Master's thesis

Copenhagen Business School

Date of submission: 15 May 2019

Time-varying Stock-Bond Correlation: Creating Minimum-Variance Portfolios Using Analyst Forecasts

Supervisor: Andreas Brøgger

Characters incl. spaces: 232,896

Number of pages: 112

Nichlas Rasmussen

Student no.: 92179

MSc in Economics and Business Administration - Finance and Investments

Magnus Blyitgen Jacobsen

Student no.: 94230

MSc in Economics and Business Administration - Accounting, Strategy and Control

Abstract

Correlation between stock and bond returns is of immense importance as it plays a vital role in investors' diversification and asset allocation decisions. The purpose of this thesis is to explain the driving forces behind volatilities and correlation of stock and bond returns and to investigate how this knowledge can be used to form portfolios that outperform traditional asset allocation strategies.

This paper takes a forwarding-looking approach by using analyst forecasts of macroeconomic variables to predict future realized volatilities and correlation of stock and bond returns. Variables from relating literature that have exhibited predictability of co-movement of returns are identified and analyst forecasts of these variables are used in predictive regressions. The factors used to forecast volatilities and correlation are implied stock market volatility, inflation rate, short and long rate, as well as change in corporate profits and change in real GDP. Mean consensus as well as dispersion in analysts' forecasts of the above-mentioned economic variables are applied in the analysis. To scrutinize the predictive power of the forecast variables we are controlling for historical data on the same variables.

An in-sample predictability analysis reveals that predictive regressions using analyst forecasts outperform, in terms of adjusted R^2, models using historical data only, but a specification including both historical and forecast variables perform even better. Additionally, the single best in-sample predictor of future realized volatilities and correlation is simply the lag, that is, the previously realized value. Results show that all mean consensus and dispersion variables are statistically significant in predicting volatilities and stock-bond correlation except mean consensus of corporate profits, and dispersion in forecasts of real GDP growth.

Next, an out-of-sample predictability analysis is conducted to examine how minimum-variance portfolios formed using analyst forecasts perform compared to two simple benchmark strategies. Our results show that a strategy using both historical data and analyst forecasts yield the best performance, which is in line with the results from the in-sample analysis. Almost all proposed strategies perform statistically better than the equally-weighted benchmark portfolio, however all strategies fail to provide evidence for statistical and economic outperformance compared to a simple moving-average strategy. Additionally, out-of-sample, the use of analyst forecasts seems to unlock some predictability, as strategies excluding the lag perform better. Several robustness checks confirm that analyst forecasts of macro variables can be used to improve investors ex-ante allocation of wealth between stocks and bonds compared to an equally-weighted strategy but fails to outperform a simple moving-average strategy.

Table of Contents

Abstract **1**

1 Introduction **5**

1.1 Background and motivation 5

1.2 Research question 7

1.3 Contribution to existing literature 8

1.4 Delimitations 8

1.5 Outline 9

2 Literature review **10**

2.1 Constant or time-varying correlation 10

2.2 Forecasting using historical data 11

2.3 Forecasting using analyst survey data 12

2.4 Determinants of stock-bond correlation reported in the literature 13

2.4.1 Using dispersion in analyst forecasts as determinants 14

3 Theory **16**

3.1 Risk reduction through diversification 16

3.2 The Markowitz portfolio optimization model 17

3.2.1 Mean-variance analysis with risky assets only 17

3.2.2 Mean-variance analysis with a risk-free and risky assets 20

3.2.3 The optimal portfolio given the investor's risk aversion 22

3.3 The capital asset pricing model 24

3.4 Risk aversion and utility functions 26

3.4.1 Utility functions used for mean-variance analysis 26

3.5 Performance measures 27

3.5.1 Sharpe ratio 27

3.5.2 Certainty equivalent 28

3.5.3 Portfolio turnover 30

3.5.4 Opportunity cost 31

4 Data and methodology 32
4.1 Data description and treatment 32
4.1.1 Stock and bond return data 32
4.1.2 Analyst forecast data 33
4.1.3 Historical data 35
4.2 In-sample and out-of-sample analysis 35
4.3 Ordinary least squares 36
4.3.1 OLS assumptions 37
4.3.2 Durbin-Watson d Test 38
4.3.3 White Test 39
4.4 Calculation of evaluation metrics 41
4.4.1 Mean return 41
4.4.2 Standard deviation 42
4.4.3 Sharpe ratio 42
4.4.4 Certainty equivalent 42
4.4.5 Portfolio turnover 43
4.4.6 Opportunity cost 43
4.5 Statistical significance testing of portfolio performance 44
5 Analysis and empirical findings 46
5.1 Stylized facts and the co-movement of stock and bond returns 46
5.1.1 Realized versus forecasted volatilities and correlation 48
5.1.2 Analyst forecasts 51
5.2 In-sample predictability of macroeconomic variables 53
5.2.1 Predictive regressions of stock-bond correlation 54
5.2.2 Predictive regressions of stock volatility 61
5.2.3 Predictive regressions of bond volatility 66
5.2.4 Performance of the predictive regressions 70
5.3 Out-of-sample predictability analysis 71
5.3.1 Hypothesis testing for difference in portfolio performance 72
5.3.2 Portfolio strategies to be considered 74
5.3.3 Benchmark strategies 74
5.3.4 Portfolios based on the historical data only 77
5.3.5 Portfolios based on analyst forecast data only 79
5.3.6 Portfolios based on historical and analyst forecast data 82
5.3.7 Other strategies considered 84
5.3.8 Overall out-of-sample predictability performance 86

5.4 Robustness check....88
5.4.1 OLS assumptions....88
5.4.2 Test for autocorrelation....88
5.4.3 Test for homoscedasticity....89
5.4.4 Comments and consequences of the OLS assumption tests....90
5.4.5 Different levels of risk aversion....91
5.4.6 Different length of estimation window....95
5.4.7 Different length of holding period and quarterly data....98

6 Conclusion and discussion 102

6.1 Concluding remarks on empirical results....102
6.1.1 Comparing findings to similar research....105
6.2 Practical validity of our findings....106
6.2.1 Practical validity of the equally-weighted benchmark strategy....106
6.2.2 Validity of using minimum-variance portfolios....107
6.2.3 Overall validity of the empirical results....109
6.3 Asset liability management: Pension funds....109
6.4 Future research....112

Bibliography i

Appendices vi

Appendix A....vi
Appendix B....vi

1 Introduction

1.1 Background and motivation

For many years it was widely accepted to assume constant negative correlation between stocks and bonds, which also serves as the argument for equally-weighted portfolios suggested by investment guru Benjamin Graham in his book from 1949, *The Intelligent Investor.* However, newer research has proved that there exist considerable time-variation in the co-movement of stock and bond returns (Gulko, 2002; Li, 2002; Connolly, Stivers and Sun, 2005).

While it may not be obvious to all investors, getting the correlation right is of immense importance as it plays a vital role in the investors' diversification and asset allocation decisions. Therefore, the purpose of this thesis is to explain the economic forces driving stock-bond correlation and investigate whether this knowledge can be used to form portfolios of stocks and bonds that outperform conventional asset allocation strategies.

Traditionally, academics and practitioners have used simplifying methods such as averages of realized correlations or complex statistical models to estimate expected correlation. All of these methods use historical data as input to the correlation estimation. This thesis takes a forward-looking rather than backward-looking approach as part of the thesis will focus on unveiling whether analyst forecasts have predictive power of the second moments of stock-bond returns.

Focusing on stocks and bonds as the only two asset classes allows us to simplify the analysis, as more asset classes significantly increases the number of inputs that must be estimated. Additionally, stocks and bonds make up, by far, the largest share of all traded financial assets globally with 45% and 36%, respectively, see figure 1. Hence, stocks and bonds are a good proxy for the entire investment opportunity set available to most investors.

