UNIVERSIDADE CATÓLICA PORTUGUESA

Stock Return Predictability and Variance Risk Premia

A frequency domain analysis

Pedro Teixeira Magalhães

Católica Porto Business School
March 2021

UNIVERSIDADE CATÓLICA PORTUGUESA

Stock Return Predictability and Variance Risk Premia

A frequency domain analysis

Final Assignment in the modality of Dissertation presented to Universidade Católica Portuguesa to obtain the degree of Master in Finance

by

Pedro Teixeira Magalhães

under the supervision of

(PhD) Gonçalo Faria and (PhD) Fábio Verona

Católica Porto Business School
March 2021

Acknowledgments

To whom to give thanks—perhaps the most peaceful part of the entire thesis.

Although the process of writing a master's degree's final work can be somewhat lonely in practice, its completion would not be possible without the support of a few people and so, here it goes:

To my parents, Rosa and Rui, and my grandmother Ana- their unconditional support and affection allowed me to achieve one of my biggest dreams;

To Gabriel, Miguel, Rafael and Ricardo- the concern within my group of unconcerned friends and their often used phrases like, "Just get it over with!" made actually finishing this inevitable;

To my housemates, Tiago, Bárbara and João- a big thank you for their help and for sharing great moments since the very beginning of this adventure;

To Diana- for her love and daily motivation from the beginning until the end of this battle;

To my Professors, colleagues and friends at UCP's Católica Porto Business School- for sharing their knowledge and guiding me through and through again during this Master's degree;

To my Professor Gonçalo Faria and Professor Fábio Verona, advisor and co-advisor respectively- for sharing ideas, critical analysis, rigor and reiterated transmission of knowledge, but fundamentally for having reinforced the passion I have for the financial world.

And last but definitely not least, a special thanks to my grandfather Filinto- for being a pillar throughout my life.

Thank you.

Resumo

O objetivo principal desta tese é analisar o poder de previsão sobre o excesso de retorno de capital *out-of-sample* do prémio de risco de variância e dos seus domínios de frequência. O prémio de risco de variância é representado através da diferença entre as expectativas de risco neutro (variância implícita) e físicas (variância realizada) da variância.

Na literatura um número considerável de variáveis apresenta uma forte performance *in-* e *out-of-sample*, sendo mesmo uma delas o prémio de risco de variância. Da mesma maneira, a decomposição em frequências de certas variáveis beneficia de um aumento do poder preditivo *out-of-sample* das mesmas.

Sendo assim, de forma a estudar o comportamento desta variável iremos decompor a série temporal em frequências.

O resultado principal desta tese é que a série original em domínio temporal e a sua componente de média frequência demonstra uma notável performance *out-of-sample* a prever o excesso de retorno de capital. Mostramos também que, apesar de a série original em domínio temporal apresentar uma melhor performance estatística (ex., maior valor de R^2 out-of-sample), em termos económicos a sua componente de média frequência oferece maiores ganhos.

Palavras-chave: Excesso de Retorno de Capital, Prémio de Risco de Variância, Previsibilidade, Domínio de Frequência

Abstract

The main objective of this thesis is to analyze the out-of-sample equity return forecasting power of the variance risk premia and its frequency components. The variance risk premia (VRP) is represented by the difference between the risk neutral (implied variance) and physical (realized variance) expectations of the variance.

In the literature, a considerable number of variables present strong in- and out-of-sample performances, being one of them the variance risk premia. Likewise, by decomposing some variables into their frequencies, their out-of-sample performances increases.

Therefore, in order to study the behavior of this variable, we decompose the time series of the variance risk premia into frequencies.

The main result of this thesis is that the original time series and its medium frequency component demonstrate a remarkable out-of-sample performance when predicting the equity excess of return. We also show that, although the time series presents a better statistical performance (i.e., a higher out-of-sample R^2), in economic terms its medium frequency component delivers higher gains.

Keywords: Equity Excess Return, Variance Risk Premia, Predictability, Frequency Domain

Index/Contents

List of Figures

List of Tables

Introduction

Knowing and understanding the dynamics of the financial markets is a crucial part of an investor's task. Therefore, it is necessary to constantly update and develop models capable of enhancing the information that is available to financial market participants, in order to construct and rebalance a successful portfolio. This is where important questions arise: What variables can predict the excess market return? And if so, is this information economically relevant?

Extensive literature has been discussing these questions, and different answers have emerged in the literature. Different variables such as consumption wealth ratio, dividend-price ratio, book-to-value ratio (e.g. Rapach and Zhou (2013)), present statistically significant values regarding the forecasting performance of market excess returns. A variable that has recently been the subject of study is the variance risk premia. As stated by Bakshi and Kapadia (2003), "a representative agent investing in an asset with unknown future variance demands a compensation for the difference between the real-world and risk-neutral expectations of the future volatility, termed Variance Risk Premia". Bollerslev et al. (2009) and Drechsler and Yaron (2011) show that the variance risk premia is a strong predictor of future aggregate market returns. Later Pyun (2018), Fan et al. (2018), Prokopczuk et al. (2017), among others, decompose the variance risk premium using different approaches and show that are several components that exhibit their own unique empirical features and distinct economic interpretations. Consequently, they show that these decompositions can deliver high in-sample and out-of-sample predictive performance, improving remarkably the return predictability of the variance risk premia. Hence, this predictor presents statistically and economic significant performance in- and out-of-sample.

Bandi et al. (2018) and Faria and Verona (2018), show that it is possible to improve accuracy of predictability exercise by using a frequency analysis approach. Motivated by these studies, our research question arises: Can the out-of-sample predictability of the US equity returns from the variance risk premium be improved by considering frequency domain techniques?

In this thesis we analyze whether a frequency domain analysis enhances the predictive power of the variance risk premia. Particularly, following a similar procedure as Faria and Verona (2018), we investigate if there is evidence that a frequency decomposition provides a better outcome than the original time series of the variance risk premia. This study thus complements previous research and can provide interest and helpful conclusions.

Our main finding is that the medium frequency component of the variance risk premium has a statistically and economic significant predictive power of the equity excess return.

The dissertation is divided as follows. Chapter 1 provides an insight of the related literature. Chapter 2 presents the data and the method used to construct the predictors. Lastly, Chapter 3 presents the main results of in-sample predictability and out-of-sample forecasting, as well as an economic analysis of the results and a sensitivity analysis.

Chapter 1

Literature Review

A vast literature studies the predictive power of several variables to forecast the equity excess return. Although many efforts were made to try to find one predictor and model that better fits in this context, this theme is still unsealed. We frame our research within six different topics in the literature: the variance risk premia, the variance risk premia as excess return predictor, variance risk premia decompositions and in-sample predictability, its out-of-sample forecasting performance, and a frequency domain analysis.

1.1 The variance risk premia

Drechsler and Yaron (2011) state that volatility plays a crucial role in determining asset valuations. In fact, volatility measures are helpful tools that allows the investor to capture how perceptions of uncertainty about economic fundamentals are manifested in prices. In other hand, volatility plays a prominent role in derivatives markets, therefore, important information about the connections between uncertainty, the dynamics of the economy, preferences, and prices arises in these markets. We focus on the derivatives-related quantity called the variance risk premia (VRP). As Bollerslev et al. (2014) state, the variance risk premia, defined as "the difference between the risk-neutral and statistical expectations of the future return variation", can be interpreted as a measure of both aggregate economic uncertainty and aggregate risk aversion. Then, it´s relevant to study the dynamics between this variable and the future stock returns.

1.2 The variance risk premia as excess return predictor

Many financial players, like investment managers, researchers and regulators, focus and work on predictive models that study the forecasting power of a large number of financial variables. Indeed, financial markets, provide crucial information regarding economic uncertainty perceptions and cash-flow risk and how they manifest themselves in asset prices. Therefore, it is understandable that many studies try to find the best variable and model using this available information to predict stock returns. In fact, the literature proves that there are a lot of variables able to forecast future market return.

