UNIVERSIDADE CATÓLICA PORTUGUESA

Time-Frequency Decomposition and Wavelet-Based Forecasting

Bond-Return Predictability

Alexandre Dias De Vasconcelos

Católica Porto Business School

UNIVERSIDADE CATÓLICA PORTUGUESA

Time-Frequency Decomposition and Wavelet-Based Forecasting

Bond-Return Predictability

Master's Final Assignment – Witten Assignment
Presented to *Universidade Católica Portuguesa*
to obtain a Master's Degree in Finance

by

Alexandre Dias De Vasconcelos

Under supervision of
(PhD) Gonçalo Faria and (PhD) Fábio Verona

Católica Porto Business School
June, 2025

Acknowlegements

I would first like to express my deepest gratitude to my supervisors Gonçalo Faria and Fabio Verona. Their steady guidance and incisive feedback for this project transformed each challenge into an opportunity for growth. I wish to express my deepest gratitude to professor Luís Pacheco. Our extensive conversations, his unfailingly insights and wise words not only introduced me to the fascinating world of finance but, more importantly, imparted a distinctive way of thinking that has reshaped my academic and personal path.

I would like to express my heartfelt gratitude to Pedro Aires, Pedro Ribeiro, Rui Santos, Marta Pimentel, António Pato, João Ferreira, Filipe Condesso, Wilson Martins and ADN group. Your friendship was the perfect antidote to the long hours of research, animated discussions, laughter-filled coffee breaks and well-timed words. I am immensely grateful to Tomás Lascasas, whose infectious sense of fun turned even the most stressful stretches of this thesis into moments worth remembering. From late-night study marathons to karaoke in the car and your knack for balancing work with genuine enjoyment kept the spirits high.

I owe a special thank you to Maria Silva, my steadfast companion through every single day of this project and, without debate, the best cook in Guarda. Your daily presence, whether elbow-to-elbow in the library infused each work session with warmth, humor. Your generosity, culinary wizardry, and constant encouragement nourished both body and spirit, making this thesis journey infinitely richer. I owe a profound debt of gratitude to my family most especially to my mother and father. Your countless sacrifices, wise words, and quiet acts of support have been the bedrock beneath every step of this journey. Finally, to everyone who asked "How's the thesis going?" and listened to the long answer. Your interest mattered more than you know. I hope the pages that follow repay your support.

Abstract

This thesis investigates whether wavelet-based frequency decomposition can enhance out-of-sample predictability and economic value in forecasting long-term government bond returns. Traditional time-domain models have often struggled to maintain forecasting accuracy when market regimes shift, motivating an exploration of multi-scale methods. Building on studies such as Faria & Verona, 2020, we apply the maximal overlap discrete wavelet transform (MODWT) to decompose various macroeconomic predictors into high-frequency, business-cycle-frequency, and low-frequency components. Out-of-sample forecasts are then generated via an expanding window approach and evaluated against a simple historical-mean benchmark. Our empirical findings reveal that certain predictors, particularly the term spread and book-to-market, exhibit substantially higher out-of-sample R^2 and Certainty Equivalent Return (CER) gains once the most relevant frequency frequencies are isolated. The mean–variance allocation framework demonstrates that wavelet-based forecasts offer notable utility improvements for moderate risk-aversion investors. However, the benefits depend strongly on portfolio weight constraints. Relaxed bounds amplify potential returns (and losses), a purely long-only setting yields modest but stable gains. Overall, these results echo the broader frequency-domain literature (Kim & In, 2005) by underscoring how wavelet-based methods can reveal valuable time-horizon-specific signals for bond-return forecasting, provided that each investor's risk profile and trading constraints are carefully considered.

Keywords: Bond returns, Wavelet decomposition, Frequency-domain forecasting, Out-of-sample prediction, Certainty Equivalent Return.

Abstrato

Esta tese investiga se a decomposição de frequência baseada em wavelets pode melhorar a capacidade de previsão fora da amostra e o valor económico na previsão de retornos de obrigações governamentais de longo prazo. Os modelos tradicionais no domínio do tempo têm frequentemente tido dificuldade em manter a precisão da previsão quando os regimes de mercado mudam, motivando uma exploração de métodos multi-escala. Com base em estudos como Faria & Verona, 2020, aplicamos a maximal overlap discrete wavelet transform (MODWT) para decompor vários preditores macroeconómicos e técnicos em componentes de alta frequência, frequência de ciclo de negócios e baixa frequência. As previsões fora da amostra são então geradas através de uma abordagem de janela deslizante expansível e avaliadas em relação a um benchmark simples de média histórica. Os nossos resultados empíricos revelam que certos preditores, particularmente o term spread e o book-to-market, exibem R^2 fora da amostra e ganhos de Retorno Equivalente de Certeza (CER) substancialmente mais altos uma vez que as frequências mais relevantes são isoladas. O framework de alocação média-variância demonstra que as previsões baseadas em wavelets oferecem melhorias notáveis de utilidade para investidores com aversão moderada ao risco. No entanto, os benefícios dependem fortemente das restrições de peso da carteira. Limites mais relaxados amplificam os retornos (e perdas) potenciais, enquanto uma configuração puramente long-only produz ganhos modestos mas estáveis. No geral, estes resultados ecoam a literatura mais ampla do domínio da frequência (Kim & In, 2005), sublinhando como os métodos baseados em wavelets podem revelar sinais valiosos específicos do horizonte de tempo para a previsão de retornos de obrigações, desde que o perfil de risco e as restrições de negociação de cada investidor sejam cuidadosamente considerados.

Palavras-chave: Retornos de obrigações, Decomposição de wavelet, Previsão no domínio da frequência, Previsão fora da amostra, Retorno Equivalente de Certeza.

Index

Table of Figures

Table of Tables

Introduction

Bond markets play a central role in global finance, serving as the backbone of sovereign and corporate funding, shaping monetary policy transmission, and acting as a key risk management tool for investors. Despite their importance, successfully predicting bond returns remains challenging. Conventional linear time-domain methods often focused on yield spreads, macroeconomic indicators, or technical signals have yielded mixed results, particularly out-of-sample, where unstable relationships or shifting market regimes can undermine previously strong in-sample evidence. These inconsistencies motivate a search for more robust techniques capable of isolating truly persistent patterns in bond returns.

A promising new direction is the frequency-domain approach, which acknowledges that economic and financial time series typically operate on multiple time scales. Some short-run variations may reflect transient liquidity conditions, policy surprises, or brief risk-on/risk-off sentiment, while medium- and long-run cycles may capture the business cycle, secular macro trends, or slow-moving risk premium components. By decomposing predictors into distinct frequency bands, investors can better identify which time scales carry genuinely useful information, and which are dominated by noise. Recent studies, (Faria & Verona, 2018, 2020), demonstrate that applying wavelet transformations to equity or bond predictors can significantly improve out-of-sample forecasting performance by revealing valuable frequency-specific signals. This thesis builds on that premise, applying wavelet-based multiresolution analysis to a range of macroeconomic indicators frequently cited for bond-return forecasting.

Book to Market Time Series

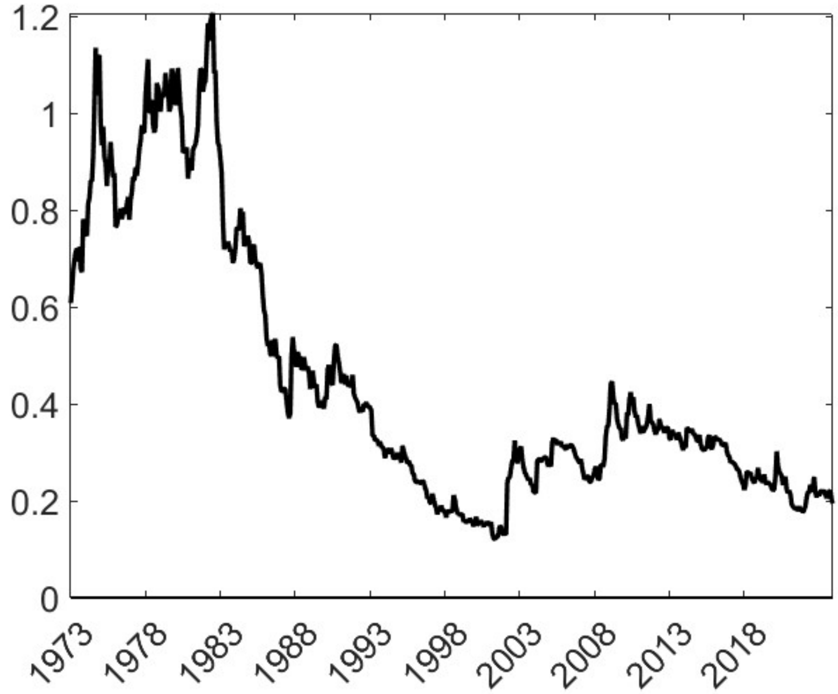

Figure 1: Original time series of predictor Book-to-market. Sample period 1973 - 2023, monthly frequency.

