,t+1}$) corresponds to the one-step-ahead forecast of the bond risk premium ($B\hat{R}P_{t+1}$) minus the risk-free rate (which is known at the beginning

of the period). The same procedure applies to the one-step-ahead equity return forecast ($\hat{R}_{e,t+1}$). In the context of the mean-variance optimization framework, the portfolio risk function $\Theta_P(\boldsymbol{\varpi})$ is set as $\Theta_P(\boldsymbol{\varpi}) = \sqrt{\boldsymbol{\varpi}' \hat{\Sigma} \boldsymbol{\varpi}}$, where $\hat{\Sigma}$ is the estimated monthly returns covariance matrix. We estimate Σ using an exponentially weighted moving average approach, setting the the decay parameter to 0.97.

To place realistic limits on the possibilities of leveraging the *APM* portfolio, we introduce some constraints on the weight vector $\boldsymbol{\varpi}$. The first constraint sets an upper bound to the sum of the portfolio weights, $\boldsymbol{\varpi}' \mathbf{I}_2 = h$, where \mathbf{I}_2 is a 2-vector of ones and h denotes the maximum leverage. The second constraint sets a lower bound l to the weight of each asset, $w_i \geq l$, with $i = b, e$ (b for bond and e for equity). We set $h = 1.5$, which means that the investor cannot borrow more than 50% of total wealth, and $l = 0$, which excludes short sales.

The *APM* portfolio return at $t+1$, $R_{p,t+1}$, is then given by:

$$R_{p,t+1} = \hat{\boldsymbol{\varpi}}' \mathbf{R}_{t+1} + (1 - \hat{\boldsymbol{\varpi}}' \mathbf{I}_2) r_f ,$$

where \mathbf{R} is the vector of realized returns of bonds (R_b) and equities (R_e) and r_f is the one-month risk-free rate. Note that if $h = 1$, the portfolio return is $R_{p,t+1} = \hat{\boldsymbol{\varpi}}' \mathbf{R}_{t+1}$.

2.3.2 Measuring the performance of the active strategy

We consider the conventional allocation of 60% to stocks and 40% to bonds as the benchmark portfolio, using six performance measures: Sharpe ratio, composite annual growth rate of returns (CAGR), tracking error, information ratio, maximum drawdown, and certainty equivalent return (CER) gain.

The reported Sharpe ratio is the one-year moving average of the portfolio's annualized Sharpe ratio. In the context of the mean-variance portfolio optimization framework, the Sharpe ratio is the traditionally reported performance metric. The tracking error is measured as the

annualized standard deviation of the *APM* monthly excess return towards the benchmark. The information ratio is measured as the annualized average *APM* monthly excess return (towards the benchmark) divided by the tracking error. Both the information ratio and the tracking error are relevant performance metrics for actively managed portfolios, as they directly inform about the merits of deviating actively from the benchmark. The maximum drawdown measures the downside risk of the strategy under analysis and gives the maximum percentage reduction in the portfolio’s cumulative return.

Power utility is given by $U(x) = \frac{x^{1-\gamma}}{1-\gamma}$, where $x = 1 + R_p$ and R_p is the portfolio return. Let \bar{U}_j , $j = \text{APM, benchmark}$ denote the average utility of an investor with access to the *APM* and benchmark portfolios, respectively. The CER is given by $CER_j = [(1 - \gamma) \bar{U}_j]^{1/(1-\gamma)} - 1$, $j = \text{APM, benchmark}$. We report the annualized utility gain, computed as $12 \cdot (CER_{\text{APM}} - CER_{\text{benchmark}})$. This can be interpreted as the annual portfolio management fee that an investor would be willing to pay for access to the *APM* portfolio instead of the benchmark portfolio.

3 Results

3.1 Out-of-sample forecasting statistical performance

As described in sub-section 2.2.2, we run four predictive models: (i) regressions using one original predictor at a time (*single_ts*); (ii) regressions using several original predictors (*multi_ts*); (iii) regressions using different frequencies from one original predictor at a time (*single_wav*); and (iv) regressions using different frequencies from different original predictors (*multi_wav*). For clarity, we only report the result of the best specification for each model (i)-(iv), i.e. the model specification that maximizes the R_{OS}^2 statistic. Results are

reported in Table 2.⁵ We highlight three main results.

First, regardless of the forecasting model considered, predictability of the *BRP* is higher than that of the *ERP*.

Second, there are common patterns across *BRP* and *ERP* forecasts. When using the information from one original predictor only (*single_ts* versus *single_wav*), there are forecasting gains from using frequency-domain information. The maximum R_{OS}^2 using the original time series of the predictors is 1.70% for the *BRP*, while it is negative for the *ERP*. When using frequency-decomposed predictors, the maximum R_{OS}^2 increases to 5.45% for the *BRP*, and is positive and statistically significant (1.77%) for the *ERP*. Likewise, there are forecasting gains when combining information from different original predictors (*single* versus *multi*), except when forecasting the *ERP* with the time series (*single_ts* versus *multi_ts*). In all other cases, there is an increase in the maximum R_{OS}^2 .