Figure 1 - Split of total capital invested in major financial asset classes

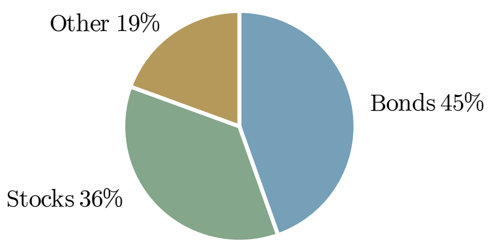

Source: (Attaluri, 2014)

The motivation to predict co-movement of stock and bond returns, alongside the related benefits to the investor, is illustrated in figure 2.

As standard deviation of portfolio returns is used as a proxy for risk, there is a clear difference between the two portfolios in the figure below. Panel A displays the realized monthly standard deviation of a naive equally-weighted portfolio and a minimum-variance portfolio, which is based on perfect foresight of the stock-bond variance-covariance matrix that ultimately provides the best possible input for portfolio formation. The later strategy represents a portfolio formed in month t using the realized volatilities and correlation of stock-bond returns from month t+1. Consider a fictional fund that invests half of its wealth into bonds and the other half into stocks, thereby ignoring the time-variation in stock-bond correlation in contrast to a fund that invests with perfect foresight of the correlation process. It is clear from the figure, that perfectly forecasting correlation between the two asset classes have great benefits both in terms of reducing the overall level but also the variation in portfolio risk.

Figure 2 - Realized volatilities and Sharpe ratios of equally-weighted and perfect foresight minimum-variance portfolios over the entire sample period Jul-1994 to Dec-2018

Panel A: Realized monthly standard deviation of equally-weighted and perfect foresight min. variance portfolios

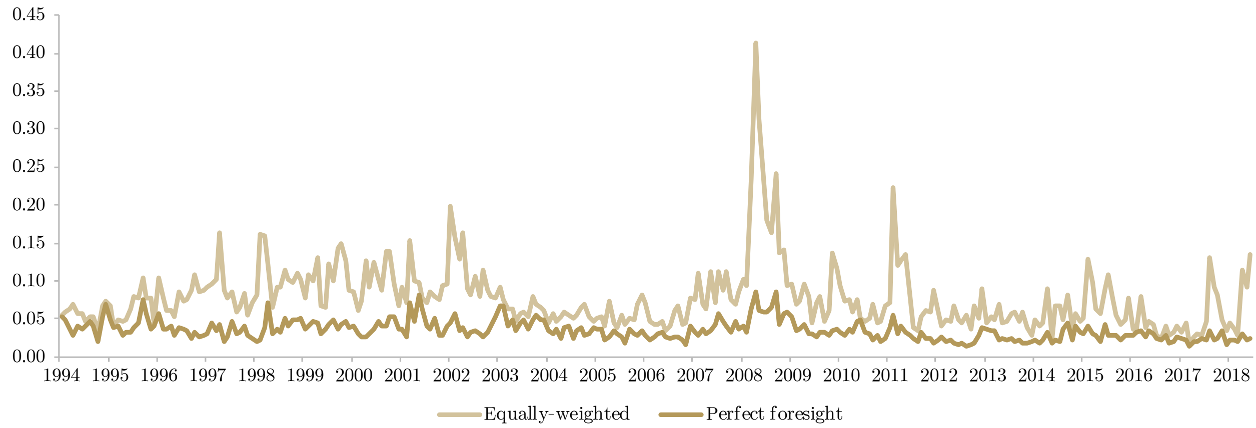

Panel B: Annualized Sharpe ratios of the two portfolio strategies and asset classes

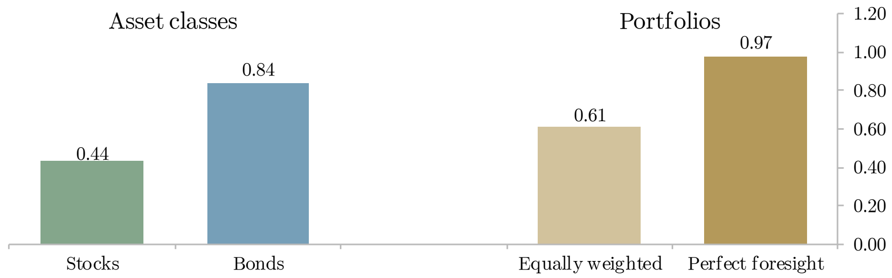

Source: Own contribution

Panel B displays Sharpe ratios of the two portfolio strategies, alongside the Sharpe ratios of the individual asset classes themselves, over the entire sample period from July 1994 to December 2018. Being able to forecast the correlation perfectly would have yielded a Sharpe ratio of 0.97, which is considerably better than the Sharpe ratio of the equally-weighted portfolio of 0.61. Interestingly, an investor would have been better off simply investing 100% of her wealth in bonds compared to the equally-weighted portfolio strategy, as it yielded a Sharpe ratio of 0.84.

1.2 Research question

This paper studies the economic factors driving the co-movement of stock and bond returns and investigates how knowledge about this process can be used to formulate superior asset allocation strategies. Therefore, the aim of the thesis is to provide an answer to the following research question:

What are the determinants behind future volatilities and correlation of stock and bond returns? Do forward-looking analyst forecasts of these variables have any explanatory power? Using information about these factors, is it possible to construct minimum-variance portfolios, which outperform benchmark asset allocation strategies?

The research question will be approached by answering several sub-questions. Firstly, previous literature will be scrutinized with the aim of finding variables that have proven to have statistical power in explaining variation in co-movement of stock and bond returns. The focus of the thesis is to use analyst forecasts of the variables found in the literature, and to investigate whether this forward-looking data adds explanatory power over and above simple historical information on the same variables. To support this part of the research question, hypothesis one will be tested.

Hypothesis 1

H_0: The collection of analyst forecast variables, $\boldsymbol{F}_t$ and $\boldsymbol{D}_t$, do not explain variation in volatilities or correlation of stock and bond returns

Where $\boldsymbol{F}_t$ *and* $\boldsymbol{D}_t$ are the collections of mean-consensus and dispersion variables, respectively. Detailed information follows in the data and methodology section.

The second part of the research question will be answered by creating minimum-variance portfolios using the predictive regressions developed to answer hypothesis one. The performance of these portfolios will be compared to benchmark portfolio strategies. Hypothesis two is tested to answer this part of the research question.

Hypothesis 2

H_0: The performance of minimum-variance portfolios formed using the predictive regressions is worse than or equal to benchmark strategies

Where the different benchmark strategies will be described in the analysis and empirical findings section and the performance measures in the data and methodology section. The results from the tests of hypothesis one and two should ultimately provide a concluding answer to the research question.

1.3 Contribution to existing literature

This study of the benefits of using analyst forecasts in determining stock-bond correlation, and in turn portfolios weights, contributes to the existing literature in several aspects. Firstly, the thesis provides an extensive literature review on factors explaining the second moments of stock and bond returns and uses these insights to investigate new specifications of predictive regressions, with combinations of regressors that have not been tested before.

Secondly, the thesis contributes with results based on most recent return data with the last observation being December 2018. Hence, the study includes the interesting low interest rate regime period that has prevailed since the latest financial crisis in 2007. Since the sample period stretches over several different macroeconomic states, the results may be considered more general than similar studies that only consider specific macroeconomic environments.

Additionally, focusing on marginal predictive power of forecast variables, the results will show whether analyst forecasts add any explanatory power over and beyond historical information. Recent studies, using analyst forecasts in predicting stock-bond correlation, have not controlled for or investigated whether forecast data is better than historical data. This thesis adds that element to the analysis. Overall, the thesis contributes to the narrow range of existing research and acts as a paper for comparison and confirmation of the findings in existing literature.