One of those variables is the variance risk premia (VRP), defined as the difference between the variance under risk-neutral probability and that under the physical probability, the one that is the object of this research. Introducing the risk factor, Merton (1973) stated that market risks play a prominent role in modern financial theory and are closely related to future stock return. Several studies (e.g.., Bakshi and Madan (2006); Bollerselev et al. (2009); Rosenberg and Engle (2002)) state that the variance risk premia, reflects investors' risk aversion, connecting then the variance risk premia with future stock return. To reinforce this relationship, Carr and Wu (2009) and Bakshi et al. (2010) show that VRP has a non-trivial dynamic over time, and Bollerslev et al. (2009), Driessen et al. (2009, 2013) and Drechsler and Yaron (2011) through a recursive utility framework, argue that this variable captures the time-varying economic uncertainty, affecting then the variation in the market risk premium.

Bollerslev et al. (2009) and Drechsler and Yaron (2011) investigate the variance risk premium for the US stock market and show that there is strong return predictability at monthly, quarterly and annual horizons. Extending these findings, Bollerslev et al. (2014) and Londono (2014) study the predictive power of variance risk premia to the international stock markets context, and present similar evidence as for the US market. Londono and Zhou (2017) further extends this investigation and build a variance risk premia measure from currency options, showing that the currency variance risk premia is also a determinant of the cross-section of stock returns. Recently Chen et al. (2016) investigate the relationships among variance risk premia, liquidity and stock returns and show that the variance risk premia can Granger-cause stock returns, concluding that VRP is a useful predictive variable for stock returns.

1.3 Variance risk premia decompositions and in-sample predictability

To develop an in-depth investigation on the predictability performance of the variance risk premia, many studies decompose it using several different methods, in order to focus on the VRP component that they believe contains more accurate information and therefore a better predictive power. Todorov (2010) analyzes the pricing of two components of variance risk due to stochastic volatility and jumps in a semiparametric framework, while Du and Kapadia (2012), Bollerslev et al. (2015) and Guo et al. (2017) show that various jump and tail risk measures have return predictability at medium-term horizons beyond the information content of VRP. Bekaert and Hoerova (2014) use a different approach and decompose the VIX index (proxy for risk-neutral variance expectation) into two parts, one that predicts market returns and the other that proxies financial instability, and show that the variance risk premia is a significant predictor of stock returns, but the conditional variance mostly is not. Bollerslev et al. (2015) provide an indirect decomposition of the conventional variance risk premia by extracting the component associated with fear. Bormetti et al. (2016) assess that the inclusion of a pricing kernel depending on both equity and variance risk premia that provides accurate pricing for S&P 500 Index options and proposes a new methodology for the VRP term structure that presents a strong predictive power on the future stock-index returns. As stated by Bormetti et al. (2016): "Our findings reveal that stock investors should not only follow the popular Wall Street's wisdom which suggests to look at the level of VIX Index but they also should look at the level of the term spread of variance risk premium". Several other studies[1] provide evidence that the decomposition of the variance risk premia enhances the predictive power of the variable.

1.4 Out-of-sample forecasting

Previous mentioned literature provided evidence that the variance risk premia presents remarkable in-sample predictability, but our main interest is on its out-of-sample forecasting performance. Several authors as Lettau and Ludvingson (2001), Campbell and Thompson (2008), and Rapach et al. (2010) show that by implementing the out-of-sample approach to test the predictive performance of different variables[2], strong evidence emerges that in fact they are capable to forecast excess stock returns. Regarding our research, recent literature supports that the variance risk premia can be used as a powerful out-of-sample predictor. Pyun (2018) studies how the market's exposure to variance risk is related to the

[1] Xing et al. (2010); Cremers and Weibnaum (2010); Cosemans (2012); Rehman and Vilkov (2012); Stilger et al. (2017).

[2] Predictors such as dividend-to-earnings ratio, consumptions-to-wealth, book-to-market ratios and term spread were used by these authors.

time-varying predictability of market returns by the VRP and proposes a model that uses the beta of contemporaneous regression of market return in changes in its variance as the predictive slope for the VRP[3]. This model can predict in a statistically and economically significant manner the one-month market return, exhibiting remarkable out-of-sample results. Fan et al. (2018) provide a model-free internally consistent decomposition of the total variance risk premia into the pure variance risk premia, and a higher order risk premia. Both these components exhibit unique empirical features and distinct economic interpretations. Overall, these decompositions jointly deliver high out-of-sample predictive performance, improving remarkably the return predictability of the total variance risk premia. Prokopczuk et al. (2017) find that the correlation risk premia proposed by Driessen et al. (2009, 2013) emerges as a strong predictor of both the market return and the realized variance, proving also that the variance risk premia predicts the market excess returns out-of-sample, exhibiting statistically significant forecasting power but also adding economic value. In this thesis a similar procedure for in- and out-of-sample predictability, as well as the economic approach, is used.

Fitting our results in the literature, Pyun (2018) proposes a methodology that generates an out-of-sample R^2 that ranges from 6% to 8%, Fan et al. achieved values of OOS R^2 that ranges from 3% to 5% and Prokopczuk et al. (2017) reported out-of-sample R^2 values of 5,50% (1-month horizon). In line with some of this findings this thesis achieves a remarkable out-of-sample R^2 for the original time series of the variance risk premium, as well as for its medium frequency component (7,077% and 6,772% respectively).

[3] Pyun (2018) shows that the slope that determines the contemporaneous relationship between market and variance risk resembles the relationship between the risk premium of the market and the market variance.

1.5 Frequency domain analysis

The frequency domain analysis provides a complementary tool for time domain analysis. In fact, frequency domain analysis surpasses a problematic issue related to the difficulty to observe occurrences of a given variable in just one frequency. This tool allows the researcher to access several different frequencies, providing then a more specific and complete analysis that contributes to the study of the significance of the different frequency levels on the behavior of the variable.

Several studies apply frequency domain analysis to correctly process all the information contained in the time series of a variable (e.g., Dew-Becker and Giglio (2016)). More recently Faria and Verona (2018) and Bandi et al. (2019) "use models where returns and predictors are linear aggregates of components operating over different frequencies, and where predictability is frequency-specific". Following this approach, and as a solid foundation to present the method used in our research, Faria and Verona (2020), focus on the out-of-sample forecasting performance of three economically motivated frequency components of the term spread of interest rates (TMS): the high-frequency component, the business cycle-frequency component, and the low-frequency component.

We analyze if there is in fact, like the TMS, a predictive power, using frequency domain techniques, of a different return predictor: the variance risk premia.

1.6 Wavelet filtering methods

Wavelet filtering methods have long been used in diverse study areas, including for finance applications. In fact, several studies use wavelet methods to forecast (out-of-sample) economic and financial time-series, such as Kilponen and Verona (2016) who forecast aggregate investment using the Tobin's q theory of investment; and Zhang et al. (2017) and Faria and Verona (2018, 2020), that focus on stock market returns predictability. The wavelet filtering method allows the decomposition of any variable, regardless of its time series properties, into different frequency bands, allowing the researcher to access information that the original time series of the variable may not provide.

Wavelet filtering method will then be explained and summarized in the method section.

Chapter 2

Data and Methodology

2.1 Data

We use monthly data from January 1990 to December 2019 to focus on the predictability of the index excess return[4] and base our empirical analysis on the aggregate S&P 500 composite index as a proxy for the aggregate market portfolio Data for VIX is from the Chicago Board of Options Exchange (CBOE), for the S&P 500 index total return from CRSP and for the 3-month treasury bill time series from the New York Federal Reserve Bank's website. In the following subsections we describe how our relevant variables are computed.

Summary statistics and correlations can be found in the Appendix section.

