

**cef.up working paper
2017-02**

**FORECASTING STOCK MARKET RETURNS
BY SUMMING THE FREQUENCY-DECOMPOSED PARTS**

**Gonçalo Faria
Fabio Verona**

Forecasting stock market returns

by summing the frequency-decomposed parts*

Gonçalo Faria[†] Fabio Verona[‡]

20 November 2017,

Forthcoming in the *Journal of Empirical Finance*

* We thank Amit Goyal for kindly providing the data on his webpage; Rossen Valkanov (editor) and an anonymous referee for insightful comments that strongly improved the paper; Gene Ambrocio, João Correia da Silva, Eleonora Granziera, Adam Gulan, Esa Jokivuolle, Mikael Juselius, Donald Percival and seminar participants at the Bank of Finland for useful comments; Teppo Tammisto and Riccardo Verona for IT support; and Gregory Moore for editorial assistance. The views expressed 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 UID/GES/00731/2016.

[†] Universidade Católica Portuguesa – Católica Porto Business School and CEFE, 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)

Abstract

We generalize the Ferreira and Santa-Clara (2011) sum-of-the-parts method for forecasting stock market returns. Rather than summing the parts of stock returns, we suggest summing some of the frequency-decomposed parts. The proposed method significantly improves upon the original sum-of-the-parts and delivers statistically and economically gains over historical mean forecasts, with monthly out-of-sample R^2 of 2.60% and annual utility gains of 558 basis points. The strong performance of this method comes from its ability to isolate the frequencies of the parts with the highest predictive power, and from the fact that the selected frequency-decomposed parts carry complementary information that captures different frequencies of stock market returns.

Keywords: predictability, stock returns, equity premium, asset allocation, frequency domain, wavelets

JEL codes: G11, G12, G14, G17

1 Introduction

Predicting stock market returns has a long tradition in finance. A reliable forecast of stock market returns is a crucial input for the computation of the cost of capital and the investment process in general. It is hardly surprising then that stock return predictability has generated an immense body of literature and that interest in the field shows no signs of abating, as recent papers by Pettenuzzo, Timmermann and Valkanov (2014), Bollerslev, Todorov and Xu (2015), Huang et al. (2015) and Rapach, Ringgenberg and Zhou (2016) confirm. This paper contributes to this literature by proposing a method defined in the joint time-frequency domain to forecast stock market returns.

Our work builds on the Ferreira and Santa-Clara (2011) sum-of-the-parts (SOP) method for forecasting stock market returns. Conceptually, the SOP method consists in decomposing the stock market return into three parts. The parts are first forecasted separately and then added together to obtain the forecast of the stock market return. The SOP method improves forecast accuracy compared to the historical mean benchmark by exploiting the different time series persistence of the three parts. Our proposed method is a generalization of the SOP method in that we modify the forecasting approach and the summing of the three parts. Specifically, we suggest summing only some of the frequency-decomposed parts – obtained using a wavelet decomposition approach (WAV) – rather than the original three parts. The wavelet decomposition is becoming an increasingly popular tool in econometric analysis and in high-frequency and low-frequency asset pricing (see Hong and Kao, 2004, Galagedera and Maharaj, 2008, Xue, Gencay and Fagan, 2013, Gencay and Signori, 2015, Bandi et al., 2016, Hasbrouck, 2017 and the references in section 2.2). We refer to our method as the SOPWAV. In a nutshell, the first step of the SOPWAV method consists of decomposing a time series into n time series components, each capturing the oscillations of the original variable within a specific frequency interval. Lower frequencies represent the long-term dynamics of the

original time series, while the higher frequencies capture short-term dynamics. Notably, the original time series can be recovered simply by summing these n frequency-decomposed components. After decomposing the different parts of stock returns into their frequency time-series components, the second step consists of isolating the best frequencies in each part and exclude noisy frequencies. Hence, by only retaining the frequency-decomposed parts that have the greatest predictive power, the SOPWAV method leads to expressive statistically and economically forecasting improvements over both the historical mean and the original SOP method.

Figure 1 provides a preview of our results. It presents, for the out-of-sample period considered in this paper (January 1950 to December 2015), the realized S&P 500 index log return (black solid line) together with the forecasts based on the historical mean (HM) of returns (black dashed line), and the SOP and SOPWAV methods (red and blue solid lines, respectively). The out-of-sample forecast from the SOPWAV method clearly tracks the dynamics of stock market returns more closely than either the SOP or HM forecasts. The strength of the frequency-specific predictability, which is at the core of the SOPWAV method, is that it allows to capture both the low-frequency dynamics of stock market returns (*i.e.* the long-run trend, as in the HM and SOP methods), as well as higher-frequency movements of the stock market only partially captured by the SOP method and not at all by the HM method.

Using the HM as a benchmark, the monthly out-of-sample R^2 for the best specification of the SOPWAV method is 2.60 %. When examining the economic significance of the SOPWAV predictive performance through an asset allocation analysis, we find that a mean-variance investor who allocates her wealth between equities and risk-free bills enjoys significant utility gains from a SOPWAV-based trading strategy. Specifically, the rate of return that she would be willing to accept instead of holding the risky portfolio is 558 basis points. Furthermore, the annualized Sharpe ratio of the strategy based on the SOPWAV method is 0.73, which is about 2.4 times the Sharpe ratio generated by the HM-based strategy. Those statistical

and economic gains are also clearly larger than the gains achieved using the original SOP method (out-of-sample R^2 of 1.30 %, utility gains of 206 basis points and Sharpe ratio of 0.53), reflecting the improved accuracy of the stock return forecasting when considering information in the joint time-frequency domain.

The rest of the paper is structured as follows. Section 2 reviews the two strands of literature on which our work builds to provide context for our contribution. Section 3 presents the data and the methodology. We elaborate on the two blocks of the SOPWAV method: the original SOP method and the discrete wavelet decomposition of the predictors. The out-of-sample forecasting procedure and asset allocation analysis are also described. Section 4 presents the out-of-sample forecasting results. In Section 5 are reported the results of some robustness tests. Section 6 concludes.

2 Literature review

Our method is rooted in two strands of literature. The first deals with the out-of-sample stock return predictability using standard time series tools. The second involves the application of wavelet methods to economic and finance topics. The following brief overview of these two areas provides context for the contribution of this paper.

2.1 Forecasting stock returns

As evidenced in the reviews of Rapach and Zhou (2013) and Harvey, Liu and Zhu (2016), the literature on predicting stock returns and equity premium is vast. Several studies discuss the in-sample predictability using predictors such as the treasury bill rate, dividend yield, dividend-price ratio, term spread, equity market volatility or the consumption-wealth ratio. This is the case for the US stock market (see *e.g.* Fama and Schwert, 1977, Campbell,

1987, Ferson and Harvey, 1991, Lettau and Ludvigson, 2001, Cochrane, 2008, Goyal and Welch, 2008 and Pastor and Stambaugh, 2009), as well as for other stock markets (see *e.g.* Cutler, Poterba and Summers, 1991, Harvey, 1991, Bekaert and Hodrick, 1992, Ferson and Harvey, 1993, Ang and Bekaert, 2007 and Hjalmarsson, 2010). Noting that predictive models require out-of-sample validation, Goyal and Welch (2008) show how poorly the above-mentioned predictors perform out-of-sample up to 2008. To overcome poor out-of-sample performance, researchers have since then turned their attention to improving the out-of-sample forecastability of stock returns, exploring two different avenues.

The first one focuses on developing and testing new predictors. For example, Bollerslev, Tauchen and Zhou (2009) test the variance risk premium, Cooper and Priestley (2009, 2013) use the output gap and the world business cycle, Rapach, Strauss and Zhou (2013) document the relevance of lagged US market returns for the out-of-sample predictability of stock returns of other industrialized countries, Li, Ng and Swaminathan (2013) study the aggregate implied cost of capital, Neely et al. (2014) consider the relevance of several technical indicators to serve as complementary predictors to the traditional set of variables, Huang et al. (2015) propose a new investor sentiment index, Moller and Rangvid (2015) study different macroeconomic variables by focusing on their fourth-quarter growth rate, and Rapach, Ringgenberg and Zhou (2016) construct an aggregate short interest position indicator.

The second avenue focuses on improving existing forecasting methods. For example, Ludvigson and Ng (2007) and Kelly and Pruitt (2013) propose using dynamic factor analysis for large data sets to summarize a large amount of information by few estimated factors, Rapach, Strauss and Zhou (2010) and Pettenuzzo and Ravazzolo (2016) suggest combining individual forecasts from different predictors, Ferreira and Santa-Clara (2011) introduce the SOP method, Dangl and Halling (2012) evaluate predictive regressions that explicitly consider the time-variation of coefficients, Pettenuzzo, Timmermann and Valkanov (2014) propose an approach to impose economic constraints on forecasts of the equity premium,

Bollerslev, Todorov and Xu (2015) decompose the predictor (the variance risk premium) into a jump and a diffusion component, Baetje and Menkhoff (2016) examine the time-instability of the standard set of economic and technical indicators, and Lima and Meng (2017) propose a quantile combination approach.

Our contribution straddles both strands of literature. In fact, the frequency decomposition of the parts of stock market returns, which is a methodological contribution in itself, also represents an expansion of the set of possible predictors, as each frequency of each part can be understood and potentially used as a new predictor.

2.2 Wavelets applications in economics and finance

Wavelets are a signal processing technique defined over a finite window in the time domain. The size of the window changes according to the frequency of interest. The high frequency features of the time series are isolated with a short window, while looking at the same signal with a large window reveals instead its low frequency features. Hence, by varying the size of the time window, it is possible to capture simultaneously both time-varying and frequency-varying features of the time series. Wavelets are thus extremely useful when the time series have structural breaks or jumps, as well as with non-stationary time series. Moreover, as wavelets allow frequency decomposition in the time domain, they are well suited to finance applications. For example, their use makes it simpler to extend standard asset pricing models in a joint time-frequency domain (see *e.g.* Ortu, Tamoni and Tebaldi, 2013).

Crowley (2007) and Aguiar-Conraria and Soares (2014) provide excellent reviews of economic and finance applications of wavelets.¹ The pioneering works in these fields are Ramsey and Lampart (1998*a,b*), in which wavelets are used to study the relationship between macroeconomic variables (consumption versus income and money supply versus income, respectively).

¹ Wavelet methods have long been popular in many fields, including astronomy, engineering, geology, medicine, meteorology and physics.