1.4 Delimitations

As mentioned earlier in the introduction, the thesis considers two asset classes only, these being stocks and bonds. The decision to do so is motivated by the fact that having only two asset classes to invest in, means that there is only one pair-wise correlation to forecast, which significantly simplifies the analysis.

Under the hypothesis that risk and hence correlation is time-varying, this paper aims at presenting a model that can predict next month's variance-covariance matrix of stock and bond returns. This matrix can be used as input to modern portfolio theory's mean-variance optimization tool as to form mean-variance efficient portfolios. However, in order to do so, it is also necessary to forecast expected returns for both stocks and bonds over the same period and this exercise has proven to be very difficult. Frameworks such as the CAPM or Fama and French's three-factor model attempt to predict equity returns but empirical tests of the models seem to reveal high statistical uncertainty and poor performance (DeMiguel, Garlappi and Uppal, 2009).

Therefore, given limited statistical validity in prediction of returns and since the emphasis of this paper is on predicting volatilities and correlation of stock and bond returns, the paper focuses on obtaining global minimum-variance portfolios (GMVP). The optimization problem of finding minimum-variance portfolios only requires the variance-covariance matrix as input. Hence, GMVPs offer a way to obtain mean-variance efficient portfolios without making any assumptions about future returns.

1.5 Outline

The remainder of the thesis will be structured as follows. Section two includes a review of the existing literature, investigating previous research methods and presents their empirical results. It is in this section that all macroeconomic variables used for the predictive regressions are discovered. Section three provides a presentation of the theories underlying this research such as mean-variance analysis and different performance measures.

Methodology used for the empirical analysis and a discussion of what and how data is treated follows in section four. Section five presents with the analysis and empirical findings. Firstly, stylized facts about stock and bond returns as well as volatilities and correlation are presented. This is followed by an in-sample analysis of the predictive power of analyst forecasts using data from the entire sample period. Next, a rolling-window regression analysis explores the out-of-sample benefits of analyst forecasts in predicting stock-bond correlation and compares the performance of portfolios formed using such a model to more traditional asset allocation strategies. Also, several robustness checks of the results are presented. Lastly, section six presents concluding remarks, a discussion of practical implications for asset allocation, and relates the results in context of the general literature.

2 Literature review

Stocks and bonds are two of the most important asset classes for ordinary investors. Hence, the correlation between stock and bond returns plays a pivotal role in the investors' diversification and asset allocation decisions. Of the investment management process (asset allocation, market timing, and security selection), research has shown that asset allocation policy is the most dominant contributor to total return and explains up to 90% of the variation in portfolio performance (Brinson, Hood and Beebower, 1986; Brinson, Singer and Beebower, 1991). This suggests that the optimal allocation of wealth between stocks and bonds is one of the most important decisions an investor faces.

Given the importance of asset allocation and that stocks and investment grade bonds account for a dominant share of all traded financial assets, one would expect that researchers have already found the answer to the immensely important question of what drives the co-movement of stock and bond returns. However, despite its great importance, the dynamics behind the correlation process seem to remain elusive. The problems and difficulties that researchers face are illustrated by Baele, Bakaert and Inghelbrecht's (2010) attempt to build a backward-looking empirical model, which tries to forecast realized stock-bond correlation. Their dynamic factor model using fundamental macroeconomic variables fails to forecast the correlation.

2.1 Constant or time-varying correlation

In many years researchers and practitioners assumed the co-movement of the two asset classes to be constant across time. In the first edition of investment guru Benjamin Graham's book, *The Intelligent Investor*, he suggests an equally-weighted portfolio of stocks and bonds based on the claim that stock-bond correlation is constant and negative. More recently, Shiller and Beltratti (1992) and Campbell and Ammer (1993) use the same framework, a dynamic present value model, to decompose the variances and covariances of monthly stock and bond returns in the United States. These two studies both implicitly assume stock-bond correlation to be constant.

Newer research has moved in a different direction by acknowledging and investigating the time-varying nature of stock-bond correlation. Scruggs and Glabadanidis (2003) tests the assumption of constant correlation empirically, by building a model which imposes a constant correlation restriction on the covariance matrix between stock and bond returns. This model is strongly rejected. Using data from the US, UK, and Germany, Andersson, Krylova and Vähämaa (2008) investigate the correlation between stocks and bonds in each of these countries. Estimating the monthly correlation from daily return observations over a time horizon of approximately 15 years shows that the correlation varies significantly over the period with sustained periods of both positive and negative correlation. Additionally, they find that the stock-bond correlation may change substantially over short periods of time, which may pose serious challenges for the asset allocation task of investment managers. In a similar fashion, Ilmanen (2003) examines historical US stock-bond correlation from 1926-2001. He finds that stocks and bonds are positively correlated most of the time but considerable time variation, including sustained periods of negative correlation, exists.

Connolly, Stivers, and Sun (2005) study time variation in co-movement of stock and bond returns using daily observations. They posit the same conclusions as the researchers above including a long list of other papers, that is, there is substantial time variation in the relation between stock and bond returns in the short-term with sustained periods of negative correlation (Gulko, 2002; Li, 2002; Fleming, Kirby and Ostdiek, 2003).

Since the literature has recognized a time-varying nature of stock-bond correlation and given the importance of asset allocation policy in portfolio performance, researchers have built models trying to forecast the pair-wise correlation of the two asset classes. The motivation is clear, if the one-period-ahead correlation can be known ex-ante, then investors are better off as they are able to form more efficient portfolios. The literature for forecasting correlation can be divided into two groups, models that employ historical information and models that take a forward-looking approach by incorporating analyst forecast data from surveys.

2.2 Forecasting using historical data

Researchers have proposed several models using historical information, some more complex than others. The simplest models include rolling historical correlation and exponential smoothing. These two methods are often used by practitioners due to their simplicity. More complex volatility models have been developed over the years and include the autoregressive conditional heteroskedasticity (ARCH) model (Engle, 1982), generalized autoregressive conditional heteroskedasticity (GARCH) model (Bollerslev, 1986), and the Dynamic Conditional Correlation (DCC) model (Engle, 2002). These statistical models for time series data are modeling variance in the error terms as a function of previously realized error terms. Due to their complexity, only the most sophisticated practitioners use these methods to model correlation.

Examples of the abovementioned complex statistical methods to model stock-bond correlation includes Cappiello, Engle and Sheppard (2006) that use Engle's DCC model to develop an econometric technique to measure risk dynamically, by finding the optimal time decay of stock-bond covariance. Additionally, de Goeij and Marquering (2004) employ daily return data on US stock and bond indices in their multivariate GARCH model, and find strong support of conditional heteroskedasticity in the correlation between stocks and bonds, that is, non-constant correlation. Their multivariate GARCH model shows good performance and reveals low covariance between stock and bond returns after bad news in the stock market and good news in the bond market.

2.3 Forecasting using analyst survey data

Recent literature has turned its focus to the impact of expectations of macroeconomic variables on the time-varying co-movement of stock and bond returns (Andersson, Krylova and Vähämaa, 2008; Baele, Bakaert and Inghelbrecht, 2010; Jivraj and Kosowski, 2011; Jivraj, 2012b). The use of expectations, through survey data instead historical data, may be more appropriate as both stocks and bonds are priced based on future expectations of several variables, not their realized values.

Intuitively, analyst forecasts of macroeconomic variables should have good predictive power as the analysts forecast a given variable based on both current and forecasted macroeconomic regimes. Hence, embedded within these variables are implied information about the future state of the macroeconomic environment. This point is illustrated by Piazzesi, Salomao and Schneider (2011), who posit that analyst forecast, ex-ante, should provide better out-of-sample predictions than the abovementioned complex statistical models, which forecast using historical information. They further argue that these models only perform well in-sample, that is, given the benefit of perfect hindsight.