Book to market High-frequency component

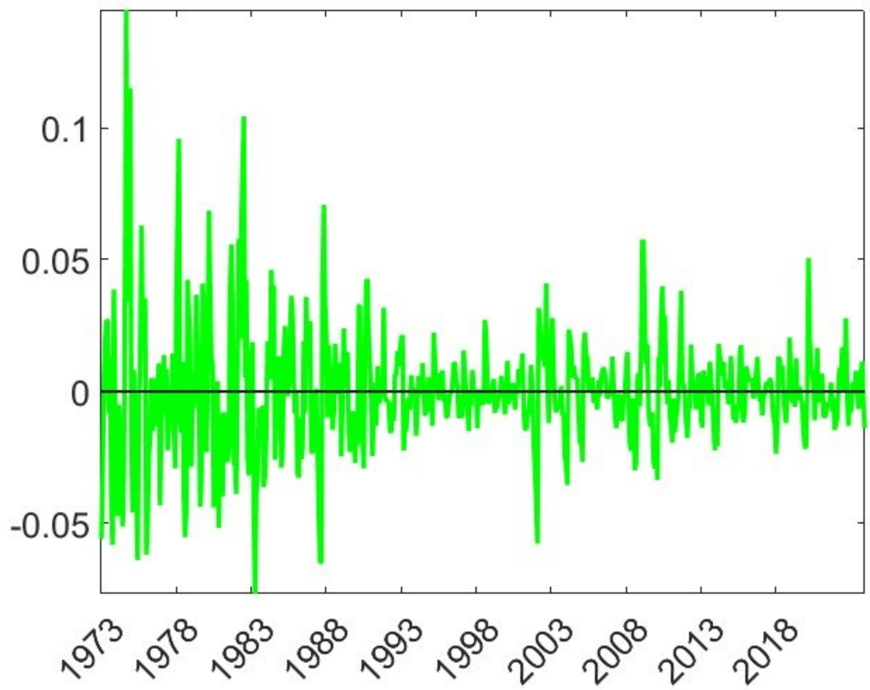

Figure 2: High-frequency component of predictor b/m. High-frequency captures oscillations of less than 16 months in the sample period of 1973 - 2023, monthly frequency

Book to Market Business-cycle frequency component

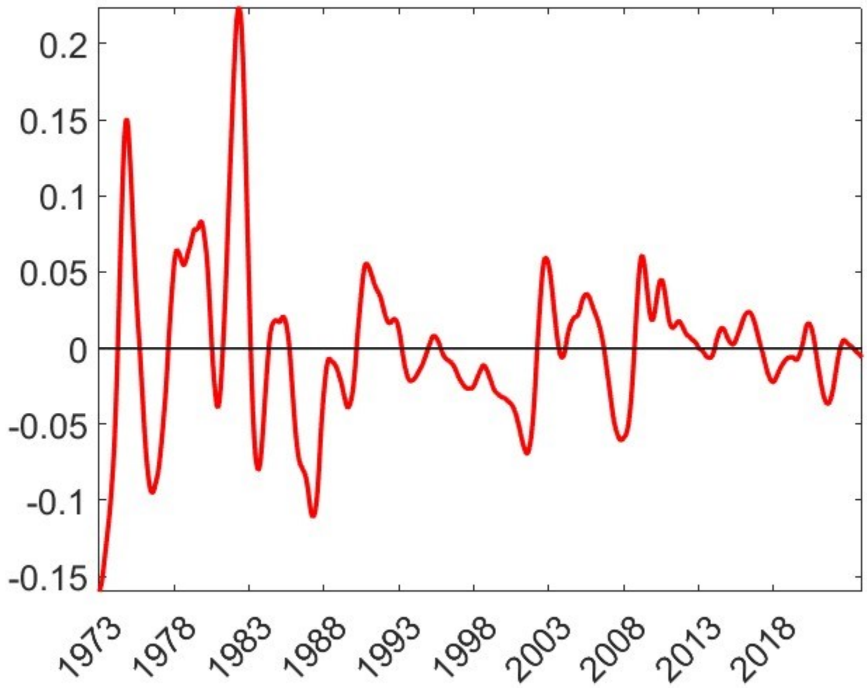

Figure 3: Business-cycle component of predictor b/m. Business-cycles captures oscillations between 16 and 128 months in the sample period of 1973 - 2023, monthly frequency.

Book to Market Low-frequency component

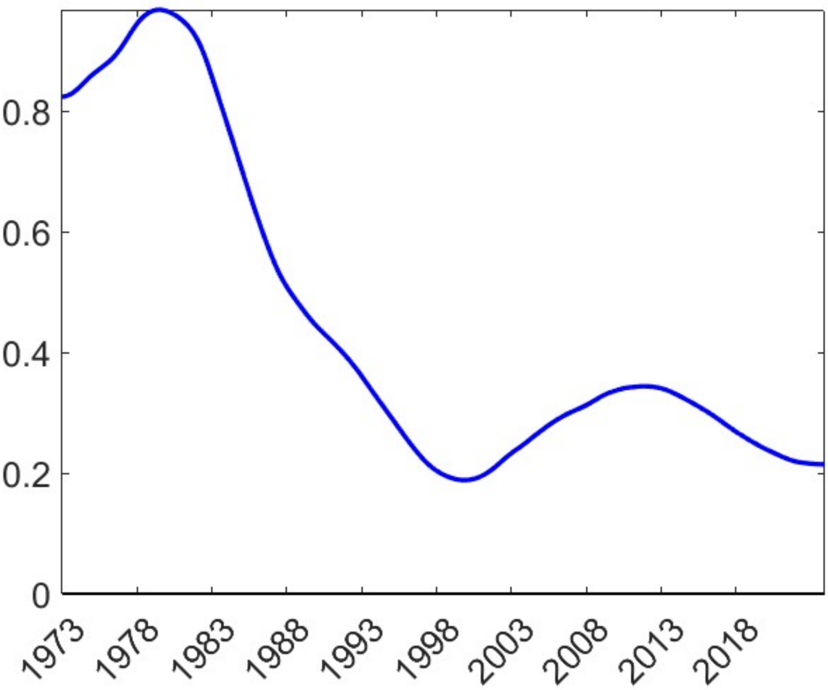

Figure 4: Low-frequency component of predictor b/m. Low-frequency captures oscillations of more than 128 months in the sample period of 1973 - 2023, monthly frequency.

Our goal is to determine whether wavelet-decomposed predictors outperform their raw time-domain equivalents in predicting long-term government bond returns out-of-sample and to assess the real-world utility of these frequency-domain forecasts by measuring Certainty Equivalent Return gains under different risk-aversion parameters and portfolio weight constraints.

We collect monthly data on bond returns and a suite of fourteen predictors, ranging from yield spreads and macro variables covering the period from January 1973 to December 2023. Each predictor is decomposed into high-frequency (HF), business-cycle frequency (BCF), and low-frequency (LF) components using the maximal overlap discrete wavelet transform. By separating out short, medium, and long-run patterns, we can pinpoint the timescales most relevant to forecasting bond excess returns. Forecasting is conducted in an expanding window, out-of-sample manner. We compare each wavelet-based predictor to its raw time-series version and to a baseline historical-mean benchmark, gauging performance via R^2_{OS}, (Campbell & Thompson, 2008, Clark & West, 2007) statistical tests. To capture economic value, we embed these forecasts in a mean–variance optimization framework, compute portfolio weights for each horizon, and quantify CER gains the premium an investor would pay to access the wavelet-based forecast rather than using the historical mean. We then see how results vary with investor risk aversion and portfolio weight constrains.

Certain predictors particularly term spread and book-to-market exhibit strong improvements when wavelet-decomposed, with some frequency bands generating substantial out-of-sample gains that surpass raw time-domain forecasts. Looser weight constraints allow deeper shorts or higher leverage, magnifying both positive and negative outcomes. Strictly long-only approaches moderate returns but reduce drawdowns, often transforming certain negative signals into mildly positive ones. Across all scenarios, wavelet decomposition shows the potential to identify hidden time horizon specific relationships that would remain invisible in standard time-domain regressions.

Following this introduction, Chapter 1 provides a detailed literature review of bond return predictability, outlining both time-domain and frequency-domain approaches, and wavelet-based forecast. Chapter 2 describes the data and methodological framework, explaining the MODWT-based decomposition, the out-of-sample forecast design, and the CER metric. In Chapter 3, we present empirical results, focusing first on statistical out-of-sample performance and then on economic performance, systematically varying investor risk aversion and weight constraints to test the robustness of the findings. Finally, Chapter 4 concludes by summarizing the key insights, identifying limitations, and suggesting avenues for future research. By dissecting bond return predictors across distinct time scales, this thesis aims to offer fresh insights into both the theoretical underpinnings of bond return predictability and the practical task of portfolio allocation. The results highlight how much value wavelet decomposition can add to real world scenarios where structural changes, noisy signals, and risk constraints are unavoidable.

Chapter 1

1. Literature review

In this chapter, we review the literature on bond return predictability, forecast in the frequency domain, and the application of wavelets in economics and finance. These topics have garnered significant attention due to their impact on improving bond return forecasting models. This review provides an in-depth exploration of the key findings and methodologies applied across the studies, laying the groundwork for the methodology utilized in this dissertation, which integrates the frequency domain with wavelet analysis for bond return prediction.

1.1 Bond Return Predictability

Under the simplest form of the expectations hypothesis, long-term bond yields represent the average of current and expected future short-term interest rates and a constant risk premium (Fama & Bliss, 1987). In such a scenario, any excess bond return beyond the short rate would largely be unpredictable. That forward-spot spreads contained information about future bond excess returns, suggesting that risk premium varied over time. Later work by Shiller (1991) confirmed that certain slope measures of the yield curve could forecast changes in yields. Collectively, these studies implied that bond markets might exhibit predictable patterns in their risk premium, at least within specific samples or horizons.