2.1.1 Variance Risk Premium (VRP)

Following the literature (eg. Bollerslev et al. (2009), Dreschler and Yaron (2011)), the variance risk premia (VRP) is defined as the difference between risk-neutral and physical expectations of variance:

$$VRP_t = E_t^{\mathbb{Q}}(\sigma_{t+1}^2) - E_t^{\mathbb{P}}(\sigma_{t+1}^2) \tag{1}$$

[4] Our sample is restricted due to the data available regarding the VIX index from CBOE. Also, we consider a period that ends in December 2019 due to the fact that there's a severe outlier in March 2020, probably due to covid-19 impact on the stock market.

where E(.) represents the expectation operator conditional on the information available at the moment, $\mathbb{Q}$ and $\mathbb{P}$ represent the risk-neutral and physical probability measures, respectively.

In order to proxy for the risk-neutral expectation of variance, following previous studies, we use the end-of-month t values of the VIX index to represent the expectation for month t + 1. This index is based on the highly liquid S&P 500 index options and is constructed through the "model-free" implied volatility approach, therefore this measure explicitly tailored replicates the risk-neutral variance of a fixed 30-day maturity, which is observable at time t. For the purpose of predictive analysis, we use the last trading day of each month observation, reflecting most recent available information.

Our physical expectation of variance is calculated by summing the daily realized variance of the S&P 500 index for each month. The daily realized variance is measured by the squared (absolute) daily index returns, where the index return is defined as the following formula: $ir_t = \left(\frac{i_t}{i_{t-1}}\right)$, where i_t represents the index return in day t. Following the risk-neutral variance procedure, the realized variance of month t represents the expectation for month t + 1, in line with Prokopczuk et al. (2017). The resulting time series are plotted in Figure 1, Figure 2 and Figure 3.

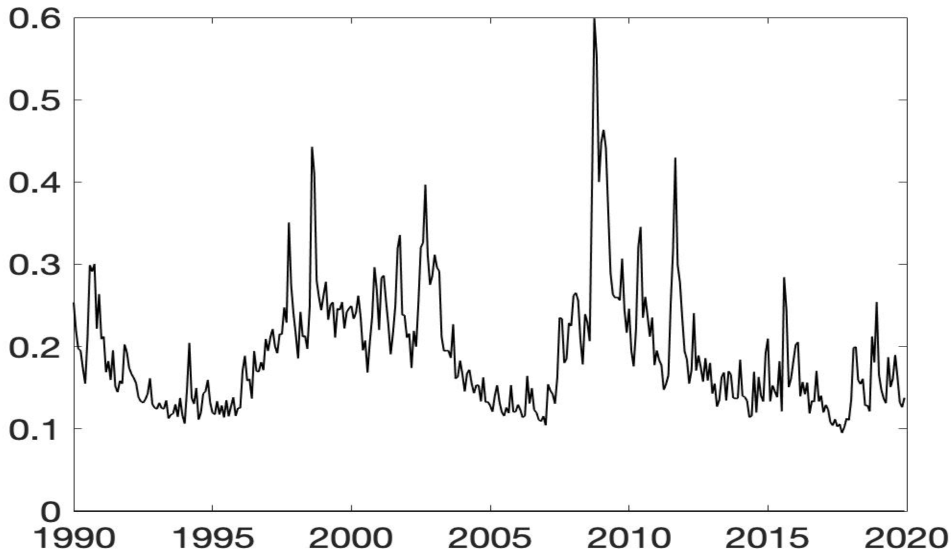

Figure 1: Time series of the implied variance.

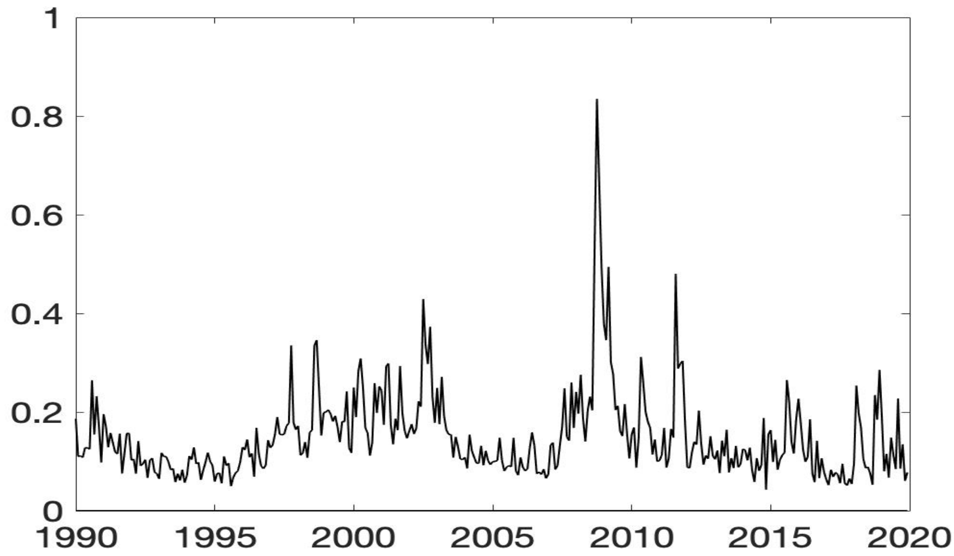

Figure 2: Time series of the realized variance.

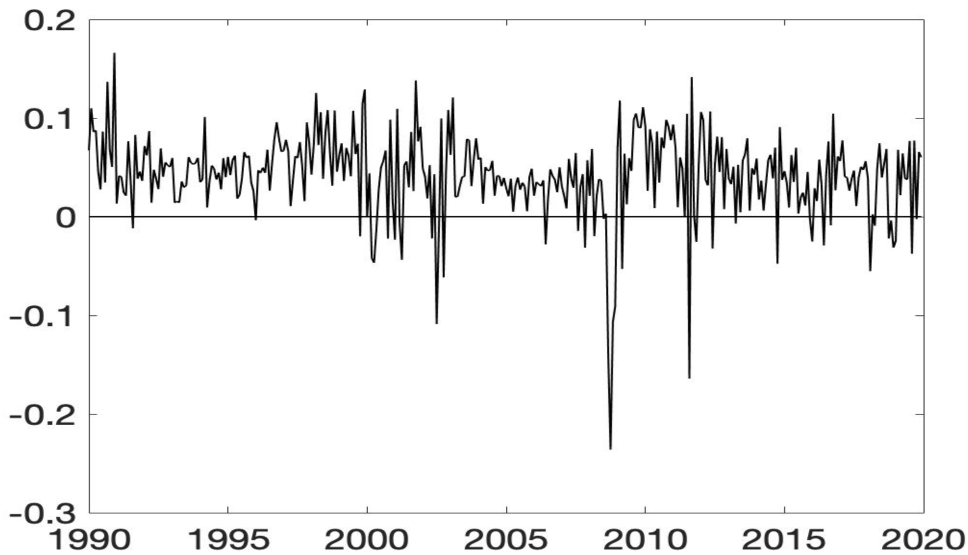

Figure 3: Time series of the variance risk premia.

2.1.2 Index Excess Return (r)

We compute the annualized monthly excess return on the S&P 500 (r_t) index by subtracting the annualized 3-month treasury bill (riskless rate) for the corresponding period from the total return on the equity index:

$$r_t = ir_t - \left((1 + Rf_t)^{1/12} - 1\right) \tag{2}$$

where ir_t represents the total return on the S&P 500 at month t and Rf_t represents the 3-month treasury bill at the end month t. The resulting time series is plotted in Figure 4.