Examples of research on time-varying stock-bond correlation that employ analyst forecast data include the work by Andersson, Krylova, and Vähämaa (2008). The researchers use data from the US, UK, and Germany and use two estimates of the cross-asset correlation. They deduce three determinants of varying correlation from asset pricing theory. These include expected inflation, economic growth expectations, and expected stock market uncertainty, where inflation and GDP are measured by analyst forecasts and stock market uncertainty is implied from option pricing. The results show that expected inflation is positively related to the time-varying correlation between stock and bond returns, which means that stock and bond prices tend to move in the same direction in periods of high inflation. The empirical findings further suggest that expected stock market uncertainty is negatively related to the co-movement of stock and bond returns. This finding is in accordance with the "flight-to-quality" phenomenon, where people shift portfolio weights to bonds instead of the riskier stocks. Lastly, they do not find any systematic relationship between expected market growth and the pair-wise correlation between stocks and bonds.

Jivraj (2012) performs a similar test on US data only. Firstly, the author performs an in-sample test of the carefully selected variables' ability to predict stock and bond returns' second moments. He finds varying degrees of predictability even after controlling for the lag, that is, the previously realized value. On the contrary, he presents out-of-sample benefits of using analyst forecasts to predict stock-bond correlation, as the empirical results suggest that investors form more efficient portfolios using the forecasted correlation compared to using a prediction based on historical data.

2.4 Determinants of stock-bond correlation reported in the literature

Several determinants of the co-movement of stock and bond returns have been proposed by researchers. This section gives an overview of the most frequently used variables within the pertaining literature.

One macroeconomic variable that may impact the stock-bond correlation, in theory, is inflation. An increasing inflation rate will cause the common discount rate to rise, which undoubtedly impacts bond holders negatively. On the other hand, the impact on stock holders is uncertain as stock returns are negatively impacted by the discount rate but positively impacted by the higher future expected cash flows. A lot of studies have used survey forecasts of inflation as an explanatory variable to stock-bond correlation (Andersson, Krylova and Vähämaa, 2008; Baele, Bakaert and Inghelbrecht, 2010; Jivraj, 2012b). In a study using US data, Ilmanen (2003) finds that in periods of high inflation, the discount rate effect dominates the cash flow effect, thereby increasing the stock-bond correlation.

Likewise, variables representing treasury rates, both short and long rates are often included in the regressions. Intuitively, the treasury rates should have some predictive power in the co-movement of stock and bond returns, which is also confirmed by the results of David and Veronesi (2008) and Viceira (2012). They find that the short rate is statistically significant in predicting stock-bond correlation and has a positive effect. Their results suggest that the short rate captures a procyclical component in the time variation of the second moments of bond returns. Viceira (2012) investigates the yield spread on long-term and short-term bonds. He argues that this serves as a proxy for the current business conditions, hence having some predictive power of the time-variation in stock and bond return co-movement.

Several researchers have included real variables, such as real GDP growth, in their analyses (Campbell and Ammer, 1993; Li, 2002; Andersson, Krylova and Vähämaa, 2008). A variable related to the growth of the economy should intuitively be able to explain some variation in security prices.

Earnings growth has also been put forward as a possible determinant (David and Veronesi, 2008; Jivraj, 2012b). The motivation behind using corporate earnings growth is related to a well-known valuation model, Gordon's growth model. In this model growth of corporate earnings is a key input, and therefore it may seem reasonable that this variable is able to explain some of the variation in the stock-bond correlation.

Not only fundamental changes in the macroeconomic environment may have an impact on the relationship between stock and bond returns. Change in investors' assessment of the current market risk and changes in the financial market dynamics may also be important factors to consider. This is evident in periods of financial turmoil, where investors demand a higher risk premium to hold stocks relative to holding bonds. A consequence of this is the "flight-to-quality" phenomenon, which is an event characterized by large portfolio shifts from stocks to bonds in periods of high market uncertainty. This mechanism decouples the returns of the two asset classes.

The pertaining literature is trying to capture this effect by including a measure of stock market uncertainty. The most frequently used measure is implied volatility from option pricing using the VIX index (Baele, Bakaert and Inghelbrecht, 2010; Kostakis, Panigirtzoglou and Skiadopoulos, 2011; DeMiguel *et al.*, 2013), but also other measures such as stock turnover have been investigated (Connolly, Stivers and Sun, 2005). In a study that focuses on stock-bond correlation around market crashes, Gulko (2002) finds that periods of negative pair-wise correlation often coincides with stock market crashes. Similarly, Connolly, Stivers and Sun (2005) suggest that option-implied stock market volatility is a good predictor of financial market turbulence, as bond returns seem to be high relative to stocks in periods with high stock market uncertainty. These two observations are in line with the "flight-to-quality" phenomenon.

Lastly, other less frequently used determinants of stock-bond correlation are a liquidity and risk aversion measures (Baele, Bakaert and Inghelbrecht, 2010), as well as dummy variables for the current business cycle (David and Veronesi, 2008).

2.4.1 Using dispersion in analyst forecasts as determinants

Motivation for using an uncertainty or dispersion measure of forecasts of macroeconomic variables follow from Li's (2002) study of stock-bond correlation in G7 countries. He finds that uncertainty about the long-term expected inflation rate plays an important role in determining the pair-wise correlation of stock and bond returns. The greater the uncertainty about the inflation rate, the stronger co-movement between stock and bond returns. Additionally, Li shows that uncertainty about other macroeconomic variables such as the real interest rate also has explanatory power, but to a lesser degree.

In a more recent study, David and Veronesi (2008) find similar results, that is, dispersion in analyst forecast of the inflation rate is able to forecast the realized covariance between the two asset classes, and even significantly so. Lastly, Jivraj (2012) shows that, dispersion in analyst forecast of both the short and long rate, as well as corporate earnings and real GDP growth all significantly explain the time-varying correlation between stocks and bonds.

3 Theory

3.1 Risk reduction through diversification

Investors typically hold more than one asset. The reason is to obtain diversification benefits that is provided from forming a portfolio. The return of a portfolio is generally less risky comparing to the return of a single security that makes up a part of the portfolio. The motivation seems clear, if one asset performs extremely badly over a given period, there is a great chance that the other assets perform relatively better, thereby limiting overall poor performance. However, the mechanics work the opposite way as well, posing a lower chance for an exceptional portfolio return compared to the chance of an exceptional return on a single security.

Although high returns are preferable, it is not the only thing that matters to investors. An investor may be willing to give up some return, if she can reduce the risk of her position considerably. Hence, it becomes a question of finding the portfolio with the most attractive risk-return trade-off. In the two-asset setting presented in this thesis, that is, only considering stock and bond indices as the investment universe, there are still considerable diversification benefits to gain. This is illustrated in figure 3, which shows that the standard deviation of the two-asset portfolio return, for some given portfolio weights, is less than the standard deviation of stocks and bonds in isolation. Deriving the variance of portfolio returns, one will find that the standard deviation of the portfolio is less than the weighted average of the asset's standard deviations. This illustrates the benefit of diversification, as the investor is able to reduce risk by forming portfolios (Munk, 2017).