More recently, wavelets methods have been used to test for the (in-sample) frequency dependence between two (or more) variables (Kim and In, 2005, Gencay, Selcuk and Whitcher, 2005, Gallegati et al., 2011, Gallegati and Ramsey, 2013 and Kilponen and Verona, 2016), and to study the comovements and lead-lag relationship between variables at different frequencies (Rua and Nunes, 2009, Rua, 2010, Aguiar-Conraria and Soares, 2011 and Aguiar-Conraria, Martins and Soares, 2012). An emerging area of research is the application of wavelet methods for forecasting purposes. Rua (2011, 2017) proposes a wavelet-based multiscale principal component analysis to forecast GDP growth and inflation and finds that significant predictive short-run improvements can be achieved with wavelets in combination with factor-augmented models. Similarly, Kilponen and Verona (2016) forecast aggregate investment using the Tobin's Q theory of investment and find that the quality of short-term forecasts is significantly enhanced when merging the wavelet approach with the proxies for Q and investment recently suggested in the literature.

In this paper, we document the statistical and economic advantages of applying wavelet decomposition in the context of forecasting stock market returns.

3 Data and methodology

We focus on the out-of-sample (OOS) predictability of monthly stock returns, proxied by the S&P500 index total return. We use monthly data from December 1927 to December 2015. Throughout the analysis, we use a set of popular predictors from the literature, which are briefly described in Appendix 1.

Our methodology to forecast stock market returns builds on two blocks: the SOP method proposed by Ferreira and Santa-Clara (2011) and the time-frequency decomposition of economic time series. We describe these blocks in sub-sections 3.1 and 3.2, respectively. The

OOS procedure and asset allocation analysis are described in sub-sections 3.3 and 3.4, respectively.

3.1 SOP decomposition of stock market returns

Ferreira and Santa-Clara (2011) propose the SOP method for forecasting stock market returns. Conceptually, this method consists in decomposing the stock market return into three parts. Each part is first forecasted separately, and then they are added together to obtain the stock market return forecast.

The stock market total return from time t to time $t+1$, R_{t+1} , can be decomposed into capital gains, CG_{t+1} , and dividend yield, DY_{t+1} :

$$1 + R_{t+1} = 1 + CG_{t+1} + DY_{t+1} = \frac{P_{t+1}}{P_t} + \frac{D_{t+1}}{P_t} , \quad (1)$$

where P_{t+1} is the stock price at time $t+1$ and D_{t+1} is the dividend per share paid between time t and $t+1$. Each component in equation (1) is then further decomposed.

Capital gains can be rewritten as

$$\begin{aligned} 1 + CG_{t+1} &= \frac{P_{t+1}}{P_t} = \frac{P_{t+1}/E_{t+1}}{P_t/E_t} \frac{E_{t+1}}{E_t} = \frac{M_{t+1}}{M_t} \frac{E_{t+1}}{E_t} = \\ &= (1 + GM_{t+1})(1 + GE_{t+1}) , \end{aligned} \quad (2)$$

where $M_{t+1} = P_{t+1}/E_{t+1}$ is the price-earnings multiple, GM_{t+1} is the price-earnings multiple growth rate, and GE_{t+1} is the earnings growth rate.

The dividend yield can be rewritten as:

$$DY_{t+1} = \frac{D_{t+1}}{P_t} = \frac{D_{t+1}}{P_{t+1}} \frac{P_{t+1}}{P_t} = DP_{t+1} (1 + GM_{t+1})(1 + GE_{t+1}) , \quad (3)$$

where $DP_{t+1} = D_{t+1}/P_{t+1}$ is the dividend-price ratio.

Substituting equations (2) and (3) in (1), the total stock market return can then be rewritten as

$$\begin{aligned} 1 + R_{t+1} &= (1 + GM_{t+1})(1 + GE_{t+1}) + DP_{t+1}(1 + GM_{t+1})(1 + GE_{t+1}) \\ &= (1 + DP_{t+1})(1 + GE_{t+1})(1 + GM_{t+1}) \ , \end{aligned} \quad (4)$$

whereby stock market return is the product of the dividend-price ratio, the growth rates of earnings and of the price-earnings multiple. Finally, by taking the logs on both sides of equation (4), the log stock return is given by the sum of the dividend-price ratio, the growth in earnings and the growth in the price-earnings multiple:

$$r_{t+1} = \log(1 + R_{t+1}) = dp_{t+1} + ge_{t+1} + gm_{t+1} \ , \quad (5)$$

where lowercase variables denote log rates.

3.2 Time-frequency decomposition of economic time series

The discrete wavelet transform (DWT) multiresolution analysis (MRA) allows the decomposition of a time series into its constituent multiresolution (frequency) components.² 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

² A detailed analysis of wavelet methods can be found in Percival and Walden (2000).

representation is given by

$$y_t = \sum_k s_{J,k} \phi_{J,k}(t) + \sum_k d_{J,k} \psi_{J,k}(t) + \sum_k d_{J-1,k} \psi_{J-1,k}(t) + \cdots + \sum_k d_{1,k} \psi_{1,k}(t) , \quad (6)$$

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

The wavelet multiresolution decomposition of y_t (equation 6) can be rewritten in a more synthetic way as

$$y_t = y_t^{S_J} + y_t^{D_J} + y_t^{D_{J-1}} + \cdots + y_t^{D_1} , \quad (7)$$

where $y_t^{S_J} = \sum_k s_{J,k} \phi_{J,k}(t)$ is the wavelet smooth component and $y_t^{D_j} = \sum_k d_{j,k} \psi_{j,k}(t)$, $j = 1, 2, \dots, J$, are the J wavelet detail components. Equation (7) shows that the original series y_t , exclusively defined in the time domain, can be decomposed in different components,

each defined in the time domain and representing the fluctuation of the original time series in a specific frequency band. In particular, for small j , the j wavelet detail components represent the higher frequency characteristics of the time series (*i.e.* its short-term dynamics). As j increases, the j wavelet detail components represent lower frequencies movements of the series. Finally, the wavelet smooth component captures the lowest frequency dynamics (*i.e.* its long-term behavior or trend).

Because of the practical limitations of DWT in empirical applications, in this paper we perform wavelet decomposition analysis by applying the maximal overlap discrete wavelet transform (MODWT) and using the Haar wavelet filter with reflecting boundary conditions.^{3,4} In our analysis, given the availability of long data series, we apply a $J = 7$ levels MRA.⁵ Thus, the wavelet decomposition delivers eight components: seven wavelet details ($y_t^{D_1}$ to $y_t^{D_7}$) and a wavelet smooth ($y_t^{S_7}$). Since we employ monthly data in our analysis, the first detail component $y_t^{D_1}$ captures oscillations between 2 and 4 months, the second detail component $y_t^{D_2}$ captures oscillations between 4 and 8 months, while detail components $y_t^{D_3}$, $y_t^{D_4}$, $y_t^{D_5}$, $y_t^{D_6}$ and $y_t^{D_7}$ capture oscillations with a period of 8-16, 16-32, 32-64, 64-128 and 128-256 months, respectively. Finally, the smooth component $y_t^{S_7}$, which we now re-denote as $y_t^{D_8}$, captures oscillations with a period longer than 256 months (21.3 years).

As an example, Figure 2 plots the time series of the (log) stock market return (top graph) and its MODWT MRA decomposition (remaining graphs). Overall, the frequency-decomposed time series exhibit significantly different dynamics that are not visible from the original time

³ Unlike the DWT, the MODWT i) is not restricted to any sample size, ii) is translation-invariant, so that it is not sensitive to the choice of the starting point of the examined time series, and iii) does not introduce phase shifts in the wavelet coefficients, so that peaks or troughs in the original time series are correctly aligned with similar events in the MODWT MRA.

⁴ Besides its simplicity and wide use (see *e.g.* Manchaldore, Palit and Soloviev, 2010 and Malagon, Moreno and Rodriguez, 2015), the Haar filter makes a neat connection to temporal aggregation as the wavelet coefficients are simply differences of moving averages (see Ortu, Tamoni and Tebaldi, 2013 and Bandi et al., 2016).

⁵ As regards the choice of J , the number of observations dictates the maximum number of frequency bands that can be used. In our case, $t_0 = 265$ is the number of observations in the in-sample period, so J is such that $J \leq \log_2 t_0 \simeq 8$.

series as all these different frequencies are, in practice, aggregated. As expected, the lower the frequency, the smoother the resulting filtered time series.

3.3 Out-of-sample forecasts

One-step-ahead OOS forecasts of stock market returns are generated using a sequence of expanding windows. We use an initial sample (December 1927 to December 1949) to make the 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 in this way until the end of the sample, thus obtaining a sequence of 792 one-step-ahead OOS forecasts. The OOS period spans the period from January 1950 to December 2015.

3.3.1 Forecast evaluation

We evaluate the forecast performance of the SOP/SOPWAV methods in terms of the Campbell and Thompson (2008) OOS R-square (R_{OS}^2). The R_{OS}^2 statistic measures the proportional reduction in the mean squared forecast error (MSFE) for the predictive method relative to a benchmark model and is given by

$$R_{OS}^2 = 1 - \frac{\sum_{t=t_0}^{T-1} (r_{t+1} - E_t r_{t+1})^2}{\sum_{t=t_0}^{T-1} (r_{t+1} - \bar{r}_t)^2}, \quad (8)$$

where $E_t r_{t+1}$ is the stock return forecast for $t+1$ from the SOP/SOPWAV methods (see sub-sections 3.3.2. and 3.3.3), \bar{r}_t is the historical mean of stock market returns up to time t , r_{t+1} is the realized stock market return in $t+1$, T is the total number of observations in the sample and t_0 is the number of observations in the initial sample. As it is standard in the literature, we choose the HM as the benchmark method. According to (8), a positive (negative) R_{OS}^2 indicates that the SOP/SOPWAV method outperforms (underperforms) the

HM in terms of MSFE.