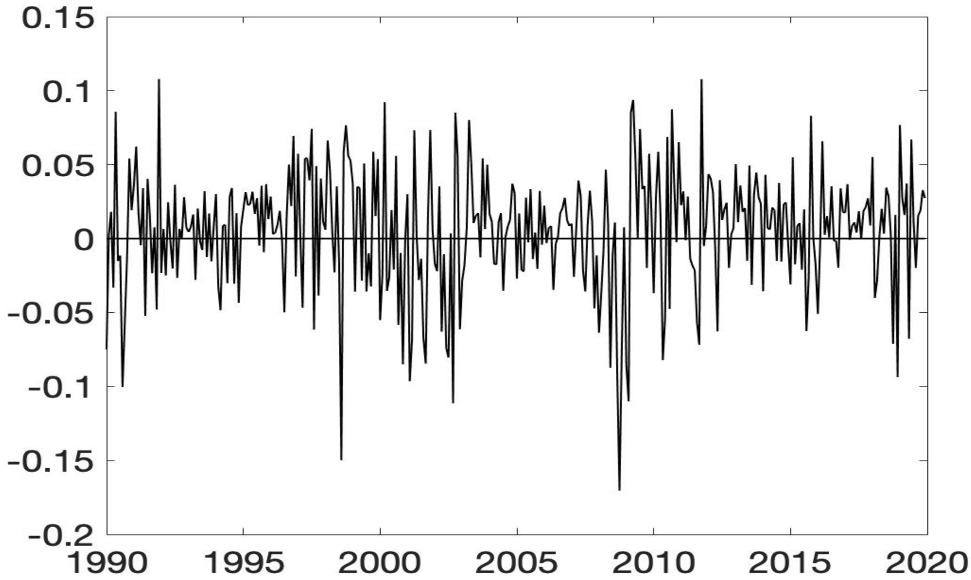

Figure 4: Time series of the equity excess return.

2.2 Method

This study aims to analyze the OOS forecasting performance of the variance risk premia original time series and its frequency components. To do so, we follow Faria and Verona (2018) method and decompose the variance risk premia using a J=5 level Maximum Overlap Discrete Wavelet Transform Multiresolution Analysis (MODWT MRA), with a Haar filter and reflecting boundary conditions, obtaining six frequency components. As we use monthly data, the first component ($VRP_t^{D_1}$) captures oscillations of the VRP_{TS} between 2 and 4 months, while components $VRP_t^{D_2}$ $VRP_t^{D_3}$ $VRP_t^{D_4}VRP_t^{D_5}$, capture oscillations of the VRP_{TS} with a period of 4-8, 8-16, 16-32, 32-64 months, respectively. The smooth component $VRP_t^{S_5}$ captures oscillations of the VRP_{TS} with a period exceeding 64 months.

In order to define our predictors, we aggregate the obtained frequencies into three components. The high-frequency component ($VRP_{HF,t}$) corresponds to

$VRP_t^{D_1}$, the medium-frequency component ($VRP_{MF,t}$) is computed as $VRP_{MF,t} = \sum_{i=2}^{5} VRP_t^{D_i}$ and the low-frequency component (VRP_{LF}) corresponds to $VRP_t^{S_5}$. Figure 5 plots the original time series of the variance risk premia, as well as the three components used in our research as predictors.

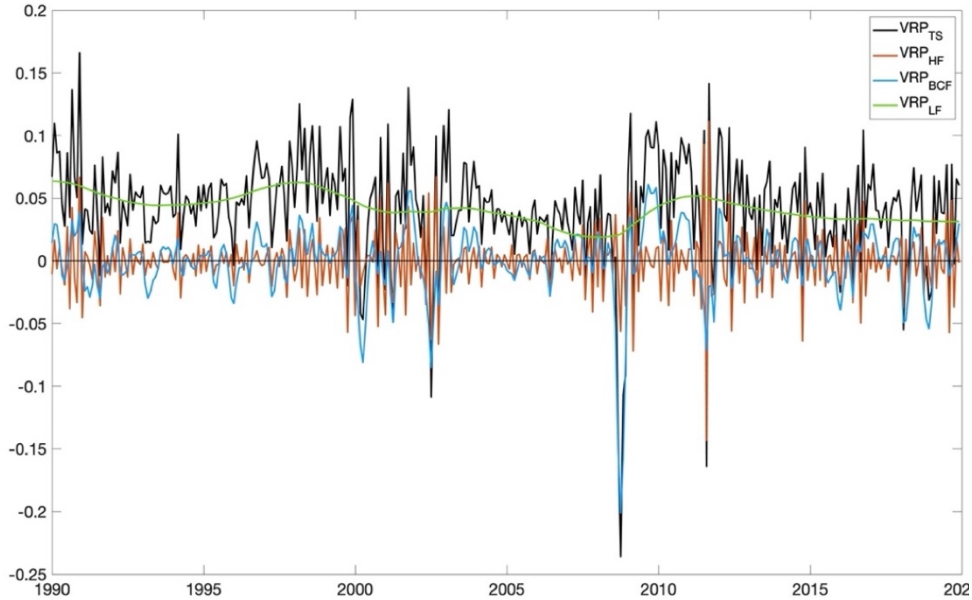

Figure 5: Time series of the variance risk premia VRP_{TS} (black line) and of its three frequency components VRP_{HF}, TMS_{MF} and VRP_{LF} obtained through wavelets decomposition capturing oscillations of the VRP_{TS} less than 4 months (orange line), between 4 and 64 months (blue line) and greater than 64 months (green line).

To better understand the discrete wavelet method used, the following sub-section will summarize this method. Table 1 presents the summary statistics and correlations.

2.2.1 Wavelet Method

Real world data or signals frequently exhibit slowly changing trends or oscillations punctuated with transients. These abrupt changes are often the most interesting parts of the data, both perceptually and in terms of the information they provide. The Fourier Transform is a powerful tool for data analysis, however, it does not represent abrupt changes efficiently. The reason for that is that the Fourier transform represents data as a sum of sine waves which are not localized in time or space. These sine waves oscillate forever, therefore, to accurately analyze signals and images that have abrupt changes we need to use a new class of functions that are well localized in time and frequency, the wavelets.

Regarding the previous information, and in order to decompose the original time series of the variance risk premia, we use the discrete wavelet transform (DWT) multiresolution analysis (MRA).

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

$$y_t = \sum_{\mathrm{k}} \mathrm{s}_{\mathrm{J,k}} \emptyset_{\mathrm{J,k}}(\mathrm{t}) + \sum_{\mathrm{k}} \mathrm{d}_{\mathrm{J,k}} \psi_{\mathrm{J,k}}(\mathrm{t}) + \sum_{\mathrm{k}} \mathrm{d}_{\mathrm{J-1,k}} \emptyset_{\mathrm{J-1,k}}(\mathrm{t}) + \cdots + \sum_{\mathrm{k}} \mathrm{d}_{1,\mathrm{k}} \psi_{1,\mathrm{k}}(\mathrm{t}) \quad (3)$$

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

The wavelet functions are generated from the two types of wavelets: the father wavelets ($\emptyset$),that work like a low-pass filter, capture the smooth and the low frequency component of the series; and the mother wavelets (ψ), that capture the high-frequency components of the series. Through scaling and translation, the wavelet functions are given as follows:

$$\emptyset_{J,k}(t) = 2^{-J/2} \emptyset(2^{-J} - k) \quad (4)$$

$$\psi_{j,k}(t) = 2^{-J/2} \psi\left(2^{-j} t - k\right), \quad (5)$$

and the wavelet coefficients are given by:

$$\mathrm{s}_{\mathrm{J,k}} = \int y(t)\, \emptyset_{\mathrm{J,k}}(t) dt \quad (6)$$

$$\mathrm{d}_{\mathrm{j,k}} = \int y_t\, \psi_{\mathrm{j,k}}(t) dt, \quad (7)$$

where j = 1, 2, 3, …, J.

Eq. (3) can be rewritten as:

$$y_t = \sum_{j=1}^{J} y_t^{D_j} + y_t^{S_J} \qquad (8)$$

where $y_t^{D_j}$, j = 1, 2, 3, …, J are the J wavelet detail components and $y_t^{S_J}$ is the wavelet smooth component. Then, the original time series y_t can be decomposed into different time series components, each defined in the time domain and representing the fluctuation of the original time series in a specific frequency band. The detail components track the high frequency features of y_t whereas the smooth component captures the low frequencies characteristics of y_t.

When it comes to the levels of frequency, a small j compressed by a wavelet function catches fast changing details (high frequencies) and, as j increases, the j wavelet detail components captures slowly changing movements of the time series (low frequencies). The wavelet's smooth component captures the lowest frequency dynamics (i.e.., its long-term behavior).