At the core, time-varying bond risk premium suggests that shocks to macroeconomic and financial conditions shift investor risk aversion or the probability distribution of future interest rates. Such an argument sets the stage for exploring how yield-curve signals, macro factors, or other variables might forecast bond returns. Yet, an important distinction arises between "in-sample" (full-sample) regression evidence and "out-of-sample" forecasting performance. Early in-sample work (Cochrane & Piazzesi, 2005) found compelling evidence that a linear combination of forward rates forecast one-year bond returns with a high R-squared. Nonetheless, subsequent tests raised doubts about whether these relationships remained stable and robust enough to be exploited out of sample.

Welch & Goyal (2008) seminal paper examined a large array of variables claimed to predict equity returns such as dividend-price ratios, earnings yields, inflation rates, and interest rate spreads and found that, in most cases, these predictors fared poorly compared to a simple historical mean forecast when tested out-of-sample. Their findings called into question the practical usefulness of many regressions that seemed significant in-sample.

In the bond context, analogous results emerged. Various term structure-based models, while showing promise on the sample used for fitting the regression, tended to have negligible or negative out-of-sample R-squared once tested on fresh data or on subsequent periods with different market conditions. Reasons for this instability might include structural breaks in monetary policy, the presence of regime shifts, changes in macroeconomic volatility, or non-linearities that a linear predictive regression cannot easily capture. As Johannes et al. (2014) point out in the equity domain, parameter uncertainty and model misspecification can lead to wide confidence intervals on out-of-sample forecasts, thus effectively nullifying any gains from seemingly strong in-sample relationships.

In response to this challenge, there are several avenues for improving out-of-sample predictions like time-varying parameters models, inclusion of macroeconomic or unspanned factors, forecast combination and machine learning. For time-varying parameter models there are empirical frameworks that allow coefficients to adapt over time (e.g., Markov-switching, rolling regressions, Bayesian updating) have been used to handle structural changes in the data. Studies like Dangl & Halling (2012) examine how time-varying parameter regressions can salvage some level of out-of-sample performance for equity returns. In the bond realm, the idea is that as monetary regimes and risk aversion shift, the mapping between a predictor (like yield spreads) and future returns changes systematically over time. The inclusion of macroeconomic or unspanned factors is a second line of research focuses on expanding the information set beyond yields alone. Ludvigson & Ng (2009) find that macro factors extracted from large panels of real and nominal variables add predictive power for bond excess returns above and beyond what is contained in the cross-section of yields alone. This concept, sometimes referred to as "unspanned macro risk", suggests that certain aspects of the economy are not captured by yield curve shapes. By adding these macro variables or factors, researchers have documented improved forecasting performance, particularly at intermediate horizons. A third strategy, relevant for both equity and bond return predictability, is forecast combination. Given that no single model consistently outperforms all others, combining forecasts from multiple models (each with a different focus, variable set, or method) can yield superior results. This approach mitigates the risk of selecting a single best model that might fail in a regime shift. Rapach et al. (2010) illustrate that combining equity premium forecasts from many predictors leads to stable gains relative to using any single predictor. Likewise, in the bond domain,Gargano et al. (2019) show that combining models that incorporate yield curve variables, macro factors, and volatility dynamics can produce meaningful out-of-sample improvements in both statistical accuracy and investor utility. Another avenue is machine learning. Recently, tree-based methods (random forests, gradient boosting) and neural networks have been applied to bond return forecasting, as in Bianchi et al. (2021). These flexible algorithms can, in principle, handle large predictor sets, interactions, and non-linearities. Early evidence suggests that combining yield curve data with macroeconomic or market-based predictors in a machine learning framework can yield large out-of-sample R-squared values, along with significant economic value for investors. Yet, these methods can be prone to overfitting if not carefully validated.

Despite these strides, achieving robust out-of-sample predictability remains challenging. The literature increasingly focuses on how to isolate stable predictive signals and avoid the large noise inherent in bond returns. A promising new direction is to exploit the frequency domain, i.e., analyzing how the predictability of bond returns may differ across distinct time horizons or cyclical frequencies. The next sections examine the general concept of frequency-domain approaches in economics, how wavelet transforms can operationalize these ideas, and evidence for whether multi-scale or wavelet-based forecasting can yield incremental gains.

1.2 Frequency Domain and Wavelets in Economics and Finance

Many economic and financial time series contain features at multiple time scales: short-term oscillations (days, weeks, or months), business cycle fluctuations (1–8 years), and low-frequency components reflecting secular trends (spanning a decade or more). Traditional time-domain approaches often assume a single dynamic process drives the data, or at least that the same regressors and functional forms apply uniformly across all frequencies. Yet, it is plausible that the cyclical drivers of bond returns differ from the secular, low-frequency drivers. For instance, high-frequency fluctuations could be influenced by short-lived monetary policy surprises, liquidity events, or flight-to-safety episodes. Meanwhile, low-frequency components might relate to long-run inflation expectations, demographic shifts, or global savings imbalances that gradually affect yields over many years.

Frequency-domain analysis attempts to disentangle these different components. A standard approach is to use Fourier transforms, which represent time-series data as sums of sinusoids with different frequencies. This reveals how variance is distributed across the frequency spectrum and can highlight cyclical patterns. However, classical Fourier methods are purely frequency-based and do not localize patterns in time. A structural break or regime shift could be missed if it only appears for a portion of the sample. One notable application of Fourier methods in finance is presented in Carr & Madan (1999), where the Fast Fourier Transform (FFT) is used to efficiently compute option prices based on the characteristic function of asset returns. Their work demonstrates the power of Fourier-based techniques in financial modeling, particularly for option valuation and pricing.

Wavelet analysis decomposes a series into components associated with specific frequency bands, each localized in time. The father wavelet captures the smooth, low-frequency part of the data (long-run trend), while the mother wavelet captures deviations at progressively higher frequencies. One advantage is that wavelets can handle non-stationarities or structural breaks more gracefully than global Fourier expansions, because wavelet basis functions can be shifted and scaled to adapt to local features (Ramsey & Lampart, 1998).

By decomposing the series, a researcher can study how different frequency components correlate with or predict other variables (e.g., bond returns). If short-term components contain mostly noise or ephemeral shocks, ignoring them might enhance the signal-to-noise ratio for forecasting. Conversely, if some medium-frequency band captures business cycle fluctuations that feed into risk premium, focusing on that band could reveal stronger relationships.

In finance, a variety of wavelet-based methods have emerged, each offering unique advantages for analyzing financial time series. Wavelet denoising is one such approach, which focuses on eliminating high-frequency noise from financial data to reduce volatility in either the predictor or target series. By filtering out these short-term fluctuations, wavelet denoising helps improve the accuracy and stability of financial models. Another widely used method is multiscale factor extraction, which involves summarizing large datasets at different frequency levels and then analyzing the predictive power of these scale-specific factors. This technique has been particularly effective in macroeconomic forecasting, as demonstrated by Rua (2017), who showed that extracting multiscale factors enhances the predictive performance of financial models. Additionally, scale-by-scale regressions provide an alternative approach where each wavelet-transformed component of the dependent variable is regressed on corresponding wavelet components of explanatory variables. This method, as explored by Gallegati & Ramsey (2013), allows researchers to capture frequency-dependent relationships between financial and macroeconomic variables, ultimately improving the robustness and interpretability of regression models.

Ferreira & Santa-Clara (2011) introduce the Sum-of-the-Parts (SOP) framework, a purely time-domain strategy that boosts return predictability by letting seasonality speak for itself. SOP carves the calendar into non-overlapping parts and estimates a separate forecasting regression for each part. When a new observation arrives, the model consults only the regression trained on that same part before rolling the forecast forward. By matching seasonal patterns with season specific parameters, SOP filters out noise created by pooling structurally different periods and has been shown to deliver appreciable out-of-sample gains and higher certainty-equivalent returns relative to a single-equation benchmark.

Faria & Verona (2018) added an extension known as the SOPWAV (Sum-of-Frequency-Parts using Wavelets) method, which integrates wavelet decomposition into the SOP framework. SOPWAV employs wavelet transformations to decompose financial time series into orthogonal frequency components before applying predictive modeling to each component. Wavelets, unlike traditional Fourier analysis, allow for both time and frequency localization, making SOPWAV particularly effective for handling non-stationary data with structural breaks. Empirical findings suggest that SOPWAV significantly improves forecast performance compared to standard SOP and single-equation predictive models, particularly in capturing long-term trends while filtering out high-frequency noise.

These approaches have proven particularly useful in asset return forecasting, where different economic forces influence returns at different horizons. The SOPWAV method has demonstrated strong performance in predicting equity risk premium, as it effectively separates short-term fluctuations from longer-term cycles. By leveraging multi-resolution analysis, SOPWAV enhances the ability of predictive models to detect persistent signals embedded in financial time series, contributing to more reliable and economically valuable forecasts.

A few studies show that standard asset pricing relations can be misleading if frequency dependence is ignored. For instance, Kim & In (2005) find that the correlation between stock returns and inflation significantly differs at long versus short horizons, a pattern that time-domain regressions gloss over. Gençay et al. (2003) demonstrate that the beta in the Capital Asset Pricing Model changes by frequency, suggesting that an asset's systematic risk is time-scale-specific. Bandi et al. (2019) document that forward/backward regressions of future excess returns on backward aggregated volatility produce a hump-shaped pattern in R-squared that peaks at intermediate horizons, highlighting that certain cyclical frequencies are more predictable.

All these findings converge on the intuition that certain frequencies contain stronger or more persistent signals for forecasting returns.