As in Ferreira and Santa-Clara (2011), we evaluate the statistical significance of the results using the *MSFE-F* statistic proposed by McCracken (2007).⁶ The *MSFE-F* statistic tests for the equality of the MSFE of the HM and the SOP/SOPWAV method forecasts as follows:

$$MSFE - F = (T - t_0) \left[\frac{\sum_{t=s_0}^{T-1} (r_{t+1} - \bar{r}_t)^2 - \sum_{t=t_0}^{T-1} (r_{t+1} - E_t r_{t+1})^2}{\sum_{t=t_0}^{T-1} (r_{t+1} - E_t r_{t+1})^2} \right] .$$

3.3.2 Forecasting with the SOP method

To forecast the aggregate stock market return, the SOP method forecasts separately each part of the stock market return as derived in equation (5). Let $E_t r_{t+1}$ denote the expected stock market return at time t for period $t+1$. Forecasting with the SOP method thus implies that

$$E_t r_{t+1} = E_t dp_{t+1} + E_t ge_{t+1} + E_t gm_{t+1} , \quad (9)$$

that is, the expected stock market return is the sum of the expected dividend-price ratio ($E_t dp_{t+1}$), the expected earnings growth ($E_t ge_{t+1}$) and the expected price-earnings multiple growth ($E_t gm_{t+1}$). To forecast the dividend-price ratio, Ferreira and Santa-Clara (2011) assume that dp follows a random walk so that the expected dividend-price ratio equals the current dividend-price ratio (*i.e.* $E_t dp_{t+1} = dp_t$). As regards the expected earnings growth, Ferreira and Santa-Clara (2011) assume that it is captured by the 20-year moving average of the growth in earnings up to time t (\overline{ge}_t), which is consistent with the view that earnings growth has a low-frequency predictable component.

The expected price-earnings multiple growth is assumed to be zero in the baseline version of the SOP. In its extended version, it is predicted by running OLS regressions with the price-earnings multiple growth gm as the dependent variable: $gm_{t+1} = a + bx_t + \varepsilon_{t+1}$, where

⁶ We thank Michael McCracken for providing us with the additional tables of critical values for evaluating the statistical significance of the R_{OS}^2 .

x_t is one of the stock market predictors commonly used in the literature. The forecast of the price-earnings multiple growth is thus given by: $E_t gm_{t+1} = \hat{a} + \hat{b}x_t$, where \hat{a} and \hat{b} are the OLS estimates of a and b , respectively. As regards the predictors x_t , we use the stock variance, the default return spread, the long-term government bond yield, the long-term government bond return, the inflation rate, the term spread, the treasury bill rate, the default yield spread, the net equity expansion, the return on equity, the dividend-payout ratio, the earnings-price ratio, the excess stock return volatility, the dividend-price ratio and the book-to-market ratio (see Appendix 1 for a brief description of these variables).

Overall, the SOP method forecast can thus be written as

$$E_t r_{t+1} = dp_t + \overline{ge}_t + \hat{a} + \hat{b}x_t .$$

3.3.3 Forecasting with the SOPWAV method

We generalize the SOP method to forecast stock market returns as given by equation (9). To forecast each part of the expected stock market return, we use the frequency-decomposed time series of: the dividend-price ratio, earnings growth, price-earnings multiple growth and the above-listed stock market predictors. We now explain in detail the SOPWAV econometric time-series method.

First, we apply the MODWT MRA decomposition ($J = 7$ levels) to all time series under analysis. As an illustration for dp , we obtain:

$$dp_t = dp_t^{D1} + dp_t^{D2} + dp_t^{D3} + dp_t^{D4} + dp_t^{D5} + dp_t^{D6} + dp_t^{D7} + dp_t^{D8} . \quad (10)$$

Note that, by using this decomposition, we are using the same amount of information as in the standard time series analysis, since the sum of the components gives exactly the original time series.

Second, we forecast separately each frequency-decomposed part of the stock market return (dp , ge and gm). To forecast the frequency-decomposed parts of dp and ge , we only use components of dp and ge , respectively. In particular, we postulate that each frequency-decomposed part of dp and ge follows an AR(1) process:

$$dp_{t+1}^{D_j} = \alpha_j + \beta_j dp_t^{D_j} + \varepsilon_{t+1} \quad \forall j = 1, \dots, 8$$

$$ge_{t+1}^{D_j} = \gamma_j + \delta_j ge_t^{D_j} + \varepsilon_{t+1} \quad \forall j = 1, \dots, 8 \quad ,$$

so that the forecast of each frequency-decomposed part is given by

$$E_t dp_{t+1}^{D_j} = \hat{\alpha}_j + \hat{\beta}_j dp_t^{D_j} \quad \forall j = 1, \dots, 8$$

$$E_t ge_{t+1}^{D_j} = \hat{\gamma}_j + \hat{\delta}_j ge_t^{D_j} \quad \forall j = 1, \dots, 8 \quad ,$$

where $\hat{\alpha}_j$, $\hat{\beta}_j$, $\hat{\gamma}_j$ and $\hat{\delta}_j$ are the OLS estimates of α_j , β_j , γ_j and δ_j , respectively.

To forecast each frequency-decomposed part of gm , we run OLS predictive regressions with the frequency component of the price-earnings multiple growth as the dependent variable and each frequency component of the x_t variable (chosen one at a time) as the independent variable:

$$gm_{t+1}^{D_j} = \eta_j + \lambda_j x_t^{D_j} + \varepsilon_{t+1} \quad \forall j = 1, \dots, 8 \quad ,$$

and then obtain the forecast of the price-earnings multiple growth as:

$$E_t gm_{t+1}^{D_j} = \hat{\eta}_j + \hat{\lambda}_j x_t^{D_j} \quad \forall j = 1, \dots, 8 \quad ,$$

where $\hat{\eta}_j$ and $\hat{\lambda}_j$ are the OLS estimates of η_j and λ_j , respectively.

Note that the frequency component at level j is used to forecast only the frequency component at the same level. This method is similar to Bandi et al. (2016) and Ortu et al. (2016), whereas it differs from Renaud, Starck and Murtagh (2002) wavelet-based (multiple resolution decomposition) forecasting method, in which the frequencies of the predictor are used to estimate the original time series of the predicted variable. Finally, as the MODWT MRA at any given point in time uses information of neighboring data points (both past and future), we recompute the frequency components at each iteration of the OOS forecasting process to make sure that we only use current and past information when making the forecasts. Hence, our SOPWAV method does not suffer from a look-ahead bias.

Third, the forecasts of dp_{t+1} , ge_{t+1} and gm_{t+1} are obtained by summing up the forecasts of their respective components (as in Rua, 2011):

$$E_t dp_{t+1} = E_t dp_{t+1}^{D_1} + E_t dp_{t+1}^{D_2} + \dots + E_t dp_{t+1}^{D_8} = \sum_{i=1}^8 \left(\hat{\alpha}_i + \hat{\beta}_i dp_t^{D_i} \right)$$

$$E_t ge_{t+1} = E_t ge_{t+1}^{D_1} + E_t ge_{t+1}^{D_2} + \dots + E_t ge_{t+1}^{D_8} = \sum_{i=1}^8 \left(\hat{\gamma}_i + \hat{\delta}_i ge_t^{D_i} \right)$$

$$E_t gm_{t+1} = E_t gm_{t+1}^{D_1} + E_t gm_{t+1}^{D_2} + \dots + E_t gm_{t+1}^{D_8} = \sum_{i=1}^8 \left(\hat{\eta}_i + \hat{\lambda}_i x_t^{D_i} \right) . \quad (11)$$

Finally, in line with the SOP method, we sum the forecasts of dp_{t+1} , ge_{t+1} and gm_{t+1} to obtain the one-step-ahead forecast of stock market returns, $E_t r_{t+1}$. For each predictor x used to forecast gm , the SOPWAV forecasting method of stock market returns is given by

$$E_t r_{t+1} = E_t dp_{t+1} + E_t ge_{t+1} + E_t gm_{t+1} = \sum_{i=1}^8 \left(\hat{\alpha}_i + \hat{\beta}_i dp_t^{D_i} \right) + \sum_{i=1}^8 \left(\hat{\gamma}_i + \hat{\delta}_i ge_t^{D_i} \right) + \sum_{i=1}^8 \left(\hat{\eta}_i + \hat{\lambda}_i x_t^{D_i} \right) . \quad (12)$$

In equation (12) we are considering all frequencies of each part (dp , ge and gm) to obtain the forecast of stock market returns. We refer to this specification as SOPWAV_ALL.

The inclusion of all frequencies of dp , ge and gm may reduce the stock return predictability by introducing noisy time-series components into the forecasting exercise. This compels a further – and, to our best knowledge, novel – step in the analysis. Namely, starting from (12), we select those frequency-decomposed parts that improve the stock return forecast and leave out those that destroy predictability. Technically, this is achieved by properly selecting the weights ω_i , $\omega_i = \{0, 1\}$, in

$$E_t r_{t+1} = \sum_{i=1}^8 \omega_i \left(\hat{\alpha}_i + \hat{\beta}_i dp_t^{D_i} \right) + \sum_{i=1}^8 \omega_{i+8} \left(\hat{\gamma}_i + \hat{\delta}_i ge_t^{D_i} \right) + \sum_{i=1}^8 \omega_{i+16} \left(\hat{\eta}_i + \hat{\lambda}_i x_t^{D_i} \right) , \quad (13)$$

to maximize the R_{OS}^2 . We refer to this specification as SOPWAV.

The economic intuition for this frequency selection in the right hand side of equation (13) results from the fact that stock market returns, in practice, may only be predictable at some frequencies (*e.g.* the idea of scale-specific predictability explored in Bandi et al., 2016), and that each part (dp , ge and gm) in equation (9) may have different predictive power at those frequencies. For example, suppose that stock market returns are only predictable at the highest frequency (scale D_1) and at the lowest frequency (scale D_8). Suppose further that the dividend price ratio dp and the earnings growth ge have relevant predictive power at both frequencies (and only at those frequencies), and that the growth of price-earnings multiple gm has no predictive power at any frequency. In such case, $\omega_1 = \omega_8 = \omega_9 = \omega_{16} = 1$ and $\omega_i = 0$ for $i = \{2, \dots, 24\} / \{8, 9, 16\}$. Thus, by only considering the relevant frequencies of the relevant parts (*i.e.* the frequency-decomposed parts that have predictive power) and excluding the remaining frequencies/parts, the SOPWAV method should improve the forecasting of stock market returns.

This approach resembles the spirit of the SOP method proposed by Ferreira and Santa-

Clara (2011): forecasting individually different parts of stock market returns to exploit their different time series characteristics and, ultimately, to improve the forecast of stock market returns. For instance, the baseline version of the SOP method (which ignores gm) assumes that dp follows a random walk and forecasts ge as the 20-year moving average of the growth in earnings (\overline{ge}). Conceptually, dp and ge are used to capture the higher and lower frequencies of stock market returns, respectively. Analytically, this is a very particular case of the SOPWAV method in (13). It corresponds to the specification where (i) for $i = 1, \dots, 8$: $\omega_i = 1$; $\hat{\alpha}_i = 0$; $\hat{\beta}_i = 1$, *i.e.*, dp is a random walk; (ii) for $i = 9, \dots, 16$, ω_i , $\hat{\gamma}_i$ and $\hat{\delta}_i$ are such that $\sum_{i=1}^8 \omega_{i+8} (\hat{\gamma}_i + \hat{\delta}_i ge_t^{D_i}) \approx \overline{ge}$; and (iii) for $i = 17, \dots, 24$, $\omega_i = 0$, *i.e.* gm is ignored. The SOPWAV method in (13) is more general and flexible (in the selection of the pair frequencies/parts) than the SOP method proposed by Ferreira and Santa-Clara (2011) and should then lead to more accurate forecasts. We support this claim with empirical evidence in section 4.