To overcome some Discrete Wavelet Transform limitations[5] we use the maximal overlap discrete wavelet transform multiresolution analysis (MODWT MRA) as our prevailing method. As Faria and Verona (2018) suggest this methodology has three main advantages:

i. Is not restricted to a particular sample size;
ii. Is translation-invariant, so that it is not sensitive to the choice of starting point for the examined time series;
iii. Does not introduce phase shifts in the wavelet coefficients.

This last feature is especially relevant in an out-of-sample forecasting exercise.

[5] See Masset (2011).

Panel A: Summary Statistics					
Variable	**Mean**	**Median**	**1st percentile**	**99th percentile**	**Std. dev.**
Excess Return (%)	0,005	0,009	-0,109	0,092	0,041
VRP_{TS} (% ann.)	0,042	0,044	-0,109	0,136	0,043
VRP_{HF} (% ann.)	0,000	0,001	-0,064	0,067	0,025
VRP_{MF} (% ann.)	0,000	0,003	-0,085	0,055	0,027
VRP_{LF} (% ann.)	0,042	0,042	0,019	0,063	0,011

Panel B: Correlations					
Variable	**Excess Return**	**VRP_{TS}**	**VRP_{HF}**	**VRP_{MF}**	**VRP_{LF}**
Excess Return	1				
VRP_{TS}	0,07	1			
VRP_{HF}	-0,16	0,73	1		
VRP_{MF}	0,22	0,80	0,26	1	
VRP_{LF}	0,12	0,34	0,00	0,14	1

Table 1: Summary statistics and correlations. Panel A reports the summary statistics for the excess return and the predictors. Panel B reports the correlations between the variables. Predictors are the original time series of the variance risk premia VRP_{TS} and the three frequency components VRP_{HF}, VRP_{MF}, and VRP_{LF} obtained through wavelets decomposition capturing oscillations of the VRP_{TS} of less than 4 months, between 4 and 64 months, and greater than 64 months, respectively. The database contains 360 monthly observations from 1990:01 to 2019:12.

Chapter 3

Empirical Results

This chapter aims to present the main results obtained from the empirical analysis and is divided as follows: In-Sample Predictability results (section 3.1), Out-of-Sample Forecasting results (section 3.2), Economic Analysis (section 3.3) and Sensitivity Analysis (section 3.4).

3.1 In-sample predictability

Let r_t be the excess return for month t. For each predictor $\varkappa_t$, the predictive regression is given by:

$$r_{t:t+1} = \alpha + \beta\varkappa_t + \varepsilon_{t:t+1} \qquad \forall t = 1,\ldots, T, \tag{8}$$

The predictive regression is estimated through ordinary least squares (OLS) to test the statistical significance of $\hat{\beta}$. Stambaugh (1999) and Campbell and Yogo (2006) infer that the predictive regression used in this model suffer some concerns about the statistical point of view. Hence, we follow Huang et al. (2015) and Rapach et al. (2016), and use a heteroscedasticity and autocorrelation-robust t-statistic and a wild bootstrapped[5] p-value to test the null hypothesis [H_0: $\beta = 0$] against the alternative hypothesis [H_1: $\beta > 0$] in the predictive regression.

[5] See Huang et al. (2015) for a detailed description of the bootstrap procedure used to compute empirical p-values.

Table 2 reports the β estimation by OLS of the predictive model (8) and the corresponding R^2 statistic (in percentage).

Predictor	$\hat{\boldsymbol{\beta}}$ (t-stat)	R^2
VRP_{TS}	0,20	4,37
	(4,16) ***	
VRP_{HF}	0,16	0,63
	(1,81) *	
VRP_{MF}	0,32	4,08
	(4,03) ***	
VRP_{LF}	0,47	1,36
	(2,43) **	

Table 2: In-sample predictive regression results. $***$, $**$, and $*$ denote significance at the 1%, 5%, and 10% levels, respectively, accordingly to wild bootstrapped p-values.

Analyzing the results from the regression we can verify that for the time series VRP_{TS} as well as for all the frequencies components the estimated $\hat{\beta}$'s are positive and statistically significant at the 10% level and above the R^2 threshold defined by Campbell and Thompson (2008) [6]. More importantly, the VRP_{TS} and the VRP_{MF} present statistical significance at the 1% level and considerable higher R^2 compared to the remaining frequencies, thus suggesting that the VRP_{TS} and its medium-frequency component are powerful in-sample predictors of equity excess return. Results obtained are in line with previous studies about in-sample predictability of the time series of the variance risk premia (such as Bollerslev et al. (2009)).

[6] Campbell and Thompson (2008) argue that a monthly R^2 of about 0,5% represents an economically relevant degree of return predictability.

3.2 Out-of-sample forecasting

An out-of-sample forecasting approach has become a crucial method to forecast financial variables and in fact many studies show that this approach provides significant results relevant in economic and financial terms, providing us with an exercise that avoids some econometric problems, like the look-ahead bias, since this method only uses information available until the moment of predictability to predict.

Although our predictors perform well in-sample, this does not guarantee that they are good out-of-sample predictors. In fact, as Goyal and Welch (2005) show that, quite often when a variable is a good in-sample predictor, its out-of-sample predictive performance is rather poor. To test that we performed an out-of-sample forecast.

The OOS forecasts are produced using a sequence of expanding windows. We start by choosing our in-sample period that ranges from January 1990 to December 2005. Then, from the information given by this period the first OOS forecast is computed. The previous sample is then increased by one observation and a new OOS forecast is produced. This procedure is used until the end of the sample. Given that, the OOS period runs from January 2016 to December 2019 (168 observations).

In order to evaluate the forecasting performance of each predictor we report the Campbell and Thompson (2008) R^2_{OS} statistics.

R^2_{OS} are calculated as follows:

$$R^2_{OS} = 100\left(1 - \frac{MSFE_{PRED}}{MSFE_{HM}}\right) \tag{9}$$

The benchmark model is the historical mean (HM) forecast $\bar{r_t}$, represented as the average excess return up to time t.

$MSFE_{PRED}$ represents the mean squared forecast error of the predictive model; $MSFE_{HM}$ represents the mean squared forecast error of the historical mean model. These values are calculated as follows:

$$MSFE_{PRED} = \sum_{t=t_0}^{T} (r_t - \hat{r}_t)^2 \tag{10}$$

$$MSFE_{HM} = \sum_{t=t_0}^{T} (r_t - \overline{r}_t)^2 \tag{11}$$

$\hat{r}_t$ represents the excess return forecast from the model using each alternative predictor.

The R^2_{OS} measures the proportional reduction in the mean squared forecast error using the predictive model compared to the historical mean. So, the predictive model used outperforms (underperforms) the historical mean model in terms of MSFE when the R^2_{OS} value is positive (negative).

To test the statistical significance of each result, we follow the literature (e.g. Rapach et al. (2016)) and use the Clark and West (2007) statistic. Using this statistic, we evaluate the MSFE parameters of the historical mean model compared with those resulted by each predictor's predictive model. In particular, we test the null hypothesis that the MSFE of the historical mean model is smaller or equal to the MSFE of the predictive model against the alternative hypothesis, that the MSFE of the historical mean model is greater than the predictive model ($H_0 : R^2_{OS} \leq 0$ against $H_1 : R^2_{OS} > 0$).

The first column of Table 3 presents, for each predictor, the corresponding R_{OS}^2 (in percentage).

Predictor	R^2_{OS}	CER gains	SR gains
VRP_{TS}	7,077***	5,003	0,308
VRP_{HF}	0,788	0,603	0,014
VRP_{MF}	6,772***	6,088	0,378
VRP_{LF}	-3,827	-0,490	-0,067

Table 3: Out-of-sample R-squares (R^2_{OS}), annualized CER gains and SR gains.

Considering the results, we can conclude that the original time series (VRP_{TS}) and the medium frequency component (VRP_{MF}) of the variance risk premia outperform the historical mean benchmark, since the respective R_{OS}^2 are positive and statistically significant, presenting remarkable values of 7,1% and 6,8% respectively. Contrariwise, the low frequency component of the variance risk premia is a poor OOS predictor of the equity excess return (negative R_{OS}^2). Although the high frequency component (VRP_{HF}) displays a positive R_{OS}^2 this value it's not statistically significant.