Figure 3 - Standard deviation of a stock-bond portfolio

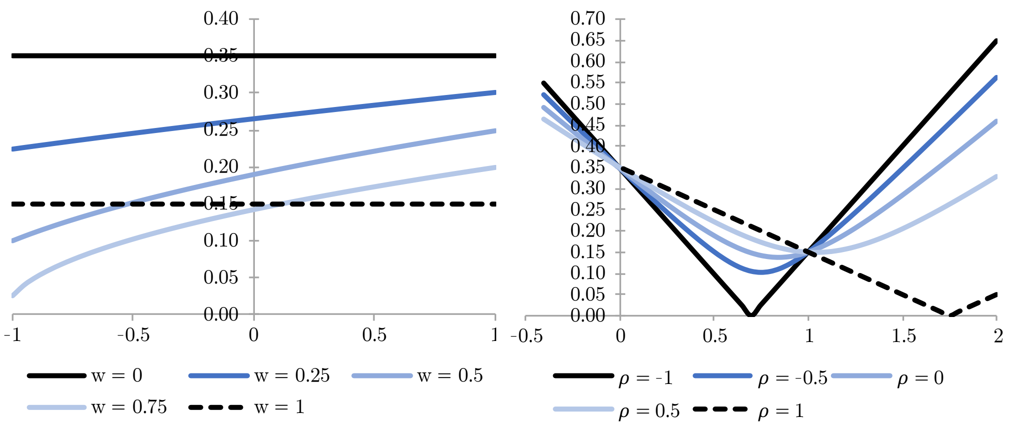

(a) Portolio std. dev. as a function of asset correlation

(b) Portolio std. dev. as a function of weight on asset 1

Source: Own contribution

The two graphs in figure 3 also underlines the importance of correlation on portfolio risk and why being able to predict the pair-wise correlation between stocks and bonds is of uttermost importance to investors.

3.2 The Markowitz portfolio optimization model

The section above illustrated the benefits of investing in a portfolio of securities to reduce the overall riskiness. Lower portfolio risk comes with lower expected returns, hence it is not only a question of minimizing variance of the investor's portfolio, but rather to find the optimal trade-off between risk and return, which in turn also depends on the investor's risk aversion. The most well-known framework to handle the optimization problem of the risk-return trade-off is called Mean-Variance Analysis, and was introduced by Nobel-laureate Harry Markowitz (Markowitz, 1952, 1959) in order to determine the optimal set of portfolio weights.

One of the assumptions and key ideas behind the framework is, that investors only consider variance and expectation of portfolio returns over a fixed period of time, when choosing the optimal portfolio. Additionally, the investor prefers as low portfolio variance as possible and as high return as possible. This, in combination, means that the investor is a mean-variance optimizer. The framework presented below is generalized to encompass portfolios of N risky assets, hence is also applicable to the simpler two-asset scenario presented in the analysis section of this thesis.

The portfolio optimization problem can be solved in three steps. First, the mean-variance optimizer needs to establish the optimal risk-return combinations available from the set of all risky assets. Next, using optimization techniques, the best available portfolio of risky assets must be identified, given the possibility of investing in a risk-free asset as well. Lastly, the final portfolio must be determined by appropriately mixing the risk-free asset and the optimal risky portfolio relative to the investor's degree of risk aversion (Bodie, Kane and Marcus, 2014).

3.2.1 Mean-variance analysis with risky assets only

In the following sections, $\boldsymbol{\pi}$ is the portfolio weight vector, which indicates what fraction of wealth is invested in each asset. $\boldsymbol{\mu}$ denotes the expected return vector and $\underline{\Sigma}$ is the variance-covariance matrix of expected rates of returns. Since all portfolio weights must sum to one, the vector must satisfy $\boldsymbol{\pi} * \mathbf{1} = 1$.

Given investors are mean-variance optimizers, they will only choose among mean-variance efficient portfolios. A portfolio is said to be mean-variance efficient if it has the lowest possible standard deviation among all portfolios with the same expected return. Hence, the first step is to trace out mean-variance efficient portfolios at all levels of mean return. This will create the mean-variance efficient frontier of risky assets, as depicted in figure 4.

Figure 4 - Mean-variance efficient frontier from 10 risky assets and global minimum and max-slope portfolios

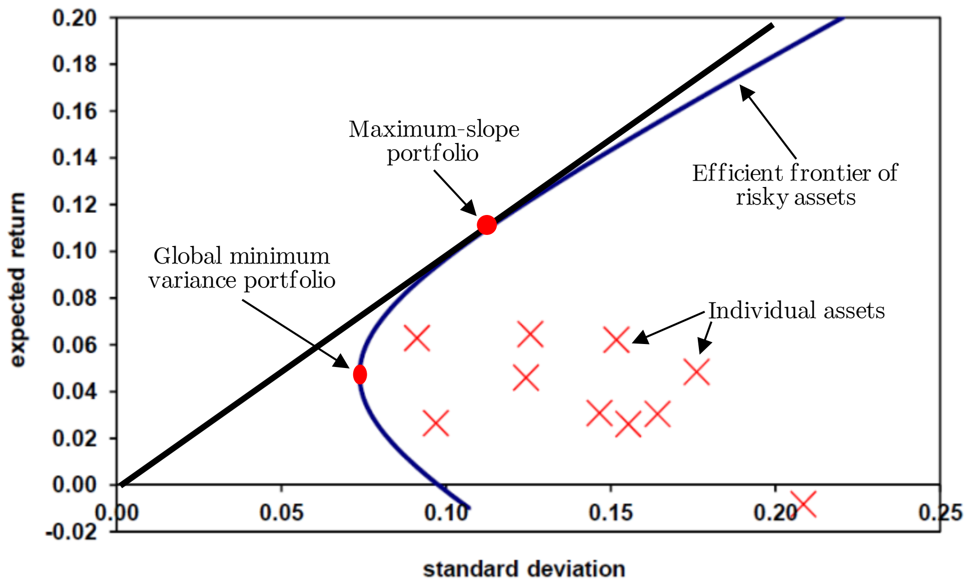

Source: Munk (2017)

There are two ways of generating this frontier. The first method involves finding a mean-variance portfolio with expected return of $\bar{\mu}$ by solving the quadratic minimization problem:

$$\min_{\pi} \boldsymbol{\pi} * \underline{\Sigma}\boldsymbol{\pi}$$

$s.t. \boldsymbol{\pi} * \boldsymbol{\mu} = \bar{\mu}$ and $\boldsymbol{\pi} * \mathbf{1} = 1$

To solve this problem and other optimization problems within mean-variance analysis, it is helpful to define some auxiliary constants, as proposed by Munk (2017). These aid the calculations of the optimal portfolio weights.

$$A = \boldsymbol{\mu} * \underline{\Sigma}^{-1}\boldsymbol{\mu} \qquad B = \mathbf{1} * \underline{\Sigma}^{-1}\boldsymbol{\mu}$$

$$C = \mathbf{1} * \underline{\Sigma}^{-1}\mathbf{1} \qquad D = AC - B^2$$

The solution to the minimization problem above is given by a mean-variance efficient portfolio with expected return $\bar{\mu}$ and portfolio weight vector:

Equation 1

$$\boldsymbol{\pi}(\bar{\mu}) = \frac{C\bar{\mu} - B}{D}\underline{\Sigma}^{-1}\boldsymbol{\mu} + \frac{A - B\bar{\mu}}{D}\underline{\Sigma}^{-1}\mathbf{1}$$

And the standard deviation of the portfolio return is:

Equation 2

$$\sigma(\bar{\mu}) = \sqrt{\frac{C\bar{\mu}^2 - 2B\bar{\mu} + A}{D}}$$

Solving the minimization problem for different values of $\bar{\mu}$ will generate efficient combinations of standard deviation and mean return and will form a hyperbola in a diagram with standard deviation on the x-axis and mean return on the y-axis. This is the efficient frontier of risky assets, see figure 4. Some people refer only to the upward-sloping part of the hyperbola as the efficient frontier, while the rest is called the minimum-variance frontier (Bodie, Kane and Marcus, 2014). The reason is that rational investors will never choose a portfolio on the downward-sloping part, since they will be able to form a portfolio with the same risk but with a higher expected return.