3.4 Asset allocation analysis

To quantify the economic value of the SOP/SOPWAV methods from an asset allocation perspective, we follow, among others, Kandel and Stambaugh (1996) and Rapach, Ringgenberg and Zhou (2016), and consider a mean-variance investor who dynamically allocates her wealth between equities and risk-free bills.⁷ The asset allocation decision is made at the end of month t , and the optimal share allocated to equities during month $t+1$ is given by

$$w_t = \frac{1}{\gamma} \frac{E_t r_{t+1} - r f_{t+1}}{\hat{\sigma}_{t+1}^2}, \quad (14)$$

⁷ We follow a standard mean-variance framework, *i.e.* without any frequency-specific consideration in its construction. It might be fairer to compare the performance of the SOPWAV and HM investors using a generalized mean-variance optimization framework that allows for frequency-specific components. This is important, for example, to investors that face risk constraints imposed at different frequencies. One way of implementing this generalization (suggested, but not explored by Chaudhuri and Lo, 2016) is to introduce a “regularization” term into the objective function (cost function), which is penalized if the resulting portfolio is excessively concentrated in a particular frequency.

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), γ is the investor's relative risk aversion coefficient, $E_t r_{t+1}$ is the SOP/SOPWAV predicted stock market return at time t for period $t+1$, and $\hat{\sigma}_{t+1}^2$ is the forecast of the variance of stock market returns. As in Ferreira and Santa-Clara (2011), we assume a relative risk aversion coefficient of two and we use all available data up to time t of past stock market returns to estimate the variance of stock returns. In order to introduce realistic limits on the possibilities of short selling and leveraging the portfolio, we also restrict the weights w_t to lie in the range between -0.5 and 1.5.

The portfolio return at time $t+1$, rp_{t+1} , is given by $rp_{t+1} = w_t r_{t+1} + (1 - w_t) rf_{t+1}$. We first compute the average utility (or certainty equivalent return, CER) of an investor that uses the allocation rule (14). The CER is given by $CER = \bar{r}_p - 0.5\gamma\sigma_{rp}^2$, where \bar{r}_p and σ_{rp}^2 are the sample mean and variance of the portfolio return, respectively. We then report the average utility gain, which is computed as the difference between the CER for an investor that uses the SOP/SOPWAV method to forecast stock returns and the CER for an investor who uses the HM forecasting strategy. The CER gain is annualized and can be interpreted as the annual management fee that an investor would be willing to pay to be exposed to a trading strategy based on the SOP/SOPWAV method instead of one based on the HM forecast (Rapach, Ringgenberg and Zhou, 2016). We also compute the Sharpe ratio (SR) of the portfolio, *i.e.* the mean portfolio return in excess of the risk-free rate divided by the standard deviation of the excess portfolio return. We then report the annualized Sharpe ratios gains, calculated as the difference between the SR of a trading strategy based on the SOP/SOPWAV method and the SR of a trading strategy based on the HM forecast.

4 Out-of-sample forecasting results

4.1 Statistical performance

As Figure 1 shows, the OOS forecast of the SOPWAV method follows the dynamics of the effective stock market returns much more closely than the HM or the SOP method forecasts. We use the R_{OS}^2 statistic as given by (8) to evaluate the statistical performance of the SOP/SOPWAV method. Results are reported in Table 1.⁸

Following Ferreira and Santa-Clara (2011), we first evaluate versions of the forecasting methods with only two parts, dp and ge (“Baseline” versions). The Ferreira and Santa-Clara (2011) SOP method outperforms the HM, delivering a positive R_{OS}^2 of 0.91. The SOPWAV method that uses *all* frequencies of both parts to obtain the forecast of stock market returns (SOPWAV_ALL) strongly underperforms the HM (R_{OS}^2 of -95.5 %). This reflects the existence of noisy components, making the forecasting exercise too imprecise. In contrast, the SOPWAV specification (13), which only uses some of the frequency-decomposed parts, delivers a positive R_{OS}^2 of 0.97.⁹ This represents an outperformance of both the HM and the SOP method. The frequencies chosen in the baseline SOPWAV method are $\{D_1, D_2, D_5, D_6, D_8\}$ for dp and $\{D_2, D_8\}$ for ge (see last eight columns in Table 1). This means that both dp and ge are useful to predict the higher frequencies (D_1 and D_2 , *i.e.* up to 8 months) and the lower frequency (D_8 , *i.e.* more than 256 months) of stock market returns, while dp also contains relevant predictability power of stock market returns at frequencies between 32 and 128 months (D_5 and D_6). All the remaining frequencies of stock market returns are not predictable when using only dp and ge as predictors.

⁸ All the simulations were run using the Matlab Wavelet Toolbox and were parallelized using the Techila technology (see <http://www.techilatechnologies.com/>).

⁹ In this simulation, the weights w_i , $i = 1, \dots, 16$ in (13) are chosen (among the 65k possible combinations of $dp_t^{D_i}$ and $ge_t^{D_i}$) in order to maximize the R_{OS}^2 (while the weights ω_i , $i = 17, \dots, 24$ are set to zero, *i.e.* gm is ignored).

Next, we evaluate the statistical performance of extended versions of the method that include the third part (gm) of stock market returns. As in Ferreira and Santa-Clara (2011), gm is forecasted using bivariate regressions of gm on different variables. The results are presented in Table 1 in lines aggregated under the “Extended” label. All extended versions of the SOP method continue to outperform the HM benchmark (*i.e.* all post positive R_{OS}^2), but in some cases the extended versions of the SOP method underperform the baseline version. The extended SOPWAV_ALL method also continues to strongly underperform the HM benchmark. Differently, the performance of the extended SOPWAV method is remarkably good. Regardless of the variable used to forecast gm , the extended SOPWAV method always outperforms (i) the HM benchmark, (ii) the baseline SOPWAV method and (iii) any version of the SOP method (baseline or extended).¹⁰ The strongest performance is achieved when the price-earnings multiple growth gm is forecasted using the long-term government yield as predictor (R_{OS}^2 of 2.60). Interestingly, the incremental statistical performance from the inclusion of gm in the baseline version of the SOPWAV (an increase in the R_{OS}^2 from 0.97 to 2.60) results exclusively from the gm predictive power of stock market returns in the frequency D_4 , *i.e.* between 16 and 32 months, which is not captured either by dp or ge .

Looking at the selected frequencies of dp , ge and gm , our analysis clearly shows that it is crucial to select the relevant frequencies (those at which stock returns are predictable) of the relevant parts (those that have predictive power at relevant frequencies) when forecasting stock market returns.¹¹

¹⁰ The predictor used to forecast gm is represented by the variable x in equation (13). For each predictor, given the weights chosen for the baseline SOPWAV method, the remaining weights ω_i , $i = 17, \dots, 24$ of the extended method result from the combination (among the 255 possible combinations of $gm_t^{D_i}$ plus the baseline) that gives the highest R_{OS}^2 . We also evaluate the extended SOPWAV method by choosing the weights combination ω_i , $i = 1, \dots, 24$, among all possible combinations of $dp_t^{D_i}$, $ge_t^{D_i}$ and $gm_t^{D_i}$, that maximizes the R_{OS}^2 . Results are broadly unchanged.

¹¹ If we use the original dp time series (in combination with the chosen frequencies for ge), then the R_{OS}^2 is 0.95. This value is very close to the one obtained by the baseline SOPWAV method (R_{OS}^2 of 0.97), which uses a subset of the frequencies of dp . Similarly, if we use \overline{ge} (in combination with the chosen frequencies of dp) instead of ge^{D_2} and ge^{D_8} (which are the chosen frequencies of ge), then the R_{OS}^2 is 0.92, which is slightly lower than the baseline SOPWAV method and close to the baseline SOP method (R_{OS}^2 of 0.91).

4.2 Asset allocation

In the previous sub-section, we show that the SOP/SOPWAV methods deliver statistically significant forecasting gains. We now quantify the economic value of these methods for stock return forecasting from an asset allocation perspective.

Columns 2 and 3 in Table 2 report the results of the CER analysis. The baseline SOPWAV method clearly outperforms the HM benchmark as well as the baseline SOP method. The performance of the extended SOPWAV method is also remarkably strong. It always outperforms the HM benchmark and any version of the SOP method (baseline or extended), and, for almost all variables used to forecast gm , it also outperforms the baseline SOPWAV method. The highest annualized CER gain for an investor who trades using the extended SOPWAV method is 558 basis points (using the long-term government bond yield to predict gm). The SOPWAV method outperforms both the HM benchmark and the SOP method also in terms of annualized Sharpe ratio gains, as reported in columns 4 and 5 in Table 2. The highest annualized SR gain (0.43) is again obtained using the extended SOPWAV method with the long-term government bond yield used to predict gm .¹²

Additionally, Figure 3 provides a dynamic perspective of the portfolio and cumulative wealth for an investor using the HM forecast, the SOP method and the SOPWAV method. The specifications of the SOP and SOPWAV methods considered are those that deliver the highest R_{OS}^2 (as reported in Table 1). In both plots, the black line is for the HM, while the blue and red lines are for the SOPWAV and SOP methods, respectively. Panel A presents the equity weights, which are constrained to lie in the range between -0.5 and 1.5. The first result that stands out is the substantially different dynamics of the equity exposure between

These findings indicate that the Ferreira and Santa-Clara (2011) assumptions regarding the first two parts of stock market returns (dp and ge) do not lead to a major statistical underperformance versus our baseline SOPWAV method. Conversely, the substantial statistical gains of the SOPWAV method versus the original SOP method are obtained when the third part of stock market returns (gm) is considered.

¹² For the OOS period under analysis, the SR of the trading strategy based on the HM forecast is 0.30.

the three portfolios. The SOPWAV method implies large and frequent changes in equity weights, whereas the equity exposure changes much more smoothly using the HM and the SOP method. Interestingly, during the two recessions in the 2000s, the equity exposure of the HM portfolio features an upward trend. In contrast, the equity exposure using the SOPWAV method is strongly reduced at the beginning of both recessions (including short selling). The second notable result from panel A is that the constraints on the weights are frequently binding for the SOPWAV portfolio, yet they are almost irrelevant for the SOP and HM portfolios.