We can conclude that, from the four predictors used, the original time series and the medium frequency component show a considerable OOS forecasting power.

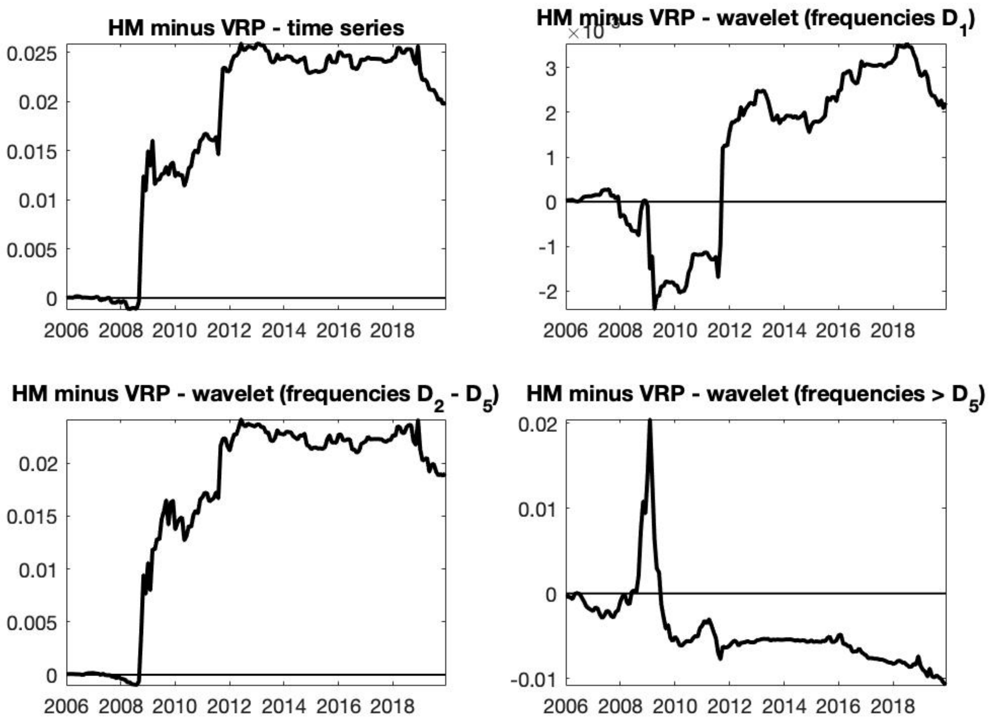

Figure 6: Cumulative squared forecast errors of the variance risk premia and its different components.

To assess the consistency over time of the OOS performance of the predictors, we look at the dynamics of the difference between the cumulative square forecasting error for the historical mean and the cumulative square forecasting error when our alternative predictors are used. We plot this differences in Fig 6.

A positive (negative) value means that the predictive model outperforms (underperforms) the historical mean model, since a positive value expresses that the cumulative square forecasting errors for the historical mean model are higher than those for the predictive model. Both the time series and the medium frequency component consistently outperform, with a similar pattern, the historical mean during the entire OOS period (excluding the first three years).

Overall, from a statistical point of view, both the time series and the medium frequency component of the variance risk premia are remarkably good predictors of the equity excess return.

3.3 Economic analysis

In order to analyze the benefits from using the predictive model and the frequency decomposition, in this section we analyze the results from an economic point of view. We perform an asset allocation analysis, considering a mean-variance investor allocating his wealth between equities and risk-free bills. At the end of month t, the investor optimally allocates:

$$w_t = \frac{1}{\gamma}\frac{\hat{r}_{t+1}}{\hat{\sigma}^2_{t+1}} \tag{12}$$

of the portfolio to equity for the period from t to t + 1. In (12) γ represents the investor's relative risk aversion coefficient, $\hat{r}_{t+1}$ is the OOS forecast of stock return at time t for the period t + 1 and $\hat{\sigma}^2_{t+1}$ is the OOS forecast of the variance of the stock return. Following the literature (e.g. Rapach et al. (2016), Fan et al. (2018) and Faria and Verona (2018)), we use a ten-year moving window of past excess returns to estimate the variance of the excess return, a relative risk aversion coefficient of three, and limit the weights w_t to be larger than -0.5 and smaller than 1.5, in order to limit the possibilities of leveraging or short-selling the portfolio.

We the compute the certainty equivalent return (CER), that can be considered as "the annual portfolio management fee that an investor is willing to pay for access to the alternative forecasting model versus the historical average forecast". If the investor uses the portfolio rule given by w_t, in (12), the average utility (CER) can be expressed as follows:

$$CER = \bar{R}_\rho - 0{,}5\gamma\sigma^2_\rho \tag{13}$$

where $\bar{R}_\rho$ represents the sample mean and σ^2_ρ the variance of the portfolio return. To understand the benefits from using the predictive model, we report annualized certainty equivalent return gains (CER gains). This value is computed as the difference between the CER for an investor that uses the predictive model to forecast excess returns and the

CER for an investor that uses the historical mean model for forecasting. A higher CER gain therefore implies larger economic benefit for the investor. The second column of table 3 presents the results for each predictor. The medium frequency component of the variance risk premia performs better than the other predictors, outperforming the time series by more than 100bp, even though time series exhibited higher R^2_{OS}.

We can also use the Sharpe ratio (SR) to analyze the performance of the predictive model and each predictor from an economic point of view, since SR is used by investors to understand the return of portfolio compared to its risk. Regarding an investor that uses the portfolio rule given by w_t , in (12), the Sharpe ratio (SR) can be expressed as follows:

$$SR = \frac{\overline{R}_\rho}{\sigma_\rho} \tag{14}$$

where $\overline{R}_\rho$ represents the sample mean and σ_ρ represents the standard deviation of the portfolio's excess return.

Similar to the CER gains analysis, SR gains are calculated as the difference between the SR for an investor that uses the predictive model to forecast excess returns and the SR for an investor that uses the historical mean model for forecasting. The right hand-side column of table 3 presents the SR gains results for each predictor. Similar to the CER gains, the medium frequency component is the one that exhibits higher economic gains.

We can infer, then, that information contained in the VRP_{MF} is economically relevant for investors that are interested in forecast and rebalance their portfolios.

3.4. Sensitivity Analysis

3.4.1 Different Wavelet Filter

We evaluate the importance of the wavelet filter used to extract the frequency components of the VRPTS using an alternative wavelet filter: the order 2 symlet. For a given support, the orthogonal wavelet with a phase response that most closely resembles a linear phase filter is called least asymmetric. Symlets are examples of least asymmetric wavelets. This filter is a modified version of the classic Daubechies db wavelets. The frequency bands are chosen to extract exactly the same frequency components in our analysis with wavelets: the high-frequency ($VRP_{SF\text{-}HF}$), the medium-frequency ($VRP_{SF\text{-}MF}$) and low-frequency ($VRP_{SF\text{-}LF}$) components.

Table 4 and Table 5 report the in-sample and out-of-sample results for this wavelet filter, respectively.

Predictor	$\widehat{\boldsymbol{\beta}}$ **(t-stat)**	**R^2**
VRP_{TS}	0,20	4,37
	(4,16) ***	
$VRP_{SF\text{-}HF}$	0,139	0,52
	(1,69) *	
$VRP_{SF\text{-}MF}$	0,28	3,52
	(3,75) ***	
$VRP_{SF\text{-}LF}$	0,44	1,38
	(2,45) **	

Table 4: In-sample predictive regression results. ∗∗∗, ∗∗, and ∗ denote significance at the 1%, 5%, and 10% levels, respectively, accordingly to wild bootstrapped p-values (symlet filter).