The second method to obtain the mean-variance frontier is to combine the global minimum variance portfolio with the maximum-slope portfolio. Both portfolios are depicted in figure 4. This property is a two-fund separation result, that is, if investors can only form portfolios from the N risky assets, then a mean-variance optimizer will choose a combination of only two special funds, the global minimum-variance and the maximum-slope portfolio (Munk, 2017). Making a large number of combinations of the two portfolios will also generate the efficient frontier of risky assets.

As mentioned above, the minimum-variance portfolio is the portfolio with the lowest standard deviation among all portfolios. The minimization problem is not dependent on expected returns of the assets, only how the assets co-vary with each other. This will be useful later in the analysis. The minimum-variance portfolio is the solution to the constrained minimization problem:

$$\min_{\pi} \boldsymbol{\pi} * \underline{\Sigma}\boldsymbol{\pi}$$

$s.t. \boldsymbol{\pi} * \mathbf{1} = 1$

The global minimum-variance portfolio is given by portfolio weight vector:

Equation 3

$$\boldsymbol{\pi}_{min} = \frac{1}{\mathbf{1} * \underline{\Sigma}^{-1}\mathbf{1}} \underline{\Sigma}^{-1}\mathbf{1}$$

And standard deviation of returns:

Equation 4

$$\sigma_{min} = \frac{1}{\sqrt{C}}$$

One would expect that assets with low volatility receive a large weight in the portfolio. However, covariances are also very important. An asset with high volatility might also receive a large weight if it has very low correlation with other low volatility assets. The reason is that such an asset provides a great potential for diversification.

The second portfolio to be mixed with the global minimum-variance portfolio is called the maximum-slope portfolio. It is the point or portfolio on the efficient frontier with the highest slope, when a line from the origin (0,0) is drawn to the all portfolios. The maximum-slope portfolio is given by the portfolio weight vector:

Equation 5

$$\boldsymbol{\pi}_{max} = \frac{1}{\mathbf{1} * \underline{\Sigma}^{-1}\boldsymbol{\mu}} \underline{\Sigma}^{-1}\boldsymbol{\mu}$$

And standard deviation of returns:

Equation 6

$$\sigma_{max} = \frac{\sqrt{A}}{|B|}$$

Irrespective of the approach being used, the result will be a hyperbola representing combinations of mean return and standard deviation of portfolios that are all mean-variance efficient. Now the investor has established the universe of all efficient portfolios of risky assets.

3.2.2 Mean-variance analysis with a risk-free and risky assets

Investors can also invest some of their wealth in a risk-free (rf) asset. The option to do so has implications for what mean-variance efficient portfolio of risky assets that will be the optimal one. A combination of an investment in the risk-free asset and an investment in a mean-variance efficient portfolio form a straight line from the point (0,rf) to the point (σ, μ), which represents the pair of the standard deviation and mean return for the risky portfolio. The slope of this line is exactly the Sharpe ratio $(\mu - rf)/\sigma$ of the risky portfolio. By extension, the line from (0,rf) that is tangent to the efficient frontier of risky assets, will form the highest Sharpe ratio. The portfolio at the point of tangency is called the tangency portfolio and the line is called the capital market line (CML), see figure 5.

Figure 5 - Mean-variance frontier, the CML, and the tangency portfolio

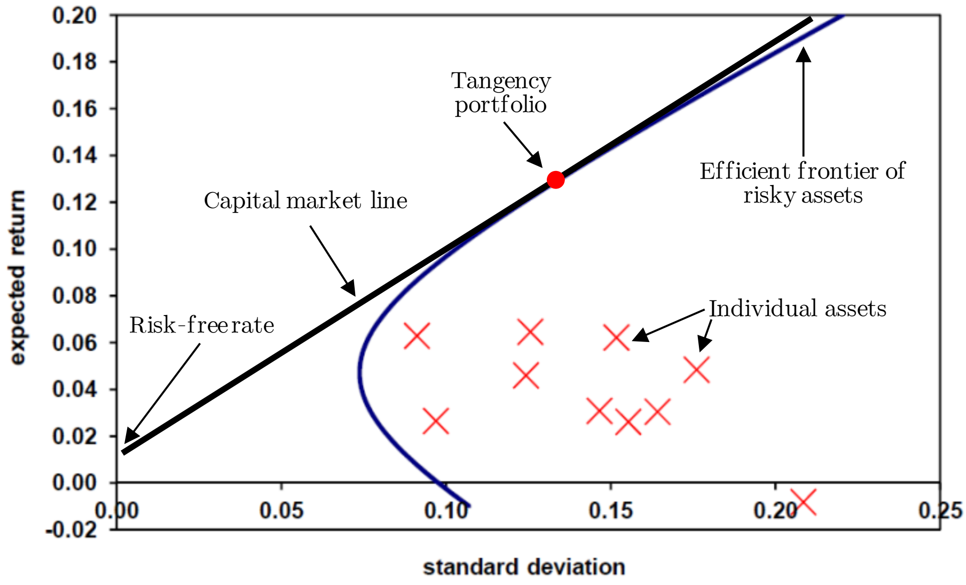

Source: Munk (2017)

The tangency portfolio is characterized by a mean-variance efficient portfolio of risky assets with portfolio weights equal to:

Equation 7

$$\boldsymbol{\pi}_{tan} = \frac{1}{B - C * rf} \underline{\Sigma}^{-1}(\boldsymbol{\mu} - rf * \mathbf{1})$$

And with a standard deviation of portfolio return equal to:

Equation 8

$$\sigma_{tan} = \frac{\sqrt{A - 2B * rf + C * rf^2}}{|B - C * rf|}$$

Since the tangency portfolio is the one maximizing the Sharpe ratio, it will include large portfolio weights on assets with large Sharpe ratios or large weights on assets with low correlation with high Sharpe ratio-assets. All combinations of the risk-free asset and the tangency portfolio make up the efficient frontier of all assets. A mean-variance optimizer will always choose a combination of the tangency portfolio and the risk-free asset, that is, a portfolio on the efficient frontier of all assets, as their preference make the investors want to be as far up North-West in the standard deviation-mean diagram in figure 5. Such a portfolio will have a weight w in the tangency portfolio and (1-w) in the risk-free assets. The combined portfolio will have an expected return and standard deviation of:

Equation 9

$$\mu(w) = w\mu_{tan} + (1-w)rf \qquad \sigma(w) = |w|\sigma_{tan}$$

If all mean-variance investors agree on the same investment opportunities, in terms of expected returns, variances, and covariances of risky assets, then they will all hold the same portfolio of risky assets in some combination with the risk-free asset, reflecting the investor's degree of risk aversion.

3.2.3 The optimal portfolio given the investor's risk aversion

The above sections concluded that a mean-variance optimizer will choose a combination of the optimal risky portfolio and the risk-free asset. Equation 9 states that such a portfolio will have a weight of w in the tangency portfolio and a weight of (1-w) in the risk-free asset. The next challenge is to determine the optimal value of w, and this depends on the investor's mean-variance preference or degree of risk aversion. In order to do so, one needs to establish an explicit expression for the optimal value of w by formalizing the mean-variance trade-off for the investor. As suggested by Munk (2017), one way to formalize this trade-off is to maximize the difference between the expected return on the portfolio minus a constant times the variance of the rate of return.