Third, when comparing the *single_wav* model with the *multi_wav* model, there are noticeable forecast improvements by using different frequencies from different original predictors (*multi_wav*) instead of using different frequencies of one original predictor (*single_wav*). The best R_{OS}^2 for the *BRP* forecast is 7.20%, while the best R_{OS}^2 for the *ERP* forecast is 3.97%.

These results indicate that using frequency-domain information helps make better forecasts of bond and equity risk premiums. Next, we analyze if these statistical gains translate to better portfolio performances.

⁵ Appendix 3 presents the results for the *single_ts* and the *single_wav* model for all original predictors. For computational reasons, we consider at most three frequencies from all possible predictors in model (iv).

3.2 Active portfolio management performance

We use the *BRP* and *ERP* forecasts from the respective *multi_wav* model to feed the vector of active views $\hat{\mathbf{R}} = (\hat{R}_{b,t+1}, \hat{R}_{e,t+1})$ driving the *APM* strategy. This is denoted as *APM_WAV*. As mentioned, the benchmark is the conventional allocation of 60% to stocks and 40% to bonds (denoted *Benchmark₆₀₋₄₀*). For comparison purposes, we also report the performance of an *APM* strategy based on the *BRP* and *ERP* forecasts obtained with the original time series of the predictors (*multi_ts*). We denote this as *APM_TS*.

Figure 3 presents the *APM_WAV*, *APM_TS*, and *Benchmark₆₀₋₄₀* portfolio weights (solid, dashed, and dotted lines, respectively). Both active strategies (*APM_WAV* and *APM_TS*) strongly deviate from the 60-40 benchmark throughout the entire sample period. Moreover, the *APM_WAV* weights seem to oscillate around the trends defined by the *APM_TS* weights. Interestingly, with the exception of the mid-nineties period, the differences between *APM_WAV* and *APM_TS* weights are most evident around and during recessions. In particular, the *APM_WAV* has relatively lower exposure to equity immediately before and during recessions. This suggests an improved equity market timing of the *APM_WAV* strategy compared to that of the *APM_TS* strategy.

In Panel A of Table 3, we report the performance measurements of the strategies. Both *APM* strategies outperform the *Benchmark₆₀₋₄₀*, with the *APM_WAV* strategy outperforming the *APM_TS*. Compared with the *Benchmark₆₀₋₄₀* performance, both *APM* strategies improve the average annual return while decreasing the maximum drawdown. This translates to higher annualized Sharpe ratios. The fact that the active deviations from the 60-40 benchmark (as illustrated in Figure 3) add value to the active investor is reflected in the annualized information ratios of 0.57 (*APM_WAV*) and 0.45 (*APM_TS*). From an utility perspective, this also translates to annualized CER gains of 1.12% (*APM_WAV*) and 0.81% (*APM_TS*).

The fact that the *APM_WAV* strategy outperforms the *APM_TS* strategy implies that

there are economic gains from using frequency-domain information in active portfolio management. To disentangle the distribution of those gains across the asset classes traded (bonds and equity), we report the performance metrics for two additional active portfolio management strategies in Panel B of Table 3. APM_Equity_WAV is based on the forecast of the $multi_wav$ ($multi_ts$) model for equity (bond) return. APM_Bond_WAV is based on the forecast of the $multi_wav$ ($multi_ts$) model for bond (equity) return. By comparing APM_Equity_WAV (APM_Bond_WAV) with APM_TS , we can assess the gains from using frequency-domain information in the forecast of equity (bond) return.

We highlight two main results. First, there are gains for both asset classes when using frequency-domain information in the forecast of their returns. The gains are quite similar in magnitude. Second, the gains are more expressive when using frequency-domain information to forecast both the return of bonds and equities.

Figure 4 shows the cumulative wealth of an investor who invests 1\$ in January 1990 and reinvests all proceeds adopting the APM_WAV strategy (solid line), the APM_TS strategy (dashed line), and the $Benchmark_{60-40}$ strategy (dotted line). From a cumulative return perspective, the active strategy APM_WAV clearly outperforms the others. By December 2018, the investor has obtained \$38.6 with the APM_WAV strategy, instead of \$28.7 with the APM_TS strategy, or \$12.4 with the $Benchmark_{60-40}$.