The main benefit of frequency-domain forecasting is improved signal extraction: if only certain frequencies are truly predictable, wavelet decomposition helps isolate them, thus improving the forecast's overall clarity. Also, wavelet-based approaches can better handle structural breaks by localizing them in time. Another advantage is interpretability: a researcher can identify at which frequencies the strongest predictive relationships occur, offering insight into the economic phenomena behind the forecast. However, there are several potential traps associated with wavelet-based approaches in financial analysis. One major concern is overfitting, as decomposing data into numerous scales or applying complex non-linear transformations can increase model complexity and introduce the risk of data-snooping, leading to spurious predictive relationships.

Cochrane & Piazzesi (2005) show that a particular tent-shaped combination of forward rates significantly predicts one-year excess returns on Treasury bonds. More recently, Bianchi et al. (2021) demonstrate that incorporating machine learning with yield curve and macro data can yield out-of-sample success.

Translating wavelet-based forecasts into trading strategies could yield even greater economic value if the wavelet approach significantly boosts predictive accuracy. Indeed, wavelet-based equity forecasting studies have reported large improvements in portfolio returns (Faria & Verona, 2018). Extrapolating from these findings, one can hypothesize that frequency-domain bond return forecasts might also deliver tangible utility gains.

The work of Faria & Verona (2020), Rua (2017), and Gargano et al. (2017) has demonstrated the potential of wavelet-based forecast for improving bond return predictions. By combining this method with macroeconomic and financial variables, these studies have opened new avenues for improving bond market forecasts, offering valuable insights for investors and policymakers alike. This dissertation extends these findings by exploring the application of wavelet-based forecast in the context of bond return predictability, with the goal of enhancing forecast accuracy and economic value.

Chapter 2

2. Data and Methodology

This chapter outlines the dataset and forecast variables employed (Section 2.1) and then describes the methodological framework in detail (Section 2.2). The methodology includes an explanation of wavelets (Section 2.2.1), the out-of-sample forecasting design (Section 2.2.2), the forecast evaluation metrics (Section 2.2.3), and concludes with a discussion of the certainty equivalent return (Section 2.2.4).

2.1 Data and predictors

The dependent variable in this study is the excess bond return, defined as one-month return on long term government bond, $Gov10Y_t$, minus the lagged risk-free rate, $R_{f,t-1}$:

$$ER_t = Gov10Y_t - R_{f,t-1}$$

For the predictive regressions, there are fourteen macroeconomic indicators derived from Goyal and Wech database. Specifically, the sample includes monthly data covering the timeframe from January 1973 to December 2023.

The predictors under consideration are:

1. AAA bond yield (AAA): Moody's AAA rated corporate bond yield, representing the highest quality credit.
2. BAA bond yield (BAA): Moody's BAA rated corporate bond yield, indicating higher credit risk than AAA.
3. Long-term government bond yield (Ity): Yield on a long-term government bond (often 10 year or more) tracked by Goyal and Welch.

4. Corporate bond return (Corpr): Aggregate return (price changes + coupon payments) on corporate bonds.
5. Dividend price ratio (d/p): Dividends divided by current stock price, commonly for the S&P 500 index.
6. Dividend yield (d/y): Annual dividend payout as a percentage of share price
7. Earnings to price ratio (e/p): Aggregate earnings per share relative to current stock price (S&P 500)
8. Book to market (b/m): Ratio of a firm's or index's book value to its market capitalization.
9. Term spread (tms): Difference between a long-term treasury yield and a short-term rate (e.g. 10Y minus T-bill).
10. Default yield spread (dfy): Yield difference between lower-grade bonds (BAA) and high-grade bonds (AAA).
11. Default return spread (dfr): Difference in total returns between corporate bonds and government bonds.
12. Inflation (infl): Monthly inflation rate to measure price-level increases.
13. Oil price changes (wtexas): Month-over-month changes in West Texas Intermediate (WTI) crude oil prices.
14. Stock to bond yield gap (ygap): Spread between the equity market's earnings yield and a bond yield.

The summary statistics appear in appendix A.

2.2 Methodology

Below, we describe how the frequency-based wavelet decomposition is carried out and how the out-of-sample forecasts and their evaluations are executed.

2.2.1 Frequency domain and Wavelets

Economic and financial time series often combine multiple cyclical and trend components that can operate at distinct frequencies over time. Traditional spectral analysis via Fourier transforms has helped isolate some of these frequency-specific behaviors (Ramsey & Lampart (1998) and Crowley (2007)). However, because the classical Fourier basis uses sine and cosine waves extending infinitely in time, identifying abrupt changes and time-localized features can be problematic (Kim & In, 2005).

Wavelet transforms provide a powerful alternative by simultaneously offering both time and frequency localization (Gençay et al. (2002), Gençay et al. (2005) and Crowley (2007)). Instead of representing a signal as an infinite sum of sines and cosines, wavelet transforms rely on short wavelets that concentrate in time and can vanish away from their center points.

A wavelet is typically defined as a function ψ of limited duration and integral zero, a mother wavelet:

$$\int \psi(t)dt = 0$$

implying it oscillates around the time axis, capturing high frequency "details." By contrast, a father wavelet ϕ integrates to one:

$$\int \phi(t)dt = 1$$

thus, representing the smooth, low-frequency part of a time series (Crowley, 2007). These two functions generate families of scaled and translated wavelets that partition the original data by frequencies (or scales). Formally, the continuous wavelet transform (CWT) projects a time series onto each wavelet at every possible scale and location. In practice, economists and financial analysts often turn to the discrete wavelet transform (DWT), which is computationally more convenient and facilitates a multiresolution analysis (MRA) (Gencay et al., 2002). In a J-level MRA using the DWT, one breaks the original time series $y_{t_{t=1}}^{N}$ into J detail components plus one smooth component:

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

In the above expression $\phi_{J,k}(t)$ and $\psi_{j,k}(t)$ are translated and scaled versions of the father and mother wavelets, respectively, here j ranges from 1 to J (the total number of scales), and k defines the location or shift in time, The wavelet function $\phi_{J,k}(t)$ and $\psi_{j,k}(t)$ are typically produced by rescaling $\phi(t)$ and $\psi(t)$ through:

$$\phi_{J,k}(t) = 2^{\frac{-J}{2}}\phi(2^{-J} - k)$$

$$\psi_{j,k}(t) = 2^{\frac{-J}{2}}\psi(2^{-j} - k)$$

The corresponding wavelet coefficients $s_{J,k}$ (for the father wavelet) and $d_{j,k}$ (for the mother wavelet) are obtained by projecting the original series y_t onto each of these scaled and translated wavelet functions:

$$s_{J,k} = \int \gamma t \phi_{J,k}(t) d(t)$$

$$d_{j,k} = \int \gamma t \psi_{j,k}(t) d(t)$$

The wavelet multiresolution decomposition of a time series y_t can be presented by:

$$y_t = y_t^{S_J} + \sum_{j=1}^{J} y_t^{D_j}$$

Where $y_t^{S_J}$ represents the lowest-frequency band (long-run trend) and each $y_t^{D_j}$ corresponds to a specific interval. With monthly data we set the number of decomposition levels. This choice satisfies at J=6. This choice satisfies the MODWT requirement $J \leq (log_2 N)$ for our 192-month initial window. The resulting bands are $y_t^{D_1}$ that isolates oscillations with periods of roughly 2-4 months, $y_t^{D_2}$ captures 4-8 months, $y_t^{D_3}$ captures 8-16, $y_t^{D_4}$ captures 16-32, $y_t^{D_5}$ captures 32-64 and $y_t^{D_6}$ capturing fluctuations from 64 to 128 months. The smooth component $y_t^{S_6}$ encompasses periods longer than 128 months (Aguiar-Conraria et al. (2012), Faria & Verona (2020) and Percival & Walden, (2000)).
We follow Faria & Verona (2018, 2020), among others, in employing the maximal overlap discrete wavelet transform (MODWT). Unlike certain DWT formulations the MODWT does not have the restriction of a specific sample size, is invariant to time shifts, so the wavelet coefficients remain essentially the same if the data are cyclically shifted. Also, it avoids phase shifts, preserving alignment of original and transformed series an important property for real-time forecasting (Faria & Verona, (2018) and Rua (2017)) These benefits make the MODWT-based MRA particularly suited to analyzing potential predictive relationships in financial applications, since sample start or end points are rarely arbitrary and time alignment is crucial for out-of-sample forecast exercises.
So, one step further is to group the detail components to form three aggregated frequency bands. The high-frequency (HF) band is the sum $y_t^{D_1} + y_t^{D_2} + y_t^{D_3}$, which reflects short-horizon movements of approximately 2–16 months. The business-cycle frequency (BCF) band is the sum of $y_t^{D_4} + y_t^{D_5} + y_t^{D_6}$ capturing

cycles from around 16 months to 128 months. Finally, the low-frequency (LF) band is the smooth part $y_t^{D_6}$, which summarizes very slow-moving trends above 128 months, often exceeding a decade. By splitting any original predictor x_t into x_t^{HF} , x_t^{BCF} and x_t^{LF} , plus the unfiltered x_t itself, one can discover which frequencies truly drive out-of-sample variations in bond returns. For example, some macro indicators may prove relevant at lower frequencies that align with economic fundamentals, while certain technical indicators exhibit short-run predictive content at higher frequencies (Kim & In, 2005; Faria & Verona, 2018).

These choices yield seven wavelet-transformed series from a single predictor (six detail components D_1,..., D_6 and one smooth component S_6), which we then aggregate into the three frequency bands.