Panel B in Figure 3 shows the corresponding log cumulative wealth for an investor trading with the HM, the SOP and the SOPWAV portfolios. In the simulation, we assume that the investor begins with \$1 and reinvests all proceeds. The wealth dynamics for an investor with the SOPWAV portfolio clearly beats the one with the SOP portfolio, which in turn outperforms the HM portfolio. The much higher rotation of the SOPWAV portfolio, as illustrated in panel A, reflects an enhanced market timing of this strategy. The divergence on the cumulative wealth between the SOPWAV portfolio and the other two portfolios is consistent throughout the entire sample period. It is interesting to note that, more recently during the 2000s, the HM-based portfolio is negatively affected by two severe drawdowns (the first occurring around the recession at the beginning of the 2000s and the second during the 2008-2009 global financial crisis). The cumulative wealth of the investor with a HM-based strategy, in fact, only returns to the level of the early 2000s in 2014. In contrast, an investor adopting a SOPWAV-based strategy quickly recovers from the (much smaller) drawdowns suffered during the 2000s recessions and benefits from a strong upward trend in her wealth until the end of the sample period.

5 Robustness tests

5.1 Sensitivity to level of investor risk aversion and variance forecast of stock market returns

In this sub-section, we evaluate the robustness of the economic performance of the SOPWAV method using (i) a different level of investor’s relative risk aversion (γ in (14)), and (ii) a different way of forecasting the variance of stock market returns when computing the optimal share allocated to equities ($\hat{\sigma}_{t+1}^2$ in (14)).

Results are reported in Table 3. First, we re-run the asset allocation analysis for an investor that has a higher level of relative risk aversion ($\gamma = 5$). For both the CER gains (columns 2 and 3) and the SR gains (columns 4 and 5), the SOPWAV method continues to outperform the HM benchmark and any version of the SOP method (baseline or extended). Second, following Rapach, Ringgenberg and Zhou (2016), we re-do the asset allocation analysis by forecasting the variance of stock market returns using a ten-year moving window of past returns. Results, reported in columns 6 to 9, show that the superior performance of the SOPWAV method is robust also in this case.

5.2 Alternative multiple in SOP method

As pointed out by Ferreira and Santa-Clara (2011), the SOP method offers flexibility as regards the choice of the variable to be included in equation (2) to decompose capital gains. For instance, instead of using the price-earnings multiple, other price multiples such as the price-dividend, the price-to-book or the price-to-sales may be used. Ferreira and Santa-Clara (2011) state that they obtain similar results using the price-earnings multiple or the price-dividend multiple. We believe it is worthwhile to compare our previous results (*i.e.* those obtained using the price-earnings multiple) with those using the price-dividend multiple for

the enlarged sample period, given the huge earnings swing between 2008 and 2009 (whereas dividends posted a rather smooth behavior during that period).¹³

Using the price-dividend multiple, equation (2) becomes

$$1 + CG_{t+1} = \frac{P_{t+1}}{P_t} = \frac{P_{t+1}/D_{t+1}}{P_t/D_t} \frac{D_{t+1}}{D_t} = \frac{M_{t+1}}{M_t} \frac{D_{t+1}}{D_t} = (1 + GMD_{t+1})(1 + GD_{t+1}) \quad , \quad (15)$$

where $M_{t+1} = P_{t+1}/D_{t+1}$ is the price-dividend multiple, GMD_{t+1} is the price-dividend multiple growth rate and GD_{t+1} is the dividend growth rate.

Accordingly, the log stock return is given by the sum of the dividend-price ratio, the growth in dividends (gd) and the growth in the price-dividend multiple (gmd): $r_{t+1} = dp_{t+1} + gd_{t+1} + gmd_{t+1}$. Table 4 shows the results obtained with the SOP and SOPWAV methods using the price-dividend multiple when decomposing stock market returns. The statistical performance, measured by the R_{OS}^2 , is documented in columns 2 and 3, while the economic performance (summarized by the CER and SR gains) is reported between columns 4 and 7. We emphasize two main results. First, both from a statistical and economic perspective, the superior performance of both methods (SOP and SOPWAV) versus the HM benchmark is robust toward the change in the multiple used. Second, the outperformance of the SOPWAV method versus the SOP continues to be evident both from a statistical point of view and from an asset allocation perspective, especially when looking at the CER gains.

5.3 Holdout analysis

Our analysis shows that it is crucial to isolate the relevant frequencies of the relevant parts to be able to enhance the OOS stock market return predictability. In line with similar analyses in the literature, the selection of the SOPWAV method specification is *ex post*, *i.e.* the choice

¹³ This large swing in earnings does not affect the results in Ferreira and Santa-Clara (2011), whose dataset ends in 2007.

of the weights ω_i in equation (13) is done by comparing the OOS performance of alternative specifications at the end of the out-of-sample period. However, in the real world the trader needs to forecast stock market returns in real time to properly rebalance her portfolio. In other words, the trader needs to know, in real time, which frequencies of which parts to use when forecasting stock market returns. It is useless for a trader to have to wait until the sample period ends to know which frequencies of which predictors work best.

Motivated by this reasoning, we run the following exercise using the extended SOPWAV method. We consider the same in-sample period as before (December 1927 to December 1949). As regards the out-of-sample period, we use the first q_0 observations as the holdout period, in which the weights in equation (13) are chosen to maximize the R_{OS}^2 during that period. We then compute the forecasts over the post-holdout out-of-sample period by using that particular choice of weights, leaving us with a total of $T - t_0 - q_0$ out-of-sample forecast errors. From a trader's perspective, the holdout period represents the period used to choose the frequencies of the parts that will be used in the subsequent real-time trading period.

We run this exercise considering four holdout periods (sequentially extended by five years): 1950–1983, 1950–1988 (that also includes the 1987 market crash), 1950–1993 and 1950–1998 (that also include the 1997 Asian crisis and the 1998 LTCM collapse). Results are reported in Table 5. There are five specifications of the extended SOPWAV method (when using SVAR, EP, RVOL, DP and BM to forecast gm) that always outperform the HM for the four holdout periods considered. Occasionally, other specifications also outperform the HM, but their performances are not robust for all holdout periods. In any case, this exercise shows that the SOPWAV method can effectively add value for a real world trader.

6 Concluding remarks

This paper proposed a new method for forecasting stock market returns (SOPWAV) by generalizing the Ferreira and Santa-Clara (2011) sum-of-the-parts (SOP) method of stock return forecasts. Specifically, we suggest a selective summing of the frequency-decomposed parts of stock market returns, instead of summing the original parts. We forecast stock market returns (S&P500 index) out-of-sample for the period running from January 1950 to December 2015. The SOPWAV method delivers statistically and economically significant gains for investors, clearly outperforming both the historical mean benchmark and the original SOP method. The results are robust to changes on the level of investor's relative risk aversion and on the variables used in the stock return decomposition. When tested from a real-time trading perspective, the results show that the SOPWAV method can also add value for traders.

Two factors account for the strong relative performance of the SOPWAV method. First, the use of wavelet decomposition makes it possible to extract and use only those frequencies in the parts that have the greatest predictive power. Second, the selected frequency-decomposed parts carry complementary information for the forecasting exercise that captures different frequencies of stock market returns. By only considering the relevant frequencies (those at which stock returns are predictable) of the relevant parts (those that have predictive power at relevant frequencies), the SOPWAV method well anticipates the future dynamics of stock market returns.

We foresee some interesting research avenues related with the proposed wavelet-based forecasting method. The granularity of analysis that this method allows is expected to be useful for forecasting other financial variables (such as volatility and correlation in equity markets or bond returns), as well as macroeconomic variables beyond GDP, inflation and investment (as already done by Rua, 2011, 2017 and Kilponen and Verona, 2016). Moreover, it would be

interesting to extend the Chaudhuri and Lo (2016) analysis of the optimal dynamic portfolio choice under the joint time-frequency domain by using wavelet methodologies, as they allow to study the implication of time-varying frequencies.

References

- Aguiar-Contraria, Luis, and Maria Joana Soares.** 2011. “Business cycle synchronization and the Euro: A wavelet analysis.” *Journal of Macroeconomics*, 33(3): 477–489.
- Aguiar-Contraria, Luis, and Maria Joana Soares.** 2014. “The Continuous Wavelet Transform: Moving Beyond Uni- And Bivariate Analysis.” *Journal of Economic Surveys*, 28(2): 344–375.
- Aguiar-Contraria, Luis, Manuel M.F. Martins, and Maria Joana Soares.** 2012. “The yield curve and the macro-economy across time and frequencies.” *Journal of Economic Dynamics and Control*, 36(12): 1950–1970.
- Ang, Andrew, and Geert Bekaert.** 2007. “Stock Return Predictability: Is it There?” *Review of Financial Studies*, 20(3): 651–707.
- Baetje, Fabian, and Lukas Menkhoff.** 2016. “Equity premium prediction: Are economic and technical indicators unstable?” *International Journal of Forecasting*, 32(4): 1193–1207.
- Bandi, Federico, Bernard Perron, Andrea Tamoni, and Claudio Tebaldi.** 2016. “The Scale of Predictability.” *Working Papers 509*.
- Bekaert, Geert, and Robert J. Hodrick.** 1992. “Characterizing Predictable Components in Excess Returns on Equity and Foreign Exchange Markets.” *The Journal of Finance*, 47(2): 467–509.

- Bollerslev, Tim, George Tauchen, and Hao Zhou.** 2009. “Expected Stock Returns and Variance Risk Premia.” *Review of Financial Studies*, 22(11): 4463–4492.
- Bollerslev, Tim, Viktor Todorov, and Lai Xu.** 2015. “Tail risk premia and return predictability.” *Journal of Financial Economics*, 118(1): 113–134.
- Campbell, John Y.** 1987. “Stock returns and the term structure.” *Journal of Financial Economics*, 18(2): 373–399.
- Campbell, John Y., and Samuel B. Thompson.** 2008. “Predicting Excess Stock Returns Out of Sample: Can Anything Beat the Historical Average?” *Review of Financial Studies*, 21(4): 1509–1531.
- Chaudhuri, S., and Andrew Lo.** 2016. “Spectral Portfolio Theory.” mimeo.
- Cochrane, John H.** 2008. “The Dog That Did Not Bark: A Defense of Return Predictability.” *Review of Financial Studies*, 21(4): 1533–1575.
- Cooper, Ilan, and Richard Priestley.** 2009. “Time-Varying Risk Premiums and the Output Gap.” *Review of Financial Studies*, 22(7): 2601–2633.
- Cooper, Ilan, and Richard Priestley.** 2013. “The World Business Cycle and Expected Returns.” *Review of Finance*, 17(3): 1029–1064.
- Crowley, Patrick M.** 2007. “A Guide To Wavelets For Economists.” *Journal of Economic Surveys*, 21(2): 207–267.
- Cutler, David M., James M. Poterba, and Lawrence H. Summers.** 1991. “Speculative Dynamics.” *Review of Economic Studies*, 58(3): 529–546.
- Dangl, Thomas, and Michael Halling.** 2012. “Predictive regressions with time-varying coefficients.” *Journal of Financial Economics*, 106(1): 157–181.