The strong performance of the APM_WAV strategy is not without its caveats. In the upper panel of Figure 5, we report the dynamics of the 3-year moving average information ratio of the APM_WAV strategy (solid line). The 3-year moving average information ratio is positive for most of the sample period, but there are periods when it is negative (i.e. generating utility losses). However, the figure also shows that the APM_WAV strategy dominates the APM_TS strategy (dotted line), as its 3-year moving average information ratio is either higher (for most of the sample) or similar. From the utility perspective, similar conclusions can be drawn by looking at the dynamics of the 3-year moving average CER gains of the

APM_WAV and *APM_TS* strategies (reported in the lower panel of Figure 5).

Overall, these results demonstrate the usefulness of frequency-domain information for active portfolio management. In the next section, we test the robustness of our findings by considering an alternative portfolio-optimization framework and other changes in the settings used so far.

4 Robustness

4.1 Alternative portfolio optimization framework

We test the robustness of the results reported so far by using the Black-Litterman model (*BLM*), which is a framework often considered in the context of *APM*. The objective of the *BLM* is to outperform the benchmark portfolio within a certain tracking error.

We use the same *BRP* and *ERP* forecasts from previous sections as the active views on stock and bond markets, treat them as inputs in a version of the *BLM* (as proposed by Da Silva, Lee, and Pornrojnangkool, 2009 and Almadi, Rapach, and Suri, 2014 and described in Appendix 4) to obtain optimal weights across assets.

We consider a power-utility maximizing investor with $\gamma = 2$ and *Benchmark*_{60–40} as the benchmark strategy. For simplicity, we assume the investor will neither leverage nor short-sell available assets ($h = 1$ and $l = 0$). The target level of the annualized tracking error of the investor is assumed to be 5.80%, i.e. the same tracking error of the *APM_WAV* strategy for an investor with $\gamma = 2$, $h = 1$ and $l = 0$. *APM_BLMWAV* and *APM_BLMTS* denote the active portfolio management strategies based on asset return forecasts from *multi_wav* and *multi_ts* methodologies used in the context of a Black-Litterman optimization framework.

The results, which are reported in Panel C of Table 3, are qualitatively similar to those in the mean-variance setting. Both *APM* strategies based on the BLM outperform the

$Benchmark_{60-40}$, achieving positive information ratios and CER gains. Similarly, using frequency-domain information in the context of the BLM still improves the performance of the strategy over the scenario where only the original time series of the predictors are used (APM_BLMWAV versus APM_BLMTS). Both the annualized information ratio and the annualized CER gain are higher (0.44 versus 0.11 and 0.68% versus 0.12%, respectively). Finally, the portfolio weights (reported in Figure 6) of the strategy using frequency-domain information are much more stable than those of the strategy using time-series information only.

4.2 Other robustness tests

In this sub-section, we briefly comment on additional robustness tests that were implemented.⁶

4.2.1 Alternative benchmarks

Instead of $Benchmark_{60-40}$, we consider two alternative benchmarks: a naive diversification rule 1/N (50% equity and 50% bonds) and an allocation of 40% equity and 60% bonds. In both cases, the information ratios and CER gains of the APM_WAV and APM_TS strategies are still positive, and the information ratios and CER gains of the APM_WAV strategy are higher than those of APM_TS strategy. Qualitatively, these results confirm that our findings are robust towards alternative benchmarks.

4.2.2 Alternative set of portfolio constraints and investor risk aversion

For a given level of risk aversion of the representative investor, the APM_WAV strategy outperforms the APM_TS strategy (and the $Benchmark_{60-40}$) in alternative scenarios with

⁶ The results are not reported here, but available upon request from the authors.

(i) no leverage or short-selling possibilities ($h = 1$ and $l = 0$), (ii) no leverage possibilities, but short-selling allowed ($h = 1$ and $l = -0.5$) and (iii) both leverage and short-selling possibilities ($h = 1.5$ and $l = -0.5$). The higher is the level of leverage and short-selling allowed, the higher is the outperformance of the *APM_WAV* strategy versus the *APM_TS* strategy. Finally, the lower is the level of risk aversion of the representative investor, the higher is the outperformance of the *APM_WAV* strategy versus the *APM_TS* strategy (everything else constant).

5 Concluding remarks

Fama and French (1989) find that different financial variables can be useful in predicting equity returns as they track different frequency components of the equity premium. In this paper, we show that using information from different frequencies of different predictors helps improve forecasts of bond and equity returns. When used in the context of active portfolio management, these forecasts lead to superior portfolio performances.

We envision several interesting research avenues related with the use of frequency-domain information for active portfolio management. Here, we only used twelve variables as possible predictors of bond and equity returns, but the same methodology can be readily applied to larger datasets, and even combined with large dimensional statistical models. It could also be worthwhile to explore the statistical and economic gains from the use of frequency-domain information in the context of forecasting models with time-varying parameters and stochastic volatility.

References

ALMADI, H., D. E. RAPACH, AND A. SURI (2014): “Return Predictability and Dynamic Asset Allocation: How Often Should Investors Rebalance?,” *The Journal of Portfolio Management*, 40(4), 16–27.