Bellow is figure 5 that represents the book to market predictor. There we can see the difference components (low-frequency, business-cycle, high-frequency and the original time series).

Book to Market and its frequency components

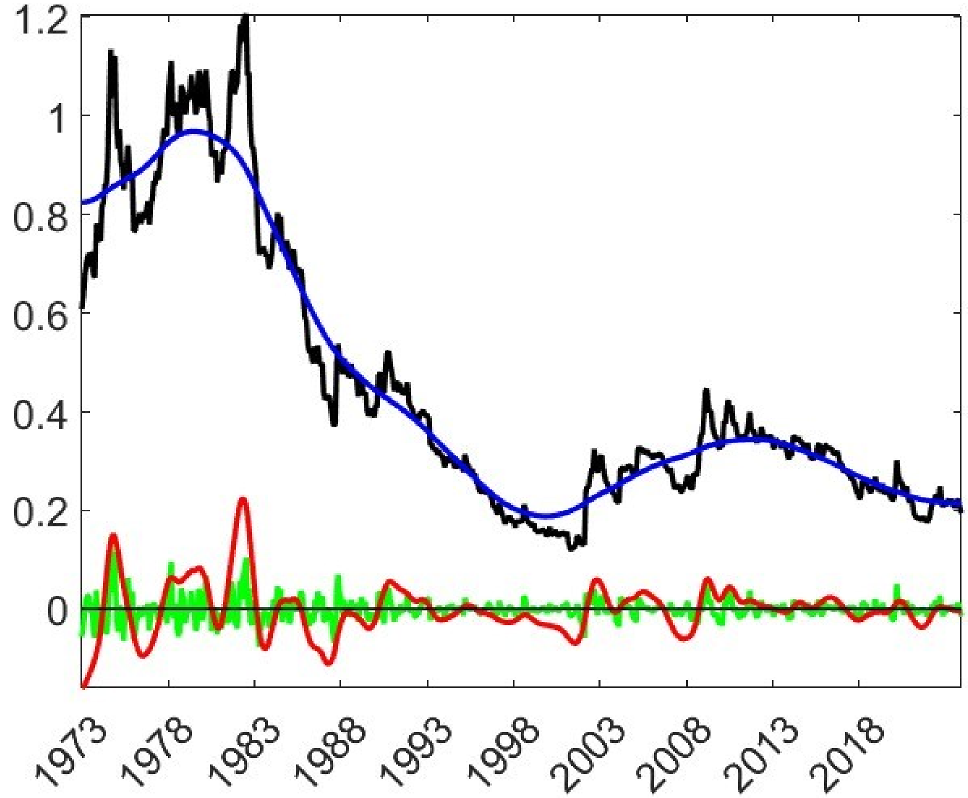

Figure 5: b/m wavelets decomposition. The original time series of each predictor (black) and its 3 frequency components obtained through wavelets decomposition: high-frequency (green), business-cycle (red) and low frequency (blue).

In this study, we adopt the Haar filter, known for its simplicity and for being interpretable in terms of local moving averages or differences. We also use reflective boundary conditions to limit distortions of the wavelet coefficients near the series boundaries, following Faria & Verona (2018).

2.2.2 Out-of-sample Forecasting

For each predictor x_t we estimate the predictive regression:

$$r_{t+1} = \alpha + \beta x_t + \varepsilon_{t+1}$$

where r_{t+1} is the excess return, α captures the mean of excess return, β captures how much the predictor alters the expectation for the next month's return and ε_{t+1} is a forecast error.

Once each predictor has been decomposed into HF, BCF, and LF components, we proceed with forecasting the long-term government bond return (LTR) in an out-of-sample (OOS) setting. Specifically, we adopt the recursive (expanding) window approach, starting with an initial sample from January 1973 to December 1989 to estimate the parameters and obtain a one-step-ahead forecast for January 1990. We then enlarge the sample by one month (i.e., now including data up to January 1990), re-estimate the model, and predict LTR for February 1990. We repeat this process up to the last available date in the dataset, yielding a sequence of one-step-ahead forecasts from January 1990 through December 2023.

This approach ensures that each forecast only uses the information that would have been available at the time of prediction, closely mirroring a realistic investment scenario in which past data become progressively available but future data remain inaccessible.

2.2.3 Forecast evaluation

To gauge the quality of the forecasts, my methodology employed the out-of-sample $R^2 (R^2_{OS})$ measure popularized by Campbell & Thompson (2008). This statistic compares the mean squared forecast error (MSFE) of the predictive regressions ($MSFE_{PRED}$) against that of a simple historical mean (HM) forecast ($MSFE_{HM}$). Formally:

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

Where $\hat{r}_{t+1}$ is the forecast from the predictive model for $t+1$, r_{t+1} is the realized value of the bond return over t to $t+1$, and r_{t+1} is the forecast produced by the historical mean approach. A positive R^2_{OS} means that our predictive regression model yields a smaller MSFE that the simple HM forecast, thus indicating an improvement over the benchmark.

To test the statistical significance of any improvement in MSFE, the Clark & West (2007) adjusted statistics is used. It is based on the null hypothesis H_0 stating that the MSFE of the historical mean is less than or equal to that of the predictive regression, versus the alternative H_1 that the predictive model indeed yields a lower MSFE. Rejection of H_0 at usual confidence levels (e.g. 1%, 5% or 10%) supports the predictive regression's superiority in MSFE terms. Formally:

$$\begin{cases} H_0 = R^2_{OS} \leq 0 \\ H_1 = R^2_{OS} > 0 \end{cases}$$

2.2.4 Certainty Equivalent Return

Beyond purely statistical comparisons, we examine the economic value of our forecasts. To do so, we consider a mean–variance–optimizing investor who allocates a portion of their wealth to the long-term bond based on each period's forecast. Specifically, at the end of month t, the share of wealth invested in the bond is:

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

where γ is the investor's coefficient of relative risk aversion (RRA), $\hat{r}_{t+1}$ is the forecasted bond return from t, $t + 1$ and $\hat{\sigma}^2_{t+1}$ is the predicted variance of bond returns over that same interval (estimated, for example, using a rolling window of past returns). Consistent with prior work (e.g., Almadi et al. (2014)), the portfolio weight w_t is typically constrained between some lower and upper bounds to avoid extremely large, short or leveraged positions.

The certainty equivalent return (CER) is understood as the guaranteed annual return an investor would accept rather than engaging in a risky strategy. Formally, if $\bar{r}_p$ is the average realized return of the portfolio and σ^2 is its variance, then:

$$CER = \bar{r}_p - \frac{1}{2} \times \gamma \times \sigma^2$$

By comparing the CER under our predictive model to the CER under the historical mean forecast, we obtain the annualized CER gain. This gain can be interpreted as the fee an investor would be willing to pay to have access to the forecasts from the predictive model, as opposed to relying solely on the historical average approach. If the predictive model yields higher CER, it implies it can deliver economic benefits to investors who use it in their allocation decisions.

Chapter 3

3. Results

This chapter presents the out-of-sample results for the predictors and compares them to a historical benchmark. The first section (3.1 Out-of-sample Performance) focusses on statistical performance, reporting the out-of-sample R_{OS}^2 for each predictor across time-domain (TS) and the wavelet-decomposed components (HF, BCF, LF). The last section (3.2 Economic performance and Performance Analysis) evaluates economic performance via Certainty Equivalent Return (CER) gains. After that we explore robustness via investor risk aversion, showing table for $\gamma = 1$, $\gamma = 3$ (baseline) and $\gamma = 5$. Lastly portfolio weight constrains, ranging from more lenient bounds ($-1 \leq w_t \leq 2$) to conservative ($0 \leq w_t \leq 1$).

3.1 Out-of-sample Performance

In Table 1 is reported the R2 for each predictor across four versions: the original time series, the high-frequency wavelet component, the business-cycle frequency component and the low-frequency component.

	R_{OS}^2			
Variable	**TS**	**HF**	**BCF**	**LF**
AAA	-1,356	-15,989	0,327	-0,557
BAA	-1,181	-20,901	-0,038	-0,496
lty	-1,417	-16,314	0,299*	-0,604
Corp	-1,114	-0,429	-23,794	-3,183
d/p	-0,640	0,158*	-0,949	-0,513
d/y	-0,380	-1,641	-1,138	-0,484
e/p	-0,266	-5,880	-0,126	-0,599
b/m	-0,169	1,750**	-0,310	-0,511
tms	1,089**	0,253	0,529*	-1,231
dfy	-0,955	-2,589	-2,065	-0,209
dfr	-0,966	-1,759	-2,858	-3,312
infl	0,609	-0,284	-5,317	-0,514
wtexas	-1,799	-0,084*	-5,450	-0,733
ygap	-0,542	-22,267	-0,619	-0,714

Table 1: Out-of-sample R-squared period 1973-2023. This table presents out-of-sample forecasting results for bond returns using fourteen macroeconomic variables. For each predictor, we report the R-squared values obtained from the untransformed time series (TS) and from three wavelet-decomposed components, namely high-frequency (HF), business-cycle frequency (BCF), and low-frequency (LF). Asterisks (*, **, ***) denote significance at the 10%, 5%, and 1% levels, respectively.

All variables, except the term spread and inflation, have negative R-squared values, so they fail to beat the historical mean in the time series. The term spread showing a result of 1,089% with a significance level of 5%. In a more moderate level inflation presents a result of 0,609%.

For the high-frequency, there are tree variables that outperform the historical mean. They are the book-to-market, that provides the strongest result of the table with 1,75% at a significance level of 5%, the dividend-to-price ratio with modest result of 0,158%, significant at 10%, and the term spread with a positive result of 0,253%.