Predictor	R^2_{OS}	CER gains	SR gains
VRP_{TS}	7,077***	5,003	0,308
VRP_{SF-HF}	0,580	0,342	-0,006
VRP_{SF-MF}	5,951***	5,791	0,370
VRP_{SF-LF}	-4,195	-0,643	-0,087

Table 5: Out-of-sample R-squares (R^2_{OS}), annualized CER gains and SR gains (symlet filter).

Although similar to the haar filter, when applied, the symlet filter, reduces the in- and out-of-sample performance of the frequency components of the variance risk premia. Thus, these results show that the haar filter enables the extraction of the frequency components of the variance risk premia with a forecasting performance superior to that obtained using this alternative filter.

3.4.2 Different Frequency Grouping

To proceed to a sensitivity analysis regarding the frequency aggregation, we group the six obtained frequencies into three components that differ from the ones used in our research: the alternative high-frequency component (VRP_{HF-A}), the alternative medium-frequency component (VRP_{MF-A}) and the low-frequency component (VRP_{LF}), that capture oscillations of the VRP_{TS} less than 16 months, between 16 and 64 months, and greater than 64 months.

Table 6 and Table 7 report both in- and out-of-sample results for the alternative components used.

Predictor	$\hat{\beta}$ (t-stat)	R^2
VRP_{TS}	0,20	4,37
	(4,16) ***	
VRP_{HF-A}	0,145	1,37
	(2,44) **	
VRP_{MF-A}	0,72	5,66
	(4,74) ***	
VRP_{LF}	0,47	1,36
	(2,44) **	

Table 6: In-sample predictive regression results. ∗∗∗, ∗∗, and ∗ denote significance at the 1%, 5%, and 10% levels, respectively, accordingly to wild bootstrapped p-values (alternative components).

Predictor	R^2_{OS}	CER gains	SR gains
VRP_{TS}	7,077***	5,003	0,308
VRP_{HF-A}	2,545**	3,339	0,210
VRP_{MF-A}	0,887	4,004	0,266
VRP_{LF}	-3,827	-0,490	-0,067

Table 7: Out-of-sample R-squares (R^2_{OS}), annualized CER gains and SR gains (alternative components).

From an in-sample analysis, the predictive regression results obtained are superior to those obtained in our research. In particular the alternative medium-frequency component presents remarkable values for $\hat{\beta}$ and R^2.

Contrariwise, regarding the out-of-sample forecasting power, the alternative medium-frequency component loses value from both a statistical and economic point of view compared to the results obtained in our research. Thus, the aggregation of the frequencies is crucial for the quality of the out-of-sample equity return forecasting exercise, and as shown, the components used in our study outperform out-of-sample those used in this sensitivity analysis.

3.4.3 Short-Sale Constraint

We consider an alternative scenario with short-sale constraint, where we impose w_t in equation (12) to be higher than zero (while keeping the upper limit w_t =1.5). Table 8 presents the results.

Predictor	R^2_{OS}	CER gains	SR gains
VRP_{TS}	7,077***	4,755	0,307
VRP_{HF}	0,788	0,885	0,039
VRP_{MF}	6,772***	5,233	0,330
VRP_{LF}	-3,827	0,742	0,104

Table 8: Out-of-sample R-squares (R^2_{OS}), annualized CER gains and SR gains (short-sale constraint).

This constraint has two opposite impacts on the predictors: it reduces the CER gains and SR gains of the VRP_{TS} and the VRP_{MF}; and increases the SR gains of the VRP_{HF} and the VRP_{LF}. These results show that the economic performance of the predictors that achieved a better performance in our research are negatively affected when and investor is restricted from short-selling.

3.4.4 Different Investor's Relative Risk Aversion

We evaluate the impact of the investor's relative risk aversion coefficient in our asset allocation analysis. We present two different scenarios from the one discussed in section 3.3: we consider a "risk-neutral" investor, with a relative risk aversion coefficient (γ in Eq. (12)) of 1; and a conservative investor, with a relative risk aversion coefficient of 5. Table 9 presents the results for both scenarios.

	Predictor	R^2_{OS}	CER gains	SR gains
	VRP_{TS}	7,077***	3,052	0,235
Scenario 1	VRP_{HF}	0,788	1,598	0,095
$\gamma = 1$	VRP_{MF}	6,772***	4,139	0,287
	VRP_{LF}	-3,827	-1,116	0,045
	VRP_{TS}	7,077***	3,833	0,295
Scenario 2	VRP_{HF}	0,788	0,454	0,012
$\gamma = 5$	VRP_{MF}	6,772***	4,689	0,359
	VRP_{LF}	-3,827	-0,999	-0,174

Table 9: Out-of-sample R-squares (R^2_{OS}), annualized CER gains and SR gains (scenario 1 and scenario 2).

The results on both scenarios show that changing the investor's relative risk aversion coefficient (γ) in Eq.12 reduces the CER gains and SR ratio of the VRP_{TS} and the VRP_{MF}. Our assumption of an investor with a relative risk aversion coefficient outperforms those used in the sensitivity analysis.

Conclusion

The variance risk premia, defined as the difference between the risk neutral (implied variance) and physical (realized variance) expectations of the variance has shown remarkable results as a strong predictor of aggregate future market returns in the literature. This variable is attractive to investors due to the fact that it can be computed using public available information. The principal limitation we came across was related to the fact that the use of the VIX index in our Variance Risk Premia calculations limited us to forecast our predictability to the one-month horizon. We also had to restrict our data sample to observations prior 2020 because this specific year presented a severe outlier that compromised our results.

We investigate the out-of-sample performance of this variable by decomposing the original time series into three different components: the high frequency component, the medium frequency component and the low frequency component. We made this decomposition in order to evaluate if there is crucial information "hidden" in the original time series (one frequency) that could enhance the out-of-sample performance of the variance risk premia.

Our empirical results show that both the original time-series and the medium frequency component of the variance risk premia are strong in- and out-of-sample equity excess return predictors.

We also observed that, although presenting better out-of-sample R_{OS}^2, the original time-series was outperformed in economic terms by its medium frequency component.

We can then conclude that the variance risk premia and its frequency components are interesting variables for many financial players to consider when making their forecasts of equity markets.

For researchers on this topic, interesting implications could arise from studying the impact of variable rare disasters occurrences (causing severe outliers on data), like the one observed (but not reported) in 2020 probably due to the outbreak of the covid 19, on the dynamics of the variance risk premia predictive power.

Bibliography

Bakshi, G., & Kapadia, N. (2003). Delta-Hedged Gains and the Negative Market Volatility Risk Premium. Review of Financial Studies, 16(2), 527–566. https://doi.org/10.1093/rfs/hhg002

Bakshi, G., & Madan, D. (2006). A theory of volatility spreads. Management Science, 52(12), 1945-1956. https://doi.org/10.1287/mnsc.1060.0579

Bakshi, G., Madan, D., & Panayotov, G. (2010). Returns of claims on the upside and the viability of U-shaped pricing kernels. Journal of Financial Economics, 97(1), 130–154. https://doi.org/10.1016/j.jfineco.2010.03.009

Bandi, F. M., Perron, B., Tamoni, A., & Tebaldi, C. (2019). The scale of predictability. Journal of Econometrics, 208(1), 120–140. https://doi.org/10.1016/j.jeconom.2018.09.008

Bekaert, G., & Hoerova, M. (2014). The VIX, the variance premium and stock market volatility. Journal of Econometrics, 183(2), 181–192. https://doi.org/10.1016/j.jeconom.2014.05.008

Bollerslev, T., Tauchen, G., & Zhou, H. (2009). Expected stock returns and variance risk premia. Review of Financial Studies, 22(11), 4463–4492. https://doi.org/10.1093/rfs/hhp008

Bollerslev, T., Marrone, J., Xu, L., & Zhou, H. (2014). Stock return predictability and variance risk premia: Statistical inference and international evidence. Journal of Financial and Quantitative Analysis, 49(3), 633–661. https://doi.org/10.1017/S0022109014000453

Bollerslev, T., Todorov, V., & Xu, L. (2015). Tail risk premia and return predictability. Journal of Financial Economics, 118(1), 113–134. https://doi.org/10.1016/j.jfineco.2015.02.010

Bormetti, G., Corsi, F., & Majewski, A. (2015). Term Structure of Variance Risk Premium and Returns' Predictability. SSRN, 1–49. https://doi.org/10.2139/ssrn.2619278

Campbell, J. Y., & Thompson, S. B. (2008). Predicting excess stock returns out of sample: Can anything beat the historical average? Review of Financial Studies, 21(4), 1509–1531. https://doi.org/10.1093/rfs/hhm055

Campbell, J. Y., & Yogo, M. (2006). Efficient tests of stock return predictability. Journal of Financial Economics, 81(1), 27–60. https://doi.org/10.1016/j.jfineco.2005.05.008

Carr, P., & Wu, L. (2009). Variance risk premiums. Review of Financial Studies, 22(3), 1311–1341. https://doi.org/10.1093/rfs/hhn038

Chen, J., Choudhry, T., Kuo, J., & Zhang, Q. (2016.). Is Liquidity the Trigger for Stock Returns? A Double Take and the Role of Investors' Risk Aversion. European Financial Management Journal.