BANDI, F., B. PERRON, A. TAMONI, AND C. TEBALDI (2019): “The Scale of Predictability,” *Journal of Econometrics*, 208(1), 120–140.

BARUNIK, J., AND L. VACHA (2015): “Realized wavelet-based estimation of integrated variance and jumps in the presence of noise,” *Quantitative Finance*, 15(8), 1347–1364.

BAXTER, M., AND R. KING (1999): “Measuring Business Cycles: Approximate Band-Pass Filters For Economic Time Series,” *Review of Economics and Statistics*, 81(4), 575–593.

BEKIROS, S., AND M. MARCELLINO (2013): “The multiscale causal dynamics of foreign exchange markets,” *Journal of International Money and Finance*, 33(C), 282–305.

BERGER, T. (2016): “Forecasting Based on Decomposed Financial Return Series: A Wavelet Analysis,” *Journal of Forecasting*, 35(5), 419–433.

BLACK, F., AND R. LITTERMAN (1992): “Global Portfolio Optimization,” *Financial Analysts Journal*, 48(5), 28–43.

CAMPBELL, J. Y., AND S. B. THOMPSON (2008): “Predicting Excess Stock Returns Out of Sample: Can Anything Beat the Historical Average?,” *Review of Financial Studies*, 21(4), 1509–1531.

CLARK, T., AND K. WEST (2007): “Approximately normal tests for equal predictive accuracy in nested models,” *Journal of Econometrics*, 138(1), 291 – 311.

CROWLEY, P. M., AND A. HUGHES HALLETT (2015): "Great moderation or "Will o' the Wisp"? A time-frequency decomposition of GDP for the US and UK," *Journal of Macroeconomics*, 44(C), 82–97.

DA SILVA, A., W. LEE, AND B. PORNROJNANGKOOL (2009): "The Black Litterman Model for Active Portfolio Management," *The Journal of Portfolio Management*, 35(2), 61–70.

DEW-BECKER, I., AND S. GIGLIO (2016): "Asset Pricing in the Frequency Domain: Theory and Empirics," *Review of Financial Studies*, 29(8), 2029–2068.

FAMA, E. F., AND K. R. FRENCH (1989): "Business conditions and expected returns on stocks and bonds," *Journal of Financial Economics*, 25(1), 23–49.

FARIA, G., AND F. VERONA (2017): "Out-of-sample time-frequency predictability of the equity risk premium," Discussion paper, mimeo, previously circulated as "Forecasting the Equity Risk Premium with Frequency-Decomposed Predictors", Bank of Finland Discussion paper 1/2017.

——— (2018): "Forecasting stock market returns by summing the frequency-decomposed parts," *Journal of Empirical Finance*, 45, 228 – 242.

——— (2019): "The yield curve and the stock market: Mind the long run," *Journal of Financial Markets*, forthcoming.

GALLEGATI, M., AND J. B. RAMSEY (2013): "Bond vs stock market's Q: Testing for stability across frequencies and over time," *Journal of Empirical Finance*, 24(C), 138–150.

GARGANO, A., D. PETTENUZZO, AND A. TIMMERMANN (2019): "Bond Return Predictability: Economic Value and Links to the Macroeconomy," *Management Science*, 65(2), 508–540.

GOYAL, A., AND I. WELCH (2008): “A Comprehensive Look at The Empirical Performance of Equity Premium Prediction,” *Review of Financial Studies*, 21(4), 1455–1508.

HARVEY, C. R., Y. LIU, AND H. ZHU (2016): “... and the Cross-Section of Expected Returns,” *Review of Financial Studies*, 29(1), 5–68.

IDZOREK, T. (2004): “A Step-by-Step Guide to the Black-Litterman Model,” Discussion paper, Zephyr Associates.

LUBIK, T. A., C. MATTHES, AND F. VERONA (2019): “Assessing U.S. aggregate fluctuations across time and frequencies,” Research Discussion Papers 5/2019, Bank of Finland.

LUDVIGSON, S. C., AND S. NG (2009): “Macro Factors in Bond Risk Premia,” *Review of Financial Studies*, 22(12), 5027–5067.

PERCIVAL, D., AND A. WALDEN (2000): *Wavelet methods for time series analysis*. Cambridge University Press.

RAPACH, D. E., AND G. ZHOU (2013): *Forecasting Stock Returns* vol. 2 of *Handbook of Economic Forecasting*, pp. 328–383. Elsevier.

SARNO, L., P. SCHNEIDER, AND C. WAGNER (2016): “The economic value of predicting bond risk premia,” *Journal of Empirical Finance*, 37, 247 – 267.

THORNTON, D. L., AND G. VALENTE (2012): “Out-of-Sample Predictions of Bond Excess Returns and Forward Rates: An Asset Allocation Perspective,” *The Review of Financial Studies*, 25(10), 3141–3168.