In the business-cycle we have three variables that outperform the historical mean- They are AAA bond yield, long-term government bond yield and the term spread. Of the three the best one is the term spread with a result of 0,529% and 10% significance level.

The low-frequency component fails to provide any positive R-squared values, therefore they all fail to beat the historical mean.

In short, wavelet decomposition sometimes unlocks hidden frequencies that outperform the benchmark. However, the results are not consistent or very significant and not all variables benefit from decomposition, for example BAA bonds yield or corporate bond return.

3.2 Economic Performance

While $\boldsymbol{R^2_{OS}}$ captures statistical gains, we also measure economic value using Certainty Equivalent Return (CER) gains. A positive CER gain means an investor using that forecast strategy achieves higher utility-adjusted returns than an investor using the historical mean forecast. In this topic we will explore different scenarios to understand the model's economic performance changes with different risk aversion and portfolio's weights.

3.2.1 CER gains for Baseline Risk Aversion

We first consider an investor with a relative risk aversion of tree ($\gamma = 3$), along with the original portfolio weight bounds ($-0,5 \leq w \leq 1,5$). The following table displays the CER gains (in percent) for each predictor under the time series (TS) and the wavelet-decomposed components (HF, BCF, LF).

	CER gains ($\gamma = 3$)			
Variable	**TS**	**HF**	**BCF**	**LF**
AAA	-2,723	-2,811	0,842	-0,111
BAA	-2,413	-3,314	0,387	-0,079
lty	0,500	-0,979	-0,955	-4,412
Corp	0,306	-0,227	-0,483	-3,006
d/p	-1,494	0,285	-1,027	0,980
d/y	-0,484	0,245	-1,182	0,976
e/p	0,006	-1,726	-0,219	1,046
b/m	0,266	1,539	-0,782	1,009
tms	0,827	0,634	-0,048	0,689
dfy	-0,965	-0,839	-1,650	0,020
dfr	-1,009	-0,918	-0,703	-2,041
infl	0,881	-0,761	-1,402	1,338
wtexas	0,221	0,601	-0,374	0,117
ygap	0,052	-1,330	-0,543	1,052

Table 2: This table displays the annualized Certainty Equivalent Return (CER) gains for individual bond-return forecasts, assuming an investor with a relative risk aversion coefficient of 3 and portfolio weight constraints set between $-0{,}5 \leq w \leq 1{,}5$. Results are evaluated in its original time-domain form (TS) and across three wavelet-decomposed frequency components—high-frequency (HF), business-cycle frequency (BCF), and low-frequency (LF).

The table above shows the results for the investor with a medium risk profile. In the time series the results are mixed with some positive and other negative. The variables that stand out for the positive are inflation and the term spread, with 0,881% and 0,827% respectably, showing that even without frequency separation they can give signals that help a moderate risk-adverse bond investor achieve higher risk-adjusted returns compared to the benchmark.

In the high-frequency observations we can find a term spread that remains positive 0,634% and the standout of the table with book-to-market presenting a significant jump from 0,266% (TS) to 1,539% (HF). This indicates that short-term variations in book-to-market may capture exploitable movements in the bond market.

The Business-cycle frequency presents the most negative results of the analysis with just AAA bonds yield positive, 0,842%, and BAA bonds yield, 0,387%.

In contrast we have low-frequency with several observations go from negative or negligible in the time series to strongly positive. The stand outs are stock to bond yield gap with 1,052%, and with a second strong appearance book-to-market with 1,009% and inflation with 1,338%. This indicates that slow moving trend in these variables may help to capture long term shifts in bond markets.

3.2.2 CER gains for Different Risk Aversion

This section is dedicated to understanding the sensitivity of the model to different levels of risk aversion. In order to represent an investor that is neutral to risk we define $\gamma = 1$, and for the risk adverse investor, $\gamma = 5$. This idea direct us to do analyze the different CER gains for each type of investor.

	CER gains ($\gamma = 1$)			
Variable	**TS**	**HF**	**BCF**	**LF**
AAA	-5,105	-5,307	0,029	-0,303
BAA	-4,253	-5,043	-0,708	-0,311
Ity	-5,275	-4,130	-0,220	-0,241
Corp	-0,497	-0,508	-1,549	-4,386
d/p	-2,441	-0,714	-2,389	0,169
d/y	-0,982	-1,399	-2,515	0,169
e/p	-0,151	-3,696	-0,330	0,169
b/m	-0,109	0,467	-1,229	0,169
tms	-1,074	0,414	-1,251	-1,165
dfy	-1,276	-1,550	-3,705	-0,315
dfr	-1,519	-1,489	-1,779	-4,047
infl	-0,191	-0,715	-3,609	0,169
wtexas	-1,410	-0,074	-1,605	-1,065
ygap	-0,035	-3,118	-1,146	0,169

Table 3: This table displays the annualized Certainty Equivalent Return (CER) gains for individual bond-return forecasts, assuming an investor with a relative risk aversion coefficient of 1 and portfolio weight constraints set between $-0{,}5 \leq w \leq 1{,}5$. Results are evaluated in its original time-domain form (TS) and across three wavelet-decomposed frequency components—high-frequency (HF), business-cycle frequency (BCF), and low-frequency (LF).

With the new scenario, presented in the table above, the results demonstrate a negative dominance in the time series, with all observation negative and some of them, like AAA bond yield, being strongly negative -5,105%. For the high-frequency the dominance of negative results is still present but this time with two variables showing positive results, but they are not very strong, the book-to-market with 0,467% and the term spread with 0,414%.

The story for the business-cycle frequency is the same as for the time series no results positive or significant. In the low-frequency is similar but slightly better with small but uniform positive gains of 0,169% for several predictors d/p, d/y, e/p, b/m, infl and ygap. This pattern suggests that the low-frequency signals are not significant for risk adverse investors.

The next table displays the second scenario, an investor who is risk adverse. This is represented with a risk aversion level of $\gamma = 5$. We still maintained a portfolio weight between $-0{,}5 \leq w \leq 1{,}5$.

	CER gains ($\gamma = 5$)			
Variable	**TS**	**HF**	**BCF**	**LF**
AAA	-1,753	-2,948	0,451	-0,759
BAA	-1,543	-3,530	-0,231	-0,672
Ity	-1,886	-2,045	0,483	-0,846
Corp	-0,433	-0,480	-1,492	-3,114
d/p	-0,964	0,004	-1,295	0,135
d/y	-0,652	-0,063	-1,407	0,117
e/p	-0,518	-1,453	-0,204	0,196
b/m	-0,277	0,988	-0,747	0,173
tms	0,298	0,354	-0,237	0,351
dfy	-0,994	-0,974	-1,734	-0,197
dfr	-0,956	-0,958	-0,957	-1,867
infl	0,336	-0,743	-1,362	0,259
wtexas	0,024	-0,070	-1,086	-0,643
ygap	-0,719	-1,163	-0,958	0,235

Table 4: This table displays the annualized Certainty Equivalent Return (CER) gains for individual bond-return forecasts, assuming an investor with a relative risk aversion coefficient of 5 and portfolio weight constraints set between $-0{,}5 \leq w \leq 1{,}5$. Results are evaluated in its original time-domain form (TS) and across three wavelet-decomposed frequency components—high-frequency (HF), business-cycle frequency (BCF), and low-frequency (LF).

This scenario overall behave like the previous one with a large majority of negative results across the board. For starter the predictors in the time series are all negative except for tree, inflation with a positive result of 0,336%, term spread with 0,298% and oil prices with a very small result of 0,024%. In the wavelets components we the high-frequency with interesting result with the book-to-market predictor making a jump to 0,988% indicating that short-term fluctuations in book-to-market can reduce drawdowns or capture profitable trades for a highly risk-averse investor. The term spread remains positive with 0,354%, a small improvement. The business-cycle frequency results are very similar to the time series, but with two catches a term spread and inflation turning negative and an improvement for AAA bond yield with 0,451% and Long-term government bond yield with 0,483% This suggests that medium-horizon changes in credit or yield variables may help a cautious investor avoid losses. Lastly the low-frequency performs the best, in term of most variables positive, but do not have strong results. The performance of them all is the term spread only with 0,351%.

When we compare the three tables for risk aversion coefficients $\gamma = 1$, $\gamma = 3$, and $\gamma = 5$, an interesting pattern emerges regarding the magnitude of positive CER gains. At low risk aversion ($\gamma = 1$), most predictors remain negative or only weakly positive, while a few (book-to-market and the term spread in their high-frequency components) yield modest gains, these improvements tend to hover below 0,5%. Because the investor is nearly risk-neutral, large positive returns would be necessary to clearly outperform the benchmark and only a small number of wavelet signals appear to provide that.

By contrast, at moderate risk aversion ($\gamma = 3$), several variables achieve notably higher positive CER values. For instance, book-to-market in its high-frequency component and inflation or earnings to price ratio in their low-frequency forms surpass 1,0%, outperforming anything seen at $\gamma = 1$. This suggests that balancing some variance penalty with the ability to profit from stable signals allows moderate investors to extract the highest utility gains from wavelet-decomposed predictors.

Finally, at high risk aversion ($\gamma = 5$), we see positive results in select frequencies for book-to-market, term spread, or certain predictors, but generally below 1,0%. In other words, while these signals still help more cautious investors, the intensifying penalty on volatility caps the upside. Hence, although strongly risk averse portfolios do avoid some negative outcomes, they also do not reach the peak positive CER gains seen with moderate risk aversion. Overall, the biggest gains from wavelet-based strategies emerge for $\gamma = 3$, indicating that a balance between return-seeking and volatility aversion most effectively exploits the frequency specific signals in these bond-return predictors.