- Fama, Eugene F., and G. Schwert.** 1977. "Asset returns and inflation." *Journal of Financial Economics*, 5(2): 115–146.
- Ferreira, Miguel A., and Pedro Santa-Clara.** 2011. "Forecasting stock market returns: the sum of the parts is more than the whole." *Journal of Financial Economics*, 100(3): 514–537.
- Ferson, Wayne, and Campbell Harvey.** 1991. "The Variation of Economic Risk Premiums." *Journal of Political Economy*, 99(2): 385–415.
- Ferson, Wayne E., and Campbell R. Harvey.** 1993. "The Risk and Predictability of International Equity Returns." *Review of Financial Studies*, 6(3): 527–566.
- Galagedera, D., and E. Maharaj.** 2008. "Wavelet timescales and conditional relationship between higher-order systematic co-moments and portfolio returns." *Quantitative Finance*, 8(2): 201–215.
- Gallegati, Marco, and James 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.
- Gallegati, Marco, Mauro Gallegati, James B. Ramsey, and Willi Semmler.** 2011. "The US Wage Phillips Curve across Frequencies and over Time." *Oxford Bulletin of Economics and Statistics*, 73(4): 489–508.
- Gencay, Ramazan, and Daniele Signori.** 2015. "Multi-scale tests for serial correlation." *Journal of Econometrics*, 184(1): 62–80.
- Gencay, Ramazan, Faruk Selcuk, and Brandon Whitcher.** 2005. "Multiscale systematic risk." *Journal of International Money and Finance*, 24(1): 55–70.
- 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, Campbell R.** 1991. “The World Price of Covariance Risk.” *The Journal of Finance*, 46(1): 111–157.
- Harvey, Campbell R., Yan Liu, and Heqing Zhu.** 2016. “... and the Cross-Section of Expected Returns.” *Review of Financial Studies*, 29(1): 5–68.
- Hasbrouck, Joel.** 2017. “High Frequency Quoting: Short-Term Volatility in Bids and Offers.” *Journal of Financial and Quantitative Analysis*, forthcoming.
- Hjalmarsson, Erik.** 2010. “Predicting Global Stock Returns.” *Journal of Financial and Quantitative Analysis*, 45: 49–80.
- Hong, Yongmiao, and Chihwa Kao.** 2004. “Wavelet-Based Testing for Serial Correlation of Unknown Form in Panel Models.” *Econometrica*, 72(5): 1519–1563.
- Huang, Dashan, Fuwei Jiang, Jun Tu, and Guofu Zhou.** 2015. “Investor Sentiment Aligned: A Powerful Predictor of Stock Returns.” *Review of Financial Studies*, 28(3): 791–837.
- Kandel, Shmuel, and Robert F Stambaugh.** 1996. “On the Predictability of Stock Returns: An Asset-Allocation Perspective.” *Journal of Finance*, 51(2): 385–424.
- Kelly, Bryan, and Seth Pruitt.** 2013. “Market Expectations in the Cross-Section of Present Values.” *Journal of Finance*, 68(5): 1721–1756.
- Kilponen, Juha, and Fabio Verona.** 2016. “Testing the Q theory of investment in the frequency domain.” Bank of Finland Research Discussion Papers 32/2016.
- Kim, Sangbae, and Francis In.** 2005. “The relationship between stock returns and inflation: new evidence from wavelet analysis.” *Journal of Empirical Finance*, 12(3): 435–444.
- Lettau, Martin, and Sydney Ludvigson.** 2001. “Consumption, aggregate wealth, and expected stock returns.” *Journal of Finance*, 56(3): 815–849.

- Lima, Luiz Renato, and Fanning Meng.** 2017. “Out-of-Sample Return Predictability: A Quantile Combination Approach.” *Journal of Applied Econometrics*, 32(4): 877–895.
- Li, Yan, David T. Ng, and Bhaskaran Swaminathan.** 2013. “Predicting market returns using aggregate implied cost of capital.” *Journal of Financial Economics*, 110(2): 419–436.
- Ludvigson, Sydney C., and Serena Ng.** 2007. “The empirical risk-return relation: A factor analysis approach.” *Journal of Financial Economics*, 83(1): 171 – 222.
- Malagon, J., D. Moreno, and R. Rodriguez.** 2015. “Time horizon trading and the idiosyncratic risk puzzle.” *Quantitative Finance*, 15(2): 327–343.
- Manchaldore, J., I. Palit, and O. Soloviev.** 2010. “Wavelet decomposition for intra-day volume dynamics.” *Quantitative Finance*, 10(8): 917–930.
- McCracken, M.W.** 2007. “Asymptotics for out of sample tests of Granger causality.” *Journal of Econometrics*, 140(2): 719 – 752.
- Mele, Antonio.** 2007. “Asymmetric stock market volatility and the cyclical behavior of expected returns.” *Journal of Financial Economics*, 86(2): 446–478.
- Moller, Stig V., and Jesper Rangvid.** 2015. “End-of-the-year economic growth and time-varying expected returns.” *Journal of Financial Economics*, 115(1): 136 – 154.
- Neely, C., D. Rapach, J. Tu, and G. Zhou.** 2014. “Forecasting the Equity Risk Premium: The Role of Technical Indicators.” *Management Science*, 60(7): 1772–1791.
- Ortu, Fulvio, Andrea Tamoni, and Claudio Tebaldi.** 2013. “Long-Run Risk and the Persistence of Consumption Shocks.” *The Review of Financial Studies*, 26(11): 2876.
- Ortu, Fulvio, Federico Severino, Andrea Tamoni, and Claudio Tebaldi.** 2016. “A Persistence-based Wold-type Decomposition for Stationary Time Series.” *Working Papers*.

- Pastor, Lubos, and Robert Stambaugh.** 2009. "Predictive Systems: Living with Imperfect Predictors." *The Journal of Finance*, 64(4): 1583–1628.
- Percival, D.B., and A.T. Walden.** 2000. *Wavelet methods for time series analysis*. Cambridge University Press.
- Pettenuzzo, Davide, Allan Timmermann, and Rossen Valkanov.** 2014. "Forecasting stock returns under economic constraints." *Journal of Financial Economics*, 114(3): 517–553.
- Pettenuzzo, Davide, and Francesco Ravazzolo.** 2016. "Optimal Portfolio Choice Under Decision-Based Model Combinations." *Journal of Applied Econometrics*, 31(7): 1312–1332.
- Ramsey, James B., and Camille Lampart.** 1998*a*. "The Decomposition of Economic Relationships by Time Scale Using Wavelets: Expenditure and Income." *Studies in Nonlinear Dynamics & Econometrics*, 3(1): 1–22.
- Ramsey, James B., and Camille Lampart.** 1998*b*. "Decomposition Of Economic Relationships By Timescale Using Wavelets: Money and Income." *Macroeconomic Dynamics*, 2(1): 49–71.
- Rapach, David E., and Guofu Zhou.** 2013. "Forecasting Stock Returns." Vol. 2 of *Handbook of Economic Forecasting*, 328–383. Elsevier.
- Rapach, David E., Jack K. Strauss, and Guofu Zhou.** 2010. "Out-of-Sample Equity Premium Prediction: Combination Forecasts and Links to the Real Economy." *Review of Financial Studies*, 23(2): 821–862.
- Rapach, David E., Jack K. Strauss, and Guofu Zhou.** 2013. "International Stock Return Predictability: What Is the Role of the United States?" *Journal of Finance*, 68(4): 1633–1662.

- Rapach, David E., Matthew C. Ringgenberg, and Guofu Zhou.** 2016. “Short interest and aggregate stock returns.” *Journal of Financial Economics*, 121(1): 46 – 65.
- Renaud, O., J.L. Starck, and F. Murtagh.** 2002. “Wavelet-based Forecasting of Short and Long Memory Time Series.” *Research Papers by the Institute of Economics and Econometrics, Geneva School of Economics and Management, University of Geneva, Institut Economie et Econometrie, Universit de Geneve.*
- Rua, Antonio.** 2010. “Measuring comovement in the time-frequency space.” *Journal of Macroeconomics*, 32(2): 685–691.
- Rua, Antonio.** 2011. “A wavelet approach for factor-augmented forecasting.” *Journal of Forecasting*, 30(7): 666–678.
- Rua, Antonio.** 2017. “A wavelet-based multivariate multiscale approach for forecasting.” *International Journal of Forecasting*, 33(3): 581–590.
- Rua, Antonio, and Luis C. Nunes.** 2009. “International comovement of stock market returns: A wavelet analysis.” *Journal of Empirical Finance*, 16(4): 632–639.
- Xue, Yi, Ramazan Gencay, and Stephen Fagan.** 2013. “Jump detection with wavelets for high-frequency financial time series.” *Quantitative Finance*, 14(8): 1427–1444.