VERONA, F. (2019): “Investment, Tobin’s Q, and Cash Flow Across Time and Frequencies,” *Oxford Bulletin of Economics and Statistics*, forthcoming.

	mean	median	min	max	std. dev.	AR(1)
<i>BRP</i> (%)	0.30	0.36	-11.2	14.4	3.10	0.05
<i>ERP</i> (%)	0.51	0.85	-22.1	16.1	4.38	0.03
DP	-3.64	-3.65	-4.52	-2.75	0.43	0.99
DY	-3.64	-3.64	-4.53	-2.75	0.43	0.99
EP	-2.84	-2.88	-4.84	-1.90	0.48	0.99
RVOL (ann.)	0.14	0.14	0.05	0.32	0.05	0.96
BM	0.47	0.35	0.12	1.21	0.28	0.99
NTIS	0.01	0.01	-0.06	0.05	0.02	0.98
LTY (%, ann.)	6.77	6.76	1.75	14.8	2.91	0.99
LTR (%)	0.69	0.72	-11.2	15.2	3.10	0.05
TMS (%, ann.)	2.09	2.24	-3.65	4.55	1.46	0.95
DFY (%, ann.)	1.09	0.95	0.55	3.38	0.46	0.96
DFR (%)	0.01	0.05	-9.75	7.37	1.49	-0.04
INFL (%)	0.32	0.30	-1.92	1.81	0.38	0.61

Table 1: Summary statistics

This table reports summary statistics for the bond risk premium (*BRP*), equity risk premium (*ERP*), and the set of predictors. *BRP* and *ERP* are measured as the difference between the return on the 10-year US Treasury bond and the return on the S&P500 index, respectively, and the return on a one-month T-bill. *BRP*, *ERP*, *LTR*, *DFR*, and *INFL* (LTY, TMS, and DFY) are measured in percent (annual percent). The set of predictors is described in Appendix 1. The sample period runs from 1973:01 to 2018:12.

	<i>single_ts</i>		<i>multi_ts</i>	
	R_{OS}^2	Predictor	R_{OS}^2	Predictors
<i>BRP</i>	1.70**	TMS	3.40***	DP, DY, TMS
<i>ERP</i>	-0.29	LTR	-0.29	LTR
	<i>single_wav</i>		<i>multi_wav</i>	
	R_{OS}^2	Predictor (frequency)	R_{OS}^2	Predictors (frequency)
<i>BRP</i>	5.45***	BM (D_1, D_2, D_5, D_7)	7.20***	BM (D_2), NTIS (D_1), TMS (D_5)
<i>ERP</i>	1.77**	INFL (D_2, D_5)	3.97***	EP (D_3), RVOL (D_5), TMS (D_7)

Table 2: Out-of-sample R-squares (R_{OS}^2)

This table reports the maximum out-of-sample R-squares (in percentage) for the bond risk premium (*BRP*) and equity risk premium (*ERP*) forecasts at monthly frequencies of four predictive models: regressions using one original predictor at a time (*single_ts*); regressions using different original predictors (*multi_ts*); regressions using the frequencies from one original predictor at a time (*single_wav*); and regressions using frequencies from different original predictors (*multi_wav*). The predictor(s) and their frequency(ies) are reported. The out-of-sample R-squares (R_{OS}^2) measures the proportional reduction in the mean squared forecast error for the predictive model relative to the forecast based on the historical mean. The one-month-ahead out-of-sample forecast of the *BRP* and the *ERP* is generated using a sequence of expanding windows. The sample period runs from 1973:01 to 2018:12. The out-of-sample forecasting period extends from 1990:01 to 2018:12 (monthly frequency). Asterisks denote significance of the out-of-sample *MSFE*-adjusted statistic of Clark and West (2007). *** and ** denote significance at the 1% and 5% levels, respectively.

	Average return	CAGR	Sharpe ratio	Maximum drawdown	Tracking error	Information ratio	CER gain
Panel A: baseline							
<i>APM_WAV</i>	14.2%	13.4%	1.28	19.1%	7.4%	0.57	1.12%
<i>APM_TS</i>	13.0%	12.3%	1.18	19.6%	7.2%	0.45	0.81%
<i>Benchmark₆₀₋₄₀</i>	9.5%	9.1%	1.13	29.1%	-	-	-
Panel B: different forecasting inputs							
<i>APM_Equity_WAV</i>	13.7%	13.0%	1.24	19.3%	7.3%	0.52	0.99%
<i>APM_Bond_WAV</i>	13.5%	12.7%	1.22	19.4%	7.3%	0.49	0.94%
Panel C: different portfolio optimization framework							
<i>APM_BLMWAV</i>	12.4%	11.7%	1.23	24.2%	5.8%	0.44	0.68%
<i>APM_BLMTS</i>	10.2%	9.6%	1.05	26.1%	5.8%	0.11	0.12%