3.2.3 CER gains with Portfolio Weight Constrains

In this final part of the Results section, we focus on how portfolio weight constraints influence the economic value (CER gains) of wavelet-based predictors. We adjust the permissible range for portfolio allocations, allowing for just long positions ($0 \leq w \leq 1$), moderate bounds ($-0{,}5 \leq w \leq 1{,}5$), or more aggressive leverage ($-1 \leq w \leq 2$). Then examine how each scenario affects an investor's ability to exploit the signals revealed through wavelet decomposition. This analysis shows whether loosening or tightening the constraints substantially alters the predictability gains observed in prior sections, thus providing insight into how practical limitations on portfolio construction might enhance or diminish the utility benefits of these frequency-based strategies.

Variable	CER gains ($-1 \leq w \leq 2$) TS	HF	BCF	LF
AAA	-2,938	-5,046	0,837	-1,095
BAA	-2,580	-5,871	0,045	-0,972
lty	-3,142	-3,380	0,831	-0,978
Corp	-0,446	-0,708	-1,753	-4,746
d/p	-1,611	0,005	-1,666	0,623
d/y	-1,023	0,055	-1,964	0,603
e/p	-0,643	-2,798	-0,340	0,699
b/m	-0,315	1,698	-1,014	0,658
tms	0,589	0,585	-0,096	0,793
dfy	-1,299	-1,516	-2,285	-0,209
dfr	-1,266	-1,348	-1,194	-2,444
infl	0,880	-1,239	-1,973	1,020
wtexas	0,359	0,417	-1,370	-0,469
ygap	-0,459	-2,125	-1,355	0,631

Table 5: This table displays the annualized Certainty Equivalent Return (CER) gains for individual bond-return forecasts, assuming an investor with a relative risk aversion coefficient of 3 and portfolio weight constraints set between $-1 \leq w \leq 2$. Results are evaluated in its original time-domain form (TS) and across three wavelet-decomposed frequency components—high-frequency (HF), business-cycle frequency (BCF), and low-frequency (LF).

Table 5, that is presented above, show us the results for a moderate risk averse investor ($\gamma = 3$) whose portfolio weights range from -1 to 2. This expanded bound allows for more aggressive leverage or deeper short positions that than a typical scenario, represented by a portfolio weight ranging from -0,5 to 1,5 (table 2). With a portfolio range between -1 and 2 the investor can short up to 100% of the capital and lever up until 200%. In theory this action should amplify the returns, but if a predictor or frequency band introduces noise or misalignments, losses can also become more severe. For the time series the only non-negative results are the term spread, inflation and oil prices. The best of the bunch is inflation with 0,88%. For the high-frequency the picture improves with the book-to-market performing strongly with a result of 1,698%. The news comes with business-cycle and the low-frequency where some frequencies are not amplified like it was expected in theory. In BCF the long-term government bond yield becomes positive with 0,831%, when in the base line scenario, it was negative. Contrary to theory, in the LF, some observations are not amplified but soften, example of this is the dividend yield, book-to-market and stock to bond yield gap.

The next table represents the most conservative portfolio with a portfolio range of $0 \leq w \leq 1$.

	CER gains ($0 \leq w \leq 1$)			
Variable	**TS**	**HF**	**BCF**	**LF**
AAA	-2,158	-1,448	0,462	-0,034
BAA	-1,915	-1,618	-0,037	-0,003
Ity	-2,230	-1,132	0,516	0,024
Corp	0,234	0,040	-0,273	-1,822
d/p	-0,873	0,192	-0,965	0,534
d/y	-0,381	-0,175	-0,989	0,529
e/p	0,089	-1,105	-0,017	0,534
b/m	0,172	0,501	-0,791	0,534
tms	0,321	0,265	-0,385	0,116
dfy	-0,561	-0,664	-1,530	-0,081
dfr	-0,574	-0,614	-0,416	-1,630
infl	0,516	-0,472	-0,873	0,584
wtexas	-0,288	0,143	-0,402	-0,084
ygap	0,124	-0,963	-0,114	0,492

Table 6: This table displays the annualized Certainty Equivalent Return (CER) gains for individual bond-return forecasts, assuming an investor with a relative risk aversion coefficient of 1 and portfolio weight constraints set between $0 \leq w \leq 1$. Results are evaluated in its original time-domain form (TS) and across three wavelet-decomposed frequency components high-frequency (HF), business-cycle frequency (BCF), and low-frequency (LF).

The results in this portfolio represent a very different picture, probably due to the fact that the investor cannot short positions and beneficiate with negative signals. We have for every frequency and even in the time series variables turning positive of negative and some smoothing with other amplify, compared to the original portfolios weights.. Examples for the time series is oil prices turning negative in this new scenario (0,221% to -0,288%), proving that without shorting the economics gains fall. An example in the high-frequency is the dividend yield with a negative result of -0,175% and previously of 0,285%, which can be explained by the presence of noise and limiting the exposed capital improves the economic gains. In the business-cycle frequency we have long-term government bond yield with another example of jumping from negative to positive, but this time with more significance (-0,955% to 0,516%), even when compared to the time series the jump is interesting because it goes from -2,23% to 0,516%, showing a modest economic gain for a conservative portfolio. For the low-frequency we see a typical case of smoothing for all variables, the one that stand one is the dividend price ratio with a modest result of 0,534% and decent improvement to the time series (-0,873%).

Adjusting portfolio weight constraints profoundly influence the economic value of wavelet-based predictors. Under the expanded range $-1 \leq w \leq 2$, investors can short more heavily or lever up to twice their capital, amplifying both positive and negative CER outcomes. While certain signals (e.g., book-to-market in high-frequency) produce substantially higher returns than in the baseline, others exhibit riskier swings that diminish utility if the model's forecasts deviate from market realities. Conversely, the strictly long-only setup $0 \leq w \leq 1$ restricts downside risk from misaligned signals but also limits the upside from potentially strong predictors or frequency bands. Notably, some previously negative or negligible results become positive in the long-only context such as long-term government bond yields in the business-cycle frequency indicating that fewer short positions and simpler allocation rules can reduce forecast noise for certain predictors. However, other signals that benefitted from shorting or leverage (like oil prices or dividend yield at certain frequencies) show diminished gains. In essence, while flexible weight bounds can magnify the returns from wavelet-based strategies, a conservative approach preserves modest but more consistent improvements. These findings underscore the importance of matching an investor’s leverage and short-selling capacity to their confidence in specific wavelet components, ultimately allowing the frequency-domain signals to be harnessed most effectively within each investor’s practical portfolio constraints.

Conclusion

This thesis sets out to investigate the predictive power of bond-return predictors through the lens of wavelet-based frequency decomposition. While traditional forecasting models often rely on time-domain analyses, wavelet methods have the potential to reveal hidden cyclic and long-term behaviors by isolating high-frequency, business-cycle, and low-frequency signals from each predictor. By combining these frequency-specific insights with a standard out-of-sample forecasting framework, we aimed to determine whether various indicators such as term spreads, bond yields, dividend ratios, and inflation could produce meaningful improvements in forecasting long-term government bond returns.

A central focus was the comparison of raw time series forecasting versus decomposed forecasting, where each predictor was split into three wavelet-based components. Empirically, we found that certain variables, such as the term spread and book-to-market, exhibited notably higher R^2_{OS} and CER gains once their most relevant frequencies were isolated. For instance, short-term fluctuations in Book-to-market often boosted forecast accuracy and improved economic value for the investor, whereas the time series of the same variable sometimes underperformed the simple historical mean benchmark. This contrast underscores the key insight that predictors may carry valuable signals only at specific frequencies, which are missed by standard approaches that treat the data as a single undifferentiated time series.

In addition to comparing frequencies, the thesis explored how out-of-sample forecast performance translates into economic value via CER gains. We used a mean-variance framework in which an investor allocates to a long-term bond based on the predicted excess return and forecast variance. The risk aversion parameter (γ) proved crucial in interpreting results. Highly risk-averse investors

($\gamma = 5$) tended to show moderate positive CER gains in only a few well-matched frequency components, since their strong variance penalty limited the benefits of more volatile signals. On the other extreme, near risk-neutral investors ($\gamma = 1$) generally required very strong return improvements to yield a net positive CER, resulting in mostly modest or negative outcomes except for a handful of short-run signals such as book-to-market or term spread in high frequency. The middle ground ($\gamma = 3$) consistently produced the most pronounced CER gains overall, suggesting that balancing a moderate penalty on volatility with a readiness to exploit stable frequency specific trends can extract the largest benefits from wavelet-based forecasts. These conclusions agree with the findings in Faria & Verona (2020), who also show that wavelet-based filtering can enhance the predictive power of bond-return models by uncovering short-term and long-term signals that traditional time-domain methods overlook.

We also tested different portfolio weight constraints, allowing short positions and leverage to varying degrees. The scenarios ranged from an aggressive $-1 \leq w \leq 2$ range where an investor could short up to 100% of capital or lever gains up to 200% to a conservative long-only constraint $0 \leq w \leq 1$. Unsurprisingly, looser constraints amplified both upsides and downsides. Certain wavelet-filtered signals, such as book-to-market in high frequency or inflation in low frequency, generated notably higher CER gains under leverage. Conversely, if a predictor misaligned with market movements, deeper short positions or stronger leverage worsened the losses. The long-only case showed some previously negative or negligible results become positive in the long-only context, indicating that fewer short positions and simpler allocation rules can reduce forecast noise for certain predictors. However, other signals that benefitted from shorting or leverage show diminished gains. This indicates that portfolio weight constraints modulate how wavelet forecasts translate into economic outcomes. An investor's practical tolerance for short sales and leverage can either unlock greater potential returns or protect against large drawdowns, depending on the reliability of the frequency band in question.