Equation 10

$$\max\left(E[r] - \frac{1}{2}\gamma Var[r]\right)$$

Where γ is a positive constant and is a proxy for risk aversion. Hence, a larger value of γ corresponds to an increase in penalty of a high portfolio variance, that is, risk. Consequently, low values of γ corresponds to investors with low risk aversion (risk lovers), and likewise a high γ resembles high risk aversion. The mean-variance preference of the investor can be represented by indifference curves in the (σ, μ) diagram. By construction, the investor is indifferent between any point on this line and any two lines correspond to different levels of satisfaction, and these cannot cross. Intuitively, investors want the highest possible indifference curve as this yields the highest expected return for any fixed value of standard deviation. Also, one should expect the indifference curves to be convex, thus becoming steeper as the risk increases. Since the investor is risk-averse by nature, she will demand increasingly higher expected return, as the risk increases. The optimal allocation between the risk-free asset and the tangency portfolio is where the indifference curve is tangent to the CML. A relatively risk-averse investor will have a steep indifference curve, making the optimal combination lie on the lower part of the CML. On the contrary, a risk-lover will have a flat indifference curve, with the optimal portfolio being tangent to the upper part of the CML, maybe even involving a short position in the risk-free asset. These two scenarios are displayed in 6.

Figure 6 - Indifference curves illustrating difference in risk aversion and choice of optimal portfolio on the capital market line

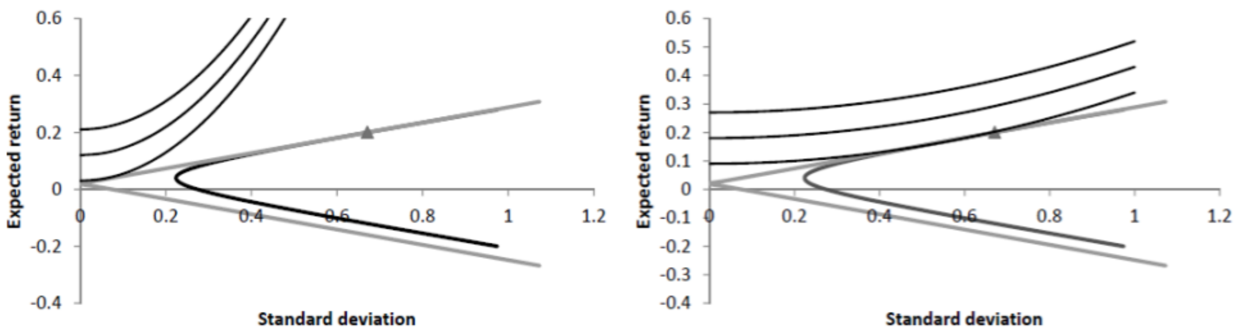

(a) High risk aversion

(b) Low risk aversion

Source: Munk (2017)

With knowledge about the mean and variance of the optimal portfolio from equation 9, the objective function of the maximization problem in equation 10 can be written as:

Equation 11

$$f(w) = \mu(w) - \frac{1}{2}\gamma\sigma(w)^2 = rf + w(\mu_{tan} - rf) - \frac{1}{2}\gamma w^2 \sigma_{tan}^2$$

To solve this, f'(w) is set to zero to find that the objective is maximized for:

Equation 12

$$w^* = \frac{\mu_{tan} - rf}{\gamma\sigma_{tan}^2}$$

Where w^* is the fraction of wealth to be invested in the tangency portfolio. The weight in the risky portfolio is decreasing in the variance and the risk aversion proxy γ, which is to be expected. Now the mean-variance optimizer has found the optimal combination of the risk-free asset and the mean-variance efficient portfolio of risky assets, given her degree of risk aversion.

3.3 The capital asset pricing model

The mean-variance analysis framework presented in the previous section requires a long list of inputs such as expected returns, volatility, and correlations. The inputs in the model were taken as given and it was not until the 1960's that an equilibrium model of asset prices was introduced. The Capital Asset Pricing Model (CAPM) was developed through the collective work of Treynor (1961), William (1964), Lintner (1965), and Mossin (1966).

The model builds upon a number of simplifying assumptions, which also were assumed in Markowitz's mean-variance analysis (Berk and DeMarzo, 2017). The first assumption states that investors can buy and sell assets at competitive prices and do so with no restrictions, that is, without incurring tax or transaction costs. Additionally, investors can lend and borrow at the risk-free rate. The second assumption is that all investors are mean-variance optimizers, so they choose to only hold efficient portfolios that yield the maximum expected return for a given level of risk. Lastly, the third assumption states that investors have homogenous expectations regarding volatilities, expected returns, and correlations of securities.

The fathers of the CAPM realized that if investors agree on the efficient frontier of risky assets and the level of the risk-free rate, then all investors will agree on the composition of the tangency portfolio. And as argued in the section above, mean-variance optimizers will always choose a combination of the tangency portfolio and the risk-free asset, hence all investors invest only in some combinations of these two instruments. In equilibrium, asset prices or equivalently returns, must be set such that investors' total demand is equal to total supply of assets. This means, that the tangency portfolio must include all assets in the economy, effectively making it equivalent to the market portfolio. From this insight, an expression for the expected return on any individual asset can be formulated.

Equation 13

$$E[r_i] = r_f + \beta_i\left[E(r_{mkt}) - r_f\right]$$

where $E[r_i]$ being the expected return on asset *i* and *rf* is the risk-free rate representing the return required when investing in a risk-free security over a period of time. The remaining term represent the return required for taking on further risk. $E(r_{mkt})$ is the expected return of the market portfolio, or equivalently the tangency portfolio. Beta in the formula is the measurement of the systematic risk of asset *i* with respect to the market portfolio. It is defined as:

Equation 14

$$\beta_i = \frac{Cov[r_i, r_{mkt}]}{Var[r_{mkt}]}$$

This implies that the expected return on a single security is a linear combination of the risk-free rate, and the product of the asset's covariance with the market portfolio and the excess market return. This relation is defined as the Security Market Line (SML) and is derived directly from the CAPM formula. In a diagram with expected returns on the y-axis and beta on the x-axis, the SML forms a straight line intercepting the vertical axis at the risk-free rate and has a slope equal to the market risk premium $E[r_{mkt} - r_f]$, see figure 7. Hence, the SML describes the relation between risk and return. It graphs individual assets' expected return as a function of their beta, and assets which are located directly on the SML, are considered to be fairly priced, while assets that are not directly on the SML are mispriced and represent an investment opportunity (Bodie, Kane and Marcus, 2014). The market portfolio will by construction have a beta equal to one.

Figure 7 - Illustration of the security market line and the capital market line. M denotes the market portfolio, whereas X and Y represent individual assets

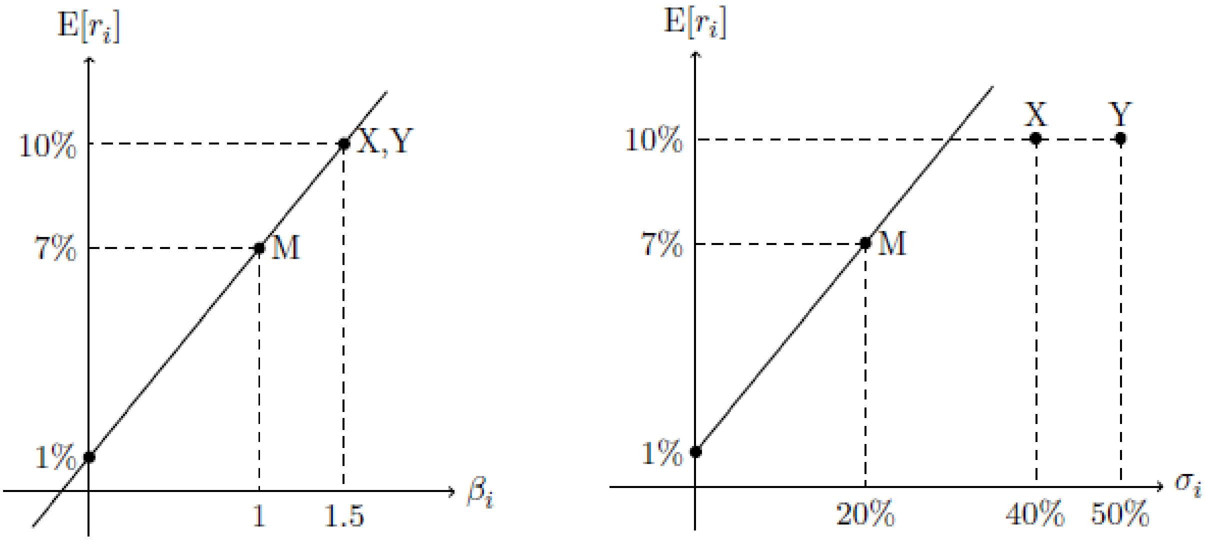

(a) Security market line

(b) Capital market line

Source: Munk (2017)

The link between the SML, described above, and the CML, which is the line that goes from (0,rf) and through the tangency portfolio, or equivalently the market portfolio, is illustrated in figure 7. Correctly priced assets according to CAPM, will lie on the SML. However, only efficient portfolios, hence combinations of the risk-free asset and the tangency portfolio will be located on the CML. As shown in the figure, individual assets X and Y carry some non-systematic risk, and are therefore not located on the efficient frontier.