Table 3: Portfolio performance statistics

This table reports the performance statistics of different portfolio strategies. The performance statistics are: average return, which is the annualized first moment of returns time series; CAGR, which is the composite annual growth rate of returns time series; Sharpe ratio, measured as the 1-year moving average of portfolio's annualized Sharpe ratio; maximum drawdown, measured as the maximum percentage reduction in the portfolio's cumulative return; tracking error, measured as the annualized standard deviation of the *APM* monthly excess return (towards the benchmark); the information ratio, measured as the annualized average *APM* monthly excess return (towards the benchmark) divided by the tracking error; CER gain, measured as the annualized increase in certainty equivalent return that a power-utility maximizing investor with relative risk aversion $\gamma = 2$ would have by having access to the *APM* portfolio instead of the benchmark portfolio. The benchmark portfolio is 60% allocation to stocks and 40% to bonds. In Panel A are presented the performance statistics for the strategies *APM_WAV* and *APM_TS*, which are the active portfolio management strategy based on asset return forecasts from *multi_wav* and *multi_ts* methodologies, respectively. In Panel B are presented the performance statistics for the strategy *APM_Equity_WAV*, which is an active portfolio management strategy based on equity (bond) return forecasts from *multi_wav* (*multi_ts*) methodology, and for the strategy *APM_Bond_WAV*, which is an active portfolio management strategy based on bond (equity) return forecasts from *multi_wav* (*multi_ts*) methodology. In Panel C are presented the performance statistics for the strategies *APM_BLMWAV* and *APM_BLMTS*, which are the active portfolio management strategy based on asset return forecasts from *multi_wav* and *multi_ts* methodologies used in the context of a Black-Litterman portfolio-optimization framework, respectively. The sample period is from 1973:01 to 2018:12. The out-of-sample forecasting period is from 1990:01 to 2018:12, monthly frequency.




Figure 1: Monthly time series of the *BRP*, the *ERP*, and their predictors

This figure plots the time series of the bond risk premium (*BRP*), equity risk premium (*ERP*), and of each of the predictors. The *BRP* and *ERP* are measured as the difference between the return on the 10-year US Treasury bond and the return on the S&P500 index, respectively, and the return on a one-month T-bill. The set of predictors is described in Appendix 1. The sample period extends from 1973:01 to 2018:12.




Figure 2: Term spread time series and wavelet decomposition

This figure plots the time series of the term spread (TMS) and the seven frequency components into which the time series is decomposed. It is applied a $J = 6$ level wavelet decomposition, which produces six wavelet details (D_1, D_2, \dots, D_6), each representing higher-frequency characteristics of the series, as well as a wavelet smooth (D_7), which captures the low-frequency dynamics of the series. The sample period runs from 1973:01 to 2018:12 (monthly frequency).



Figure 3: *APM_WAV*, *APM_TS*, and balanced $Benchmark_{60-40}$ portfolio weights

This figure plots the *APM_WAV*, *APM_TS*, and balanced $Benchmark_{60-40}$ portfolio weights (solid, dashed and dotted lines, respectively), rebalanced on a monthly basis. *APM_WAV* and *APM_TS* stand for the active portfolio management strategy based on asset return forecasts from *multi_wav* and *multi_ts* methodologies, respectively. The sample period is from 1973:01 to 2018:12. The out-of-sample forecasting period runs from 1990:01 to 2018:12 (monthly frequency). Gray bars denote NBER-dated recessions.




Figure 4: Cumulative wealth for *APM_WAV*, *APM_TS*, and *Benchmark*₆₀₋₄₀ investors. This figure represents the cumulative wealth of an investor who begins with \$1 and reinvests all proceeds on a monthly basis, adopting an *APM_WAV*, *APM_TS*, and *Benchmark*₆₀₋₄₀ strategy (solid, dashed, and dotted lines, respectively). The *APM_WAV* and *APM_TS* active portfolio management strategies are based on asset return forecasts from *multi_wav* and *multi_ts* methodologies, respectively. The sample period extends from 1973:01 to 2018:12. The out-of-sample forecasting period runs from 1990:01 to 2018:12 (monthly frequency). Gray bars denote NBER-dated recessions.




Figure 5: 3-year moving average information ratios and CER gains

The upper figure plots the 3-year moving average information ratio for the *APM_WAV* and *APM_TS* strategies relative to the *Benchmark₆₀₋₄₀*. The lower figure plots the 3-year moving average annualized CER gain for the *APM_WAV* and the *APM_TS* strategies. The sample period is from 1973:01 to 2018:12. The out-of-sample forecasting period runs from 1990:01 to 2018:12 (monthly frequency). Gray bars denote NBER-dated recessions.