From a broader perspective, this study highlights how wavelet decomposition can transform otherwise unremarkable or marginal bond-return predictors into valuable signals once their relevant time scales are isolated. It also clarifies that investor-specific factors, like risk aversion and allowable portfolio weights, can either magnify or mute the resulting benefits. Thus, the key takeaway is that bond return forecasting can profit from sophisticated frequency domain techniques, but only if the implementation is aligned with the investor's objectives and risk preferences.

Looking forward, there remain open avenues for further exploration. One potential direction is to combine wavelet components across multiple predictors, effectively blending the strongest signals in high-frequency, business-cycle-frequency, and low-frequency ranges from several variables. Further research could explicitly model trading frictions, bid-ask spreads, and liquidity constraints, as discussed by Welch & Goyal (2008) in the equity context and Almadi et al. (2014) for dynamic asset allocation. Accounting for realistic transaction costs will clarify whether wavelet-based signals retain net economic value once rebalancing costs are factored in, especially at higher frequencies.

The present thesis demonstrates that wavelets, when carefully integrated with standard asset allocation frameworks, can enhance both predictive accuracy and economic value for bond-return forecasts, contingent on an investor's risk appetite and feasible allocation constraints.

Bibliography

Aguiar-Conraria, L., Martins, M. M. F., & Soares, M. J. (2012). The yield curve and the macro-economy across time and frequencies. *Journal of Economic Dynamics and Control*, *36*(12), 1950–1970. https://doi.org/10.1016/j.jedc.2012.05.008

Almadi, H., Rapach, D. E., & Suri, A. (2014). Return predictability and dynamic asset allocation: How often should investors rebalance? *Journal of Portfolio Management*, *40*(4). https://doi.org/10.3905/jpm.2014.40.4.016

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

Bianchi, D., Büchner, M., & Tamoni, A. (2021). Bond Risk Premiums with Machine Learning. *Review of Financial Studies*, *34*(2). https://doi.org/10.1093/rfs/hhaa062

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). https://doi.org/10.1093/rfs/hhm055

Carr, P., & Madan, D. (1999). Option valuation using the fast Fourier transform. *The Journal of Computational Finance*, *2*(4). https://doi.org/10.21314/jcf.1999.043

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

Cochrane, J. H., & Piazzesi, M. (2005). Bond risk premia. In *American Economic Review* (Vol. 95, Issue 1, pp. 138–160). https://doi.org/10.1257/0002828053828581

Crowley, P. M. (2007). A guide to wavelets for economists. *Journal of Economic Surveys*, *21*(2), 207–267. https://doi.org/10.1111/j.1467-6419.2006.00502.x

Dangl, T., & Halling, M. (2012). Predictive regressions with time-varying coefficients. *Journal of Financial Economics*, *106*(1), 157–181. https://doi.org/10.1016/j.jfineco.2012.04.003

Fama, E. F., & Bliss, R. R. (1987). The Information in Long-Maturity Forward Rates. *The American Economic Review*, *77*(4).

Faria, G., & Verona, F. (2018). Forecasting stock market returns by summing the frequency-decomposed parts. *Journal of Empirical Finance*, *45*. 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*. https://doi.org/10.1016/j.finmar.2019.100508

Ferreira, M. A., & Santa-Clara, P. (2011). Forecasting stock market returns: The sum of the parts is more than the whole. *Journal of Financial Economics*, *100*(3), 514–537. https://doi.org/10.1016/j.jfineco.2011.02.003

Gallegati, M., & Ramsey, J. B. (2013). Bond vs stock market's Q: Testing for stability across frequencies and over time. *Journal of Empirical Finance*, *24*, 138–150. https://doi.org/10.1016/j.jempfin.2013.10.003

Gargano, A., Pettenuzzo, D., & Timmermann, A. (2019). Bond Return Predictability: Economic Value and Links to the Macroeconomy. *Management Science*, *65*(2), 508–540. https://doi.org/10.1287/mnsc.2017.2829

Gençay, R., Selçuk, F., & Whitcher, B. (2003). Systematic risk and timescales. *Quantitative Finance*, *3*(2), 108–116. https://doi.org/10.1088/1469-7688/3/2/305

Gençay, R., Selçuk, F., & Whitcher, B. (2005). Multiscale systematic risk. *Journal of International Money and Finance*, *24*(1), 55–70. https://doi.org/10.1016/j.jimonfin.2004.10.003

Gençay, R., Whitcher, B., & Selçuk, F. (2002). *Robustness of systematic risk across time scales*. https://www.researchgate.net/publication/267264906

Johannes, M., Korteweg, A., & Polson, N. (2014). Sequential learning, predictability, and optimal portfolio returns. *Journal of Finance*, *69*(2), 611–644. https://doi.org/10.1111/jofi.12121

Kim, S., & In, F. (2005). The relationship between stock returns and inflation: New evidence from wavelet analysis. *Journal of Empirical Finance*, *12*(3), 435–444. https://doi.org/10.1016/j.jempfin.2004.04.008

Ludvigson, S. C., & Ng, S. (2009). Macro factors in bond risk premia. *Review of Financial Studies*, *22*(12). https://doi.org/10.1093/rfs/hhp081

Percival, D. B., & Walden, A. T. (2000). Wavelet Methods for Time SeriesAnalysis. In *Wavelet Methods for Time SeriesAnalysis*. https://doi.org/10.1017/cbo9780511841040

Ramsey, J. B., & Lampart, C. (1998). Decomposition of economic relationships by timescale using wavelets: Money and income. *Macroeconomic Dynamics*, *2*(1). https://doi.org/10.1017/s1365100598006038

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

Rua, A. (2017). A wavelet-based multivariate multiscale approach for forecasting. *International Journal of Forecasting*, *33*(3), 581–590. https://doi.org/10.1016/j.ijforecast.2017.01.007

Shiller, R. J. (1991). Yield spreads and interest rate movements: A bird's eye view. *Review of Economic Studies*, *58*(3), 495–514. https://doi.org/10.2307/2298008

Welch, I., & Goyal, A. (2008). A comprehensive look at the empirical performance of equity premium prediction. *Review of Financial Studies*, *21*(4). https://doi.org/10.1093/rfs/hhm014

Appendix

Summary Statistics					
Variable	**Mean**	**Medium**	**1st percentile**	**99th percentile**	**Standard Deviation**
AAA	0,0714	0,0715	0,0243	0,1461	0,0287
BAA	0,0822	0,0802	0,0328	0,1680	0,0307
Ity	0,0633	0,0623	0,0079	0,1384	0,0309
Corp	0,0067	0,0066	-0,0638	0,0813	0,0280
d/p	0,0276	0,0215	0,0115	0,0572	0,0125
d/y	0,0277	0,0218	0,0115	0,0573	0,0126
e/p	0,0624	0,0537	0,0119	0,1376	0,0290
b/m	0,4511	0,3336	0,1322	1,1256	0,2796
tms	0,0192	0,0200	-0,0189	0,0437	0,0151
dfy	0,0108	0,0095	0,0057	0,0268	0,0044
dfr	0,0002	0,0005	-0,0482	0,0400	0,0153
infl	0,0032	0,0030	-0,0054	0,0129	0,0038
wtexas	0,0097	0,0016	-0,2216	0,3226	0,0989
ygap	-2,9377	-2,9685	-4,4574	-2,0760	0,4500

Table 7: Reports the mean, median, first-percentile, 99th-percentile and standard deviation of each monthly predictor used in the analysis. The sample period ranges from 1973 to 2023. The number of observations is 612.

AI Declaration

During the preparation of my written work/thesis, Master's Final Assessment, Chat GPT was used for the following tasks: generating summaries, data analysis, refining language, code analysis, reviewing content and academic research with the prompts used listed at the end of the document in the Prompts List section. After using this tool/service, I reviewed and edited the content as necessary, and I take full responsibility for the content of the work presented.

I also declare that I am aware of and respect the Artificial Intelligence Rules of Conduct of Católica Porto Business School.

Prompts List Section

- “Find the error in this MATLAB code”
- “Explain me, with the maximum detail, the following MATLAB code”
- “what the error in this one?”
- "explain me in this paragraph like i am sone one who do not know nothing about the topic”
- "Don’t do anything yet I’m just uploading files to future research”
- "Is there any academic papers that talk about the Fourier methods?"
- "find me one that I can citate in the text I send you earlier"
- "rewrite this to be in full text with an appropriate academic English”
- "do a small resume of the paper”
- "does this paper talk about bonds?"
- “rewrite the methodology, but do not cut or summarize anything. Just correct the English”
- “Make it more clear”
-“Can you articulate better English in these sentences”
- "I will give you the methodology and you will say the points that need to be corrected or improved ok?"
- “Find new research papers”
- “Can you give examples of good predictors to test in the model that I gave you above?”
- “What do you think about oil prices as a predictor?”
- “And the stock to bond yield gap?”
- “And the default yield spread?”
- “Can you analyze the results for the MATLAB code (it's in the pdf)?”
- “what do you think about comparation the results for different levels of risk adversion?”
-“ this is my thesis, what is your opinion on it?”