The SML and CML has the following implications. Firstly, the market portfolio is efficient. Hence, the portfolio with the highest expected return for any given level of risk is located on the CML. Also, risk premia on any investment is proportional to its beta with the market. Therefore, the linear relationship between risk and expected return can be described by the SML.

3.4 Risk aversion and utility functions

The concept of risk aversion goes back centuries, at least as far back as 1738, where mathematician Daniel Bernoulli studied expected utility from a coin-toss game called the St. Petersburg Paradox (Bodie, Kane and Marcus, 2014). Hence, for a very long time, it has been widely accepted that investors are risk averse and demand a risk premium on risky assets.

The preferences of an investor are typically represented by utility functions as already presented in section 3.2.3. It is assumed that a decision-maker can assign a welfare or utility to a certain level of wealth. If the original wealth at the beginning is denoted by W_0 then wealth at the end period, W, can be described as:

Equation 15

$$W = W_0(1 + r)$$

Where r is the portfolio return over the period. Naturally, the return and in turn the end-period wealth depends on the chosen portfolio, hence both the rate of return and ultimo wealth are random variables. A utility function $u(W)$ is then a function that assigns a level of satisfaction or welfare to all possible levels of wealth. The goal of the investor is to maximize the expected value of utility $E[u(W)]$ when picking a portfolio.

Utility functions are assumed to exhibit a number of characteristics (Munk, 2017). The first being that $u'(W) > 0$. This means that the function is increasing, which can be translated into the investor being greedy, that is, the more wealth the better. The second characteristic is that the function is concave, so that $u''(W) < 0$, which in turn means that marginal utility $u'(W)$ is decreasing in wealth. This translates into the fact that an investor appreciates a dollar more highly when poor than when rich. It also means that the investor is risk averse, so that she rejects any risky gamble where expected profit is zero or negative.

3.4.1 Utility functions used for mean-variance analysis

The negative exponential utility function of wealth is often used in mean-variance analysis, as it is consistent with the assumed characteristics described above, also given the assumed mean-variance criterion put forward in equation 10. The negative exponential utility function can be described as:

Equation 16

$$u(W) = -e^{-kW}$$

Where $k > 0$ to represent a greedy and risk averse investor, since $u'(W) = ke^{-kW}$ and $u''(W) = -k^2 e^{-kW}$. This utility function is often referred to as CARA because it exhibits constant absolute risk aversion k. Supposing returns, and hence wealth, are normally distributed, so preferences only depend on mean and variance, it can be shown that:

Equation 17

$$W \sim N(\mu, \sigma^2) \quad \rightarrow \quad E[u(W)] = E[-e^{-kW}] = -E[e^{-kW}] = -e^{-kW_0(1+\mu-\frac{1}{2}kW_0\sigma^2)}$$

Since the function $-e^{-kW_0(1+x)}$ is increasing in x, we have that the portfolio maximizing $E[u(W)]$ also maximizes the assumed criterion $\mu - \frac{1}{2}kW_0\sigma^2$. Hence, the formulized mean-variance trade-off of the investor in equation 10 is consistent with a negative exponential utility function. The risk aversion constant γ is equal to the investors absolute risk aversion k times the initial wealth W_0. This means that an investor with a negative exponential utility function optimizes her utility by investing a fraction of $w^* = \frac{\mu_{tan} - rf}{kW_0\sigma_{tan}^2}$ in the tangency portfolio and the rest in the risk-free asset. This is the exact same conclusion from equation 12 in the previous section.

Other utility functions exists but the negative exponential utility functions is attractive since it is mathematically tractable with normally distributed returns (Munk, 2017).

3.5 Performance measures

In order to evaluate the attractiveness of investment strategies, a multitude of performance measures have been introduced over time. This section describes the performance measures that is used in the analysis. Specifically, this section will describe the Sharpe ratio, the certainty equivalent, portfolio turnover, and opportunity cost.

3.5.1 Sharpe ratio

As touched upon in section 3.3 "The capital asset pricing model", investors are concentrated with obtaining high return investments. However, it is not the case that a high return is unconditionally better than a low one. While excess returns describes the return above the risk free rate, it does not contain any information about the relevant risk. In order to isolate the quality of an investment strategy, various risk-reward measures have been introduced (Pedersen, 2015).

Soon after the introduction of the CAPM, William F. Sharpe introduced one such risk-reward measure (Sharpe, 1966). The performance measure is similar to the Treynor measure (Treynor, 1965) that was introduced one year earlier. The important difference between the two measures is, that Sharpe's measure includes the non-systematic risk. Sharpe argues that diversification is of substantial importance for the performance of portfolios for which reason the non-systematic risk should be included when assessing portfolio managers.

Sharpe′s measure was originally named "reward-to-variability ratio" and was introduced as a measure to evaluate mutual fund performance. Today, it is one of the most utilized performance measures, the generally accepted name has changed to "Sharpe ratio", and the application has extended much further than mutual fund performance. The Sharpe ratio of a portfolio is defined as the investment reward per unit of risk and is formally written:

Equation 18

$$SR = \frac{R_p - R_f}{\sigma(R_p - R_f)}$$

The ratio is visually represented by the slope of straight line in figure 4.

The Sharpe ratio is considered a superior measure for evaluating portfolios than simply comparing average returns. The reason being that the Sharpe ratio corrects for risk and is invariant to leverage.

3.5.2 Certainty equivalent

As described in section 3.3 "The capital asset pricing model", investors are under normal circumstances encountered with the trade-off between high return and low risk. The reason being that investors on aggregate are risk averse and therefore require a premium on risky investments. As a result, in order to be attractive for investors, investment opportunities that carry a high amount of risk must sell cheaper than safe investment opportunities with similar expected return. The market consolidates the combined risk aversion and securities are priced thereafter. Individual investors, however, may have a different attitude toward risk and use the risk-aversion difference to optimize their own utility. The level of risk aversion is therefore important when evaluating investments and we have therefore included it in the form of the performance measure certainty equivalent.

Suppose that a person is given the binary choice of 1) gaining USD 100 for certain versus 2) be given a 50% chance of gaining USD 200 and 50% chance of gaining USD 0. If the person is risk averse, she will prefer the first choice. Now lower the value of the first choice until the person is indifferent. Say indifference is reached when the first choice is lowered to USD 75. This is then defined as the certainty equivalent of the above gamble for the particular person

(Hershey and Schoemaker, 2008). The USD 25 in value that the person accepts to forgo as a consequence of her risk aversion is called the opportunity cost of risk.

To make a more general statement about the nature of the certainty equivalent, consider an investor whose utility is a function of her wealth and risk aversion. Let Wdescribe the investor's wealth and γ describe the investor's level of risk aversion, where γ is defined by John W. Pratt's (1964) formulation of