Figure 6: *APM_BLMWAV*, *APM_BLMTS*, and balanced *Benchmark₆₀₋₄₀* portfolio weights

This figure plots the *APM_BLMWAV*, *APM_BLMTS*, and balanced *Benchmark₆₀₋₄₀* portfolio weights (solid, dashed, and dotted lines, respectively), rebalanced on a monthly basis. The *APM_BLMWAV* and *APM_BLMTS* active portfolio management strategies are based on asset return forecasts from *multi_wav* and *multi_ts* methodologies used in the context of a Black-Litterman portfolio-optimization framework. The sample period extends from 1973:01 to 2018:12. The out-of-sample forecasting period runs from 1990:01 to 2018:12 (monthly frequency). Gray bars denote NBER-dated recessions.

Appendix 1. Predictors of equity and bond risk premiums

- 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).
- Excess stock return volatility (RVOL): calculated using a 12-month moving standard deviation estimator.
- Book-to-market ratio (BM): ratio of book value to market value for the DJIA.
- 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.
- 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.

Appendix 2. Maximal overlap discrete wavelet transform

Discrete wavelet transform (DWT) multiresolution analysis (MRA) allows the decomposition of a time series into its constituent multiresolution (frequency) components. There are two types of wavelets: father wavelets (ϕ), which capture the smooth and low frequency part of the series, and mother wavelets (ψ), which capture the high frequency components of the series, where $\int \phi(t) dt = 1$ and $\int \psi(t) dt = 0$.

Given a time series y_t with a certain number of observations N , its wavelet multiresolution representation is given by

$$y_t = \sum_k s_{J,k} \phi_{J,k}(t) + \sum_k d_{J,k} \psi_{J,k}(t) + \sum_k d_{J-1,k} \psi_{J-1,k}(t) + \cdots + \sum_k d_{1,k} \psi_{1,k}(t) , \quad (5)$$

where J represents the number of multiresolution levels (or frequencies), k defines the length of the filter, $\phi_{J,k}(t)$ and $\psi_{j,k}(t)$ are the wavelet functions, and $s_{J,k}$, $d_{J,k}$, $d_{J-1,k}, \dots, d_{1,k}$ are the wavelet coefficients.

The wavelet functions are generated from the father and mother wavelets through scaling and translation as follows

$$\begin{aligned} \phi_{J,k}(t) &= 2^{-J/2} \phi(2^{-J}t - k) \\ \psi_{j,k}(t) &= 2^{-j/2} \psi(2^{-j}t - k) , \end{aligned}$$

while the wavelet coefficients are given by

$$\begin{aligned} s_{J,k} &= \int y_t \phi_{J,k}(t) dt \\ d_{j,k} &= \int y_t \psi_{j,k}(t) dt , \end{aligned}$$

where $j = 1, 2, \dots, J$.

Due to the practical limitations of DWT in empirical applications, we perform wavelet decomposition analysis here by applying the maximal overlap discrete wavelet transform (MODWT).

Appendix 3. Out-of-sample R-squares for all predictors

Predictor	<i>ERP</i>			<i>BRP</i>		
	<i>single_ts</i>	<i>single_wav</i>	Frequency	<i>single_ts</i>	<i>single_wav</i>	Frequency
DP	-1.87	-0.59	D_6	-0.75	2.89**	D_1
DY	-1.96	-0.24	D_1	-0.42	-0.01	D_1
EP	-1.05	0.77**	D_3	-0.60	-0.50	D_7
RVOL	-0.73	-0.30	D_1	-0.26	-0.09	D_4
BM	-0.53	0.50**	D_5, D_6	-0.15	5.45***	D_1, D_2, D_5, D_7
NTIS	-3.05	0.15	D_1	-3.19	0.53	D_1
LTY	-0.32	0.08	D_6	-1.81	0.76**	D_4
LTR	-0.29	0.85**	D_7	-0.36	-4.16	D_7
TMS	-0.76	1.70***	D_7	1.70**	1.35**	D_5
DFY	-2.82	-0.86	D_6	-1.14	-0.25	D_7
DFR	-1.84	0.07	D_1	-1.13	-0.82	D_6
INFL	-0.61	1.77**	D_2, D_5	-0.77	-0.44	D_1

Table 4: Out-of-sample R-squares (R^2_{OS})

This table reports the out-of-sample R-squares as percentages for bond risk premium (*BRP*) and equity risk premium (*ERP*) forecasts at monthly frequencies of regressions using one original predictor at a time (*single_ts*) and regressions using the frequencies of one original predictor at a time (*single_wav*). The list of predictors is described in Appendix 1. The out-of-sample R-squares (R^2_{OS}) measures the proportional reduction in the mean squared forecast error for the predictive model relative to the forecast based on the historical mean (HM). The one-month-ahead out-of-sample forecast of the *BRP* and the *ERP* is generated using a sequence of expanding windows. The sample period is from 1973:01 to 2018:12. The out-of-sample forecasting period is from 1990:01 to 2018:12, monthly frequency. Asterisks denote the significance of the out-of-sample *MSFE*-adjusted statistic of Clark and West (2007). *** and ** denote significance at the 1% and 5% levels, respectively.