Model		SOP	SOPWAV ALL	SOPWAV	Frequencies of dp , ge and gm								
					D_1	D_2	D_3	D_4	D_5	D_6	D_7	D_8	
Baseline		0.91***	-95.50	0.97***	dp	1	1	0	0	1	1	0	1
					ge	0	1	0	0	0	0	0	1
Extended	SVAR	0.06	-83.92	1.42***	gm	0	0	0	0	0	1	1	0
	DFR	0.99***	-69.06	1.33***		1	1	0	1	0	0	0	0
	LTY	0.81***	-90.60	2.60***		0	0	0	1	0	0	0	0
	LTR	0.98***	-84.68	1.73***		0	0	0	0	1	0	0	0
	INFL	0.71***	-85.91	1.78***		0	0	0	0	1	1	1	0
	TMS	1.13***	-90.47	1.83***		0	0	0	0	1	0	0	0
	TBL	0.92***	-87.51	2.41***		0	0	1	0	1	0	0	0
	DFY	0.88***	-59.96	1.55***		1	0	0	0	1	1	1	0
	NTIS	0.82***	-85.09	1.50***		0	1	0	0	0	1	1	0
	ROE	0.85***	-120.20	2.30***		0	0	1	0	1	0	0	0
	DE	1.30***	-108.68	1.56***		0	0	0	0	0	1	1	0
	EP	1.22***	-292.80	1.30***		0	0	0	0	1	1	0	0
	RVOL	1.10***	-83.08	1.86***		1	1	1	0	1	1	0	0
	DP	0.59***	-95.30	1.22***		0	0	0	0	1	1	0	0
BM	0.36**	-113.72	1.25***		0	0	0	0	1	0	0	0	

Table 1: Out-of-sample R-squares (R_{OS}^2)

This table reports the out-of-sample R-squares (R_{OS}^2 , in percentages) for stock market return forecasts at monthly (non-overlapping) frequencies from the SOP, SOPWAV_ALL and SOPWAV methods (equations (9), (12) and (13) respectively). The three methods are considered in a specification where only dp and ge are used to forecast stock market returns (baseline versions) and in an extended version where gm is additionally considered and estimated using bivariate regressions on different predictors. It is also reported the frequency-specific components of dp , ge and gm that maximize the R_{OS}^2 for the full forecasting period (1950:01 to 2015:12) of the SOPWAV method (equation 13) in the baseline version and in each of the extended versions. The R_{OS}^2 measures the proportional reduction in the mean squared forecast error for the predictive method relative to the forecast based on the historical mean (HM). The one-month-ahead out-of-sample forecast of stock market return is generated using a sequence of expanding windows. The sample period is from 1927:12 to 2015:12. The out-of-sample forecasting period runs from 1950:01 to 2015:12. Asterisks denote significance of the out-of-sample $MSFE-F$ statistic of McCracken (2007). *** and ** denote significance at the 1% and 5% level, respectively.

Model		CER gains		SR gains	
		SOP	SOPWAV	SOP	SOPWAV
Baseline		1.61	3.01	0.18	0.28
Extended	SVAR	0.87	2.85	0.06	0.25
	DFR	1.90	3.10	0.21	0.22
	LTY	1.62	5.58	0.18	0.43
	LTR	1.60	4.13	0.13	0.28
	INFL	1.49	3.98	0.14	0.33
	TMS	1.93	4.06	0.20	0.33
	TBL	1.70	4.65	0.21	0.36
	DFY	1.35	2.82	0.13	0.23
	NTIS	1.51	3.37	0.11	0.29
	ROE	1.61	3.67	0.16	0.35
	DE	2.06	2.91	0.23	0.30
	EP	1.87	3.27	0.21	0.31
	RVOL	1.86	4.14	0.24	0.29
	DP	0.93	3.10	0.18	0.29
BM	0.51	3.37	0.03	0.29	

Table 2: Out-of-sample CER and Sharpe ratios gains

The second and third columns report the annualized certainty equivalent return (CER) percentage gains for an investor allocating her wealth between equities and risk-free bills according to the rule (14), using stock return forecasts from the SOP method (equation 9) and from the SOPWAV method (equation 13) using frequency components listed in Table 1. Both methods are considered in a specification where only dp and ge are used to forecast stock market returns (baseline versions), as well as in an extended version where gm is additionally considered and estimated using bivariate regressions on different predictors. Columns 4 and 5 give the annualized Sharpe ratios (SR) gains for an investor using the SOP and SOPWAV method, respectively, versus the SR obtained when using forecasts based on the HM. The investor is assumed to have a relative risk aversion coefficient of two and the equity weight in the portfolio is constrained to a range between -0.5 and 1.5. The out-of-sample forecasting period runs from 1950:01 to 2015:12.

Model	$\gamma = 5$				$\hat{\sigma}^2 : 10y MA$				
	CER gains		SR gains		CER gains		SR gains		
	SOP	SOPWAV	SOP	SOPWAV	SOP	SOPWAV	SOP	SOPWAV	
Baseline	0.59	1.23	0.16	0.24	1.43	2.52	0.09	0.15	
Extended	SVAR	-0.08	1.28	-0.07	0.24	1.26	2.31	0.08	0.14
	DFR	0.65	1.43	0.17	0.18	2.05	1.98	0.14	0.11
	LTY	0.59	2.97	0.16	0.40	1.56	6.40	0.10	0.42
	LTR	0.58	2.07	0.09	0.25	1.56	3.26	0.09	0.20
	INFL	0.53	1.72	0.11	0.31	1.68	4.09	0.11	0.26
	TMS	0.72	1.67	0.19	0.29	1.90	4.13	0.11	0.26
	TBL	0.63	2.12	0.20	0.34	1.86	4.29	0.13	0.26
	DFY	0.46	1.23	0.11	0.19	1.79	2.31	0.11	0.13
	NTIS	0.56	1.40	0.08	0.25	1.75	2.80	0.10	0.16
	ROE	0.59	1.74	0.15	0.33	1.16	2.93	0.07	0.18
	DE	0.85	1.40	0.21	0.28	2.00	2.93	0.13	0.19
	EP	0.69	1.44	0.19	0.28	1.54	2.87	0.10	0.18
	RVOL	0.68	1.73	0.23	0.22	2.24	2.66	0.16	0.15
	DP	0.29	1.28	0.15	0.25	1.74	2.84	0.17	0.17
BM	0.11	1.38	0.00	0.25	-0.19	2.68	-0.02	0.15	

Table 3: Out-of-sample CER and Sharpe ratios gains

This table reports the annualized certainty equivalent return (CER) percentage gains and the annualized Sharpe ratios (SR) gains for an investor allocating her wealth between equities and risk-free bills according to the rule (14), using stock return forecasts from the SOP method (equation 9) and from the SOPWAV method (equation 13) using frequency components listed in Table 1. Both methods are considered in a specification where only dp and ge are used to forecast stock market returns (baseline versions), as well as in an extended version where gm is additionally considered and estimated using bivariate regressions on different predictors. In columns 2 to 5 are reported results assuming an investor with a relative risk aversion coefficient of five. Columns 6 to 9 show the reported results for variance of stock market returns in (14) forecasted using a ten-year window of past returns. The equity weight in the portfolio is constrained to a range between -0.5 and 1.5. The out-of-sample forecasting period runs from 1950:01 to 2015:12.

Model	R_{OS}^2		CER gains		SR gains		
	SOP	SOPWAV	SOP	SOPWAV	SOP	SOPWAV	
Baseline	0.78***	0.97***	1.32	1.46	0.16	0.15	
Extended	SVAR	0.16**	1.82***	0.73	2.28	0.07	0.25
	DFR	0.79***	1.24***	1.52	2.45	0.20	0.19
	LTY	0.60***	2.37***	1.16	3.75	0.14	0.29
	LTR	0.89***	2.28***	1.55	3.22	0.15	0.23
	INFL	0.66***	1.31***	1.30	2.12	0.14	0.19
	TMS	0.97***	1.03***	1.64	1.40	0.13	0.07
	TBL	0.74***	1.24***	1.36	1.90	0.16	0.16
	DFY	0.60***	1.14***	1.05	1.72	0.12	0.13
	NTIS	0.83***	1.48***	1.65	2.47	0.12	0.25
	ROE	0.26**	1.69***	0.60	2.81	0.06	0.27
	DE	0.50***	1.01***	0.82	1.60	0.10	0.14
	EP	0.90***	0.99***	1.33	1.39	0.24	0.14
	RVOL	0.85***	1.46***	1.42	2.10	0.19	0.15
	DP	0.60***	1.07***	0.90	1.54	0.22	0.16
	BM	0.20**	1.34***	0.14	1.91	0.00	0.15

Table 4: Out-of-sample R-squares (R_{OS}^2), CER and Sharpe ratios gains
This table reports the out-of-sample R-squares (R_{OS}^2 , in percentages), the annualized certainty equivalent return (CER) percentage gains and the annualized Sharpe ratios (SR) gains for an investor allocating her wealth between equities and risk-free bills according to the rule (14), using stock return forecasts from the SOP and from the SOPWAV methods using the price-dividend multiple in the capital gain decomposition (equation 15). The R_{OS}^2 measures the proportional reduction in the mean squared forecast error for the predictive method relative to the forecast based on the historical mean (HM). Both methods are considered in a specification where only dp and gd are used to forecast stock market returns (baseline versions), as well as in an extended version where gmd is additionally considered and estimated using bivariate regressions on different predictors. The investor is assumed to have a relative risk aversion coefficient of two and the equity weight in the portfolio is constrained to lie between -0.5 and 1.5. The out-of-sample forecasting period runs from 1950:01 to 2015:12. Asterisks denote significance of the out-of-sample $MSFE-F$ statistic of McCracken (2007). *** and ** denote significance at the 1% and 5% level, respectively.

Initial in-sample period	1927-1949	1927-1949	1927-1949	1927-1949
Holdout out-of-sample period	1950-1983	1950-1988	1950-1993	1950-1998
Forecasting out-of-sample period	1984-2015	1989-2015	1994-2015	1999-2015
SVAR	0.66**	0.77**	1.29**	0.96**
DFR	-3.41	-4.23	-4.13	-4.91
LTY	1.07***	0.85**	-1.67	-2.21
LTR	-15.46	-0.78	-0.87	-1.55
INFL	0.51*	0.89**	-2.05	-3.58
TMS	-0.02	0.59*	-1.76	-2.46
TBL	0.64**	0.46*	-2.05	-3.04
DFY	-4.03	-5.35	-5.85	-8.27
NTIS	-1.00	-0.62	-0.12	-1.84
ROE	-1.07	-0.57	-2.06	-1.13
DE	-11.83	-15.43	-19.44	-16.77
EP	1.09***	1.57***	2.01***	1.97***
RVOL	0.83**	0.97**	2.19***	1.61***
DP	0.89**	1.40***	1.79***	1.65***
BM	0.66**	1.42***	1.77***	1.81***

Table 5: Out-of-sample R-squares (R_{OS}^2)

This table reports the out-of-sample R-squares (R_{OS}^2 , in percentages) for stock market return forecasts at monthly (non-overlapping) frequencies from the extended SOPWAV method (equation 13). The R_{OS}^2 measures the proportional reduction in the mean squared forecast error for the predictive method relative to the forecast based on the historical mean (HM). The one-month-ahead out-of-sample forecast of stock market return is generated using a sequence of expanding windows. The sample period is from 1927:12 to 2015:12. There are four different “holdout” out-of-sample periods in order to choose the weights in equation (13), with the respective subsequent forecasting period. Asterisks denote significance of the out-of-sample $MSFE-F$ statistic of McCracken (2007). ***, ** and * denote significance at the 1 %, 5 %, and 10 % level, respectively.

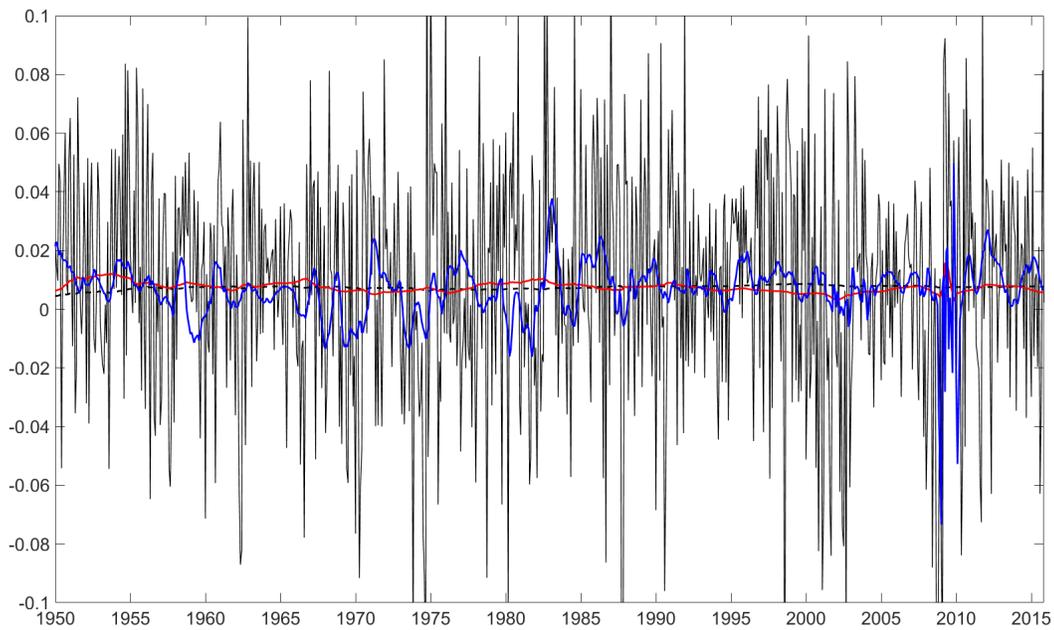


Figure 1: Realized and predicted stock market returns

The black solid line corresponds to the (log) realized stock market return as proxied by the (log) S&P 500 index return. The blue and red lines plot the one-month ahead out-of-sample predictive regression forecast for the stock market return based on the SOPWAV and SOP methods, respectively. The black dashed line represents the one-month-ahead out-of-sample stock market return forecast based on the historical mean (HM) of returns. The sample period runs from 1950:01 to 2015:12, monthly frequency.

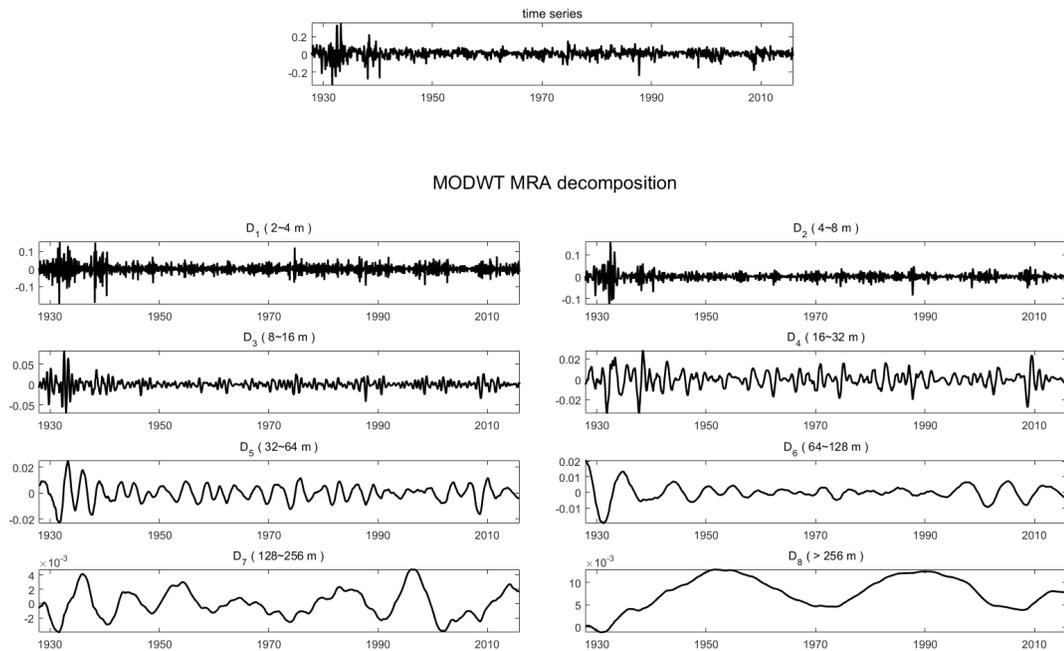
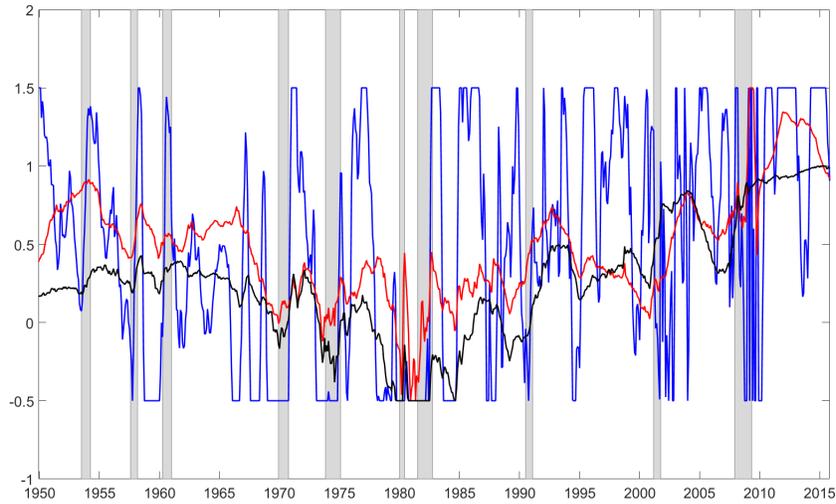


Figure 2: Stock market return, time series and MODWT MRA decomposition
The time series of the (log) stock market return as proxied by the S&P 500 index returns is presented in the top graph. The eight frequency-specific components (D_i , $i = 1, \dots, 8$) in which the stock market return time series is decomposed, using the $J = 7$ levels wavelet decomposition, are presented in the remaining graphs. The sample period runs from 1927:12 to 2015:12, monthly frequency.

Panel A. Equity weights



Panel B. Log cumulative wealth

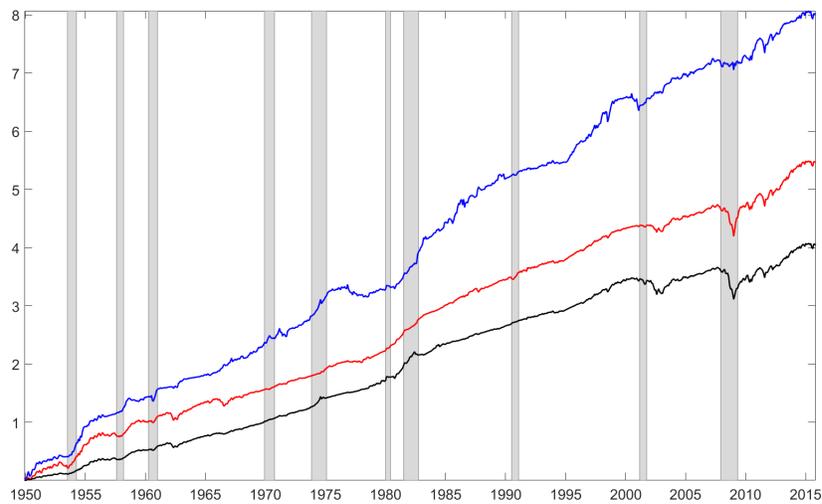


Figure 3: Equity weights and log cumulative wealth

Panel A plots the dynamics of the equity weight for a mean-variance investor allocating her wealth each month between equities and risk-free bills according to the rule (14), using stock return forecasts based on the SOPWAV and SOP methods, and the HM benchmark. The equity weight is constrained to lie between -0.5 and 1.5. The blue and red lines represent the equity weight for the investor using the SOPWAV and SOP method specification with highest out-of-sample R-squares (as reported in Table 1), respectively. The black line represents the equity weight for the investor using the HM benchmark. The Panel B delineates the corresponding log cumulative wealth for the investor, assuming she begins with \$1 and reinvests all proceeds. The blue, red, and black lines represent the log cumulative wealth for the investor using the SOPWAV, SOP, and HM methods, respectively. Gray bars denote NBER-dated recessions. The investor is assumed to have a relative risk aversion coefficient of two. The sample period runs from 1950:01 to 2015:12, monthly frequency.

Appendix 1. Variables used in the analysis

1. Stock variance (SVAR): sum of squared daily stock market returns on the S&P 500
2. Default return spread (DFR): difference between long-term corporate bond and long-term bond returns
3. Long-term yield (LTY): long-term government bond yield
4. Long-term return (LTR): long-term government bond return
5. Inflation rate (INFL): growth in the consumer price index with a one-month lag
6. Term spread (TMS): difference between the long-term government bond yield and the T-bill
7. Treasury bill rate (TBL): three-month Treasury bill rate
8. Default yield spread (DFY): difference between BAA- and AAA-rated corporate bond yields
9. Net equity expansion (NTIS): ratio of 12-month moving sums of net issues by NYSE-listed stocks to NYSE market capitalization.
10. Return on equity (ROE): ratio of 12-month moving sums of earnings to book value of equity for the S&P 500
11. Dividend-payout ratio (DE): difference between the log of dividends (12-month moving sums of dividends paid on the S&P 500) and the log of earnings (12-month moving sums of earnings on the S&P 500)
12. Earnings-price ratio (EP): difference between the log of earnings (12-month moving sums of earnings on the S&P 500) and the log of prices (S&P 500 index price).

13. Excess stock return volatility (RVOL): computed using a 12-month moving standard deviation estimator, as in Mele (2007)
14. 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 price)
15. Book-to-market (BM): ratio of book value to market value for the Dow Jones Industrial Average