Clark, T. E., & West, K. D. (2007). Approximately normal tests for equal predictive accuracy in nested models. Journal of Econometrics, 138(1), 291–311. https://doi.org/10.1016/j.jeconom.2006.05.023

Cosemans, M. (2012). The Pricing of Long and Short Run Variance and Correlation Risk in Stock Returns. SSRN Electronic Journal. https://doi.org/10.2139/ssrn.1825934

Cremers, M., & Weinbaum, D. (2010). Deviations from Put-Call Parity and Stock Return Predictability. Journal of Financial and Quantitative Analysis, 45(2), 335-367. https://doi.org/10.1017/S002210901000013X

Dew-Becker, I., & Giglio, S. (2016). Asset pricing in the frequency domain: Theory and empirics. Review of Financial Studies, 29(8), 2029–2068. https://doi.org/10.1093/rfs/hhw027

Drechsler, I., & Yaron, A. (2011). What's vol got to do with it. Review of Financial Studies, 24(1), 1–45. https://doi.org/10.1093/rfs/hhq085

Driessen, J., Maenhout, P. J., & Vilkov, G. (2009). The price of correlation risk: Evidence from equity options. Journal of Finance, 64(3), 1377–1406. https://doi.org/10.1111/j.1540-6261.2009.01467.x

Driessen, J., Maenhout, P. J., & Vilkov, G. (2013). Option-Implied Correlations and the Price of Correlation Risk. SSRN Electronic Journal. https://doi.org/10.2139/ssrn.2359380

Du, J., & Kapadia, N. (2012). Tail and Volatility Indices from Option Prices. Working Paper, University of Massachusetts, Amherst, 1(November 2010).

Goyal, A. & Welch, I. (2005). A Comprehensive Look at The Empirical Performance of Equity Premium Prediction. Review of Financial Studies, 21(4), 1455-1508. https://doi.org/10.1093/rfs/hhm014

Fan, Z., Xiao, X., & Zhou, H. (2018). Variance Risk Premium, Higher Order Risk Premium, and Expected Stock Returns. SSRN Electronic Journal. https://doi.org/10.2139/ssrn.3120260

Faria, G., & Verona, F. (2018). Forecasting stock market returns by summing the frequency-decomposed parts. Journal of Empirical Finance, 45(December 2017), 228–242. https://doi.org/10.1016/j.jempfin.2017.11.009

Faria, G., & Verona, F. (2020). The yield curve and the stock market: Mind the long run. Journal of Financial Markets, 50, 100508. https://doi.org/10.1016/j.finmar.2019.100508

Faria, G., & Verona, F. (2020). Time-frequency forecast of the equity premium. Quantitative Finance. https://doi.org/10.1080/14697688.2020.1820071

Guo, H., Wang, K., & Zhou, H. (2017). Good Jumps, Bad Jumps, and Conditional Equity Premium. SSRN Electronic Journal, October 2014. https://doi.org/10.2139/ssrn.2516074

Huang, D., Jiang, F., Tu, J., & Zhou, G. (2015). Investor sentiment aligned: A powerful predictor of stock returns. Review of Financial Studies, 28(3), 791–837. https://doi.org/10.1093/rfs/hhu080

Kilponen, J., & Verona, F. (2016). Testing the Q theory of investment in the frequency domain. Bank of Finland Research Discussion Paper, 32.

Lettau, M., & S. Ludvigson, 2001. Consumption, Aggregate Wealth, and Expected Stock Returns. The Journal of Finance 56 (3), 815-49. https://doi.org/10.1111/0022-1082.00347

Londono, J. M. (2014). The Variance Risk Premium Around the World. SSRN Electronic Journal, June 2010. https://doi.org/10.2139/ssrn.2517020

Londono, J. M., & Zhou, H. (2017). Variance risk premiums and the forward premium puzzle. Journal of Financial Economics, 124(2), 415–440. https://doi.org/10.1016/j.jfineco.2017.02.002

Masset, P. (2011). Analysis of Financial Time-Series Using Fourier and Wavelet Methods. SSRN Electronic Journal, October, 1–36. https://doi.org/10.2139/ssrn.1289420

Merton, R. C. (1973). An Intertemporal Capital Asset Econometrica: Journal of the Econometric Society, 41(5), 867–887.

Prokopczuk, M., Tharann, B., & Wese Simen, C. (2017). Predicting the Equity Market with Option Implied Variables. SSRN Electronic Journal. https://doi.org/10.2139/ssrn.3070353

Pyun, S. (2019). Variance risk in aggregate stock returns and time-varying return predictability. Journal of Financial Economics, 132(1), 150–174. https://doi.org/10.1016/j.jfineco.2018.10.002

Rapach, D. E., Strauss, J. K., & Zhou, G. (2010). Out-of-sample equity premium prediction: Combination forecasts and links to the real economy. Review of Financial Studies, 23(2), 821–862. https://doi.org/10.1093/rfs/hhp063

Rapach, D. E., Strauss, J. K., & Zhou, G. (2013). International stock return predictability: What is the role of the united states? Journal of Finance, 68(4), 1633–1662. https://doi.org/10.1111/jofi.12041

Rapach, D. E., Ringgenberg, M. C., & Zhou, G. (2016). Short interest and aggregate stock returns. Journal of Financial Economics, 121(1), 46–65. https://doi.org/10.1016/j.jfineco.2016.03.004

Rehman, Z., & Vilkov, G. (2011). Risk-Neutral Skewness: Return Predictability and Its Sources. SSRN Electronic Journal, November 2008. https://doi.org/10.2139/ssrn.1301648

Rosenberg, J. V., & Engle, R. F. (2002). Empirical pricing kernels. Journal of Financial Economics, 64(3), 341–372. https://doi.org/10.1016/S0304-405X(02)00128-9

Rua, A. (2011). A wavelet approach for factor-augmented forecasting. Journal of Forecasting, 30(7), 666–678. https://doi.org/10.1002/for.1200

Stambaugh, R. F. (1999). Predictive regressions. Journal of Financial Economics, 54(3), 375–421. https://doi.org/10.1016/S0304-405X(99)00041-0

Stilger, P. S., Kostakis, A., & Poon, S. H. (2017). What does risk-neutral skewness tell us about future stock returns? Management Science, 63(6), 1814–1834. https://doi.org/10.1287/mnsc.2015.2379

Todorov, V. (2010). Variance risk-premium dynamics: The role of jumps. Review of Financial Studies, 23(1), 345–383. https://doi.org/10.1093/rfs/hhp035

Xing, Y., Zhang, X., & Zhao, R. (2010). What does the individual option volatility smirk tell us about future equity returns? Journal of Financial and Quantitative Analysis, 45(3), 641–662. https://doi.org/10.1017/S0022109010000220

Zhang, K., Gençay, R., & Ege Yazgan, M. (2017). Application of wavelet decomposition in time-series forecasting. Economics Letters, 158, 41–46. https://doi.org/10.1016/j.econlet.2017.06.010