Appendix 4. Implemented version of the Black-Litterman model

There are N assets and K active investment views ($N = K = 2$: bonds and stocks). μ is an $N \times 1$ vector of expected excess returns: BRP and ERP forecasts for bonds and stocks, respectively. τ is a scaling parameter (which we set to unity as in Almadi, Rapach, and Suri, 2014), Σ is an $N \times N$ covariance matrix, P is a $K \times N$ matrix whose elements in each row represent the weight of each asset in each of the K -view portfolios, Ω is a matrix representing the confidence in each view, Q is a $K \times 1$ vector of expected excess returns of the K -view portfolios, and Π is a $N \times 1$ vector of the equilibrium excess returns of the assets.

The original Black-Litterman model (BLM) of expected excess returns in Black and Litterman (1992) is given by:

$$\mu = [(\tau \Sigma)^{-1} + P' \Omega^{-1} P]^{-1} [(\tau \Sigma)^{-1} \Pi + P' \Omega^{-1} Q] ,$$

which by applying the Matrix Inversion Lemma can be rewritten as follows (Da Silva, Lee, and Pornrojnangkool, 2009):

$$\mu = \Pi + \Sigma P' \left[\frac{\Omega}{\tau} + P \Sigma P' \right]^{-1} (Q - P \Pi) = \Pi + G , \quad (6)$$

where G is the term that captures the deviations of expected excess returns from the equilibrium due to active investment views. Equation (6) summarizes the key idea behind the BLM model: the expected excess return will be different from the equilibrium excess return if and only if investor views differ from equilibrium views.

The construction of the actively managed portfolios consists in two steps. First, we compute the posterior expected excess return vector, μ_{t+1} , and posterior return covariance matrix, Σ_{t+1} . We start from the selected vector of excess return forecasts (\widehat{BRP}_{t+1} and \widehat{ERP}_{t+1}) obtained from predictive regression models explained in section 2.2. We generate an exponen-

tially weighted moving average estimate of the monthly return covariance matrix $V = \hat{\Sigma}_{t+1}$.

We set the decay parameter to 0.97, which is frequently used for monthly series.

To set matrix Ω , we follow the suggestion of Da Silva, Lee, and Pornrojnangkool (2009) and use:

$$\frac{\Omega}{\tau} = \text{diag}(\text{diag}(PVP')) .$$

Adopting the approach of Idzorek (2004), the posterior return covariance matrix is given by:

$$\Sigma_{t+1} = [(\tau V)^{-1} + (P'\Omega^{-1}P)]^{-1} .$$

From expression (6) for the expected returns, by setting (i) the vector of the equilibrium excess returns of assets as $\Pi = 0$, as in Da Silva, Lee, and Pornrojnangkool (2009), (ii) using the vector of *BRP* and *ERP* forecasts as matrix Q and (iii) using the posterior return covariance matrix Σ_{t+1} , it is obtained the posterior expected excess return vector, μ_{t+1} .

The second step for the construction of the portfolio consists in using μ_{t+1} and Σ_{t+1} to obtain the portfolio weights. Recall that the objective function of an active asset manager is to maximize the return of the portfolio with a penalty on the square of tracking error towards the relevant benchmark:

$$\begin{aligned} \max \quad & (\boldsymbol{\varpi}_A + \boldsymbol{\varpi}_B)' \boldsymbol{\mu} - \lambda \boldsymbol{\varpi}_A' \Sigma \boldsymbol{\varpi}_A \\ \text{s.t.} \quad & \boldsymbol{\varpi}_A' \mathbf{1} = 0 \end{aligned} \tag{7}$$

where $\boldsymbol{\varpi}_A$ and $\boldsymbol{\varpi}_B$ are the vectors of active positions and benchmark portfolio weights, respectively. The parameter λ is given by $\lambda = \frac{1}{2TE} \sqrt{\Theta' \Sigma \Theta}$, with TE representing the tracking error (set to a constant annualized value of 5.80% as explained in section 4.1) and matrix Θ

is:

$$\Theta = \Sigma^{-1} \left(I - 1 \frac{1' \Sigma^{-1}}{1' \Sigma^{-1} 1} \right) \mu .$$

The active weights ϖ_A are given by $\varpi_A = \frac{\Theta}{2\lambda}$. Thus, total weights are $\varpi = \varpi_A + \varpi_B$. We assume the investor will neither leverage nor short-sell available assets (following the notation in the paper, $h = 1$ and $l = 0$). We further assume that the investor rebalances the portfolio at the same monthly frequency as the forecast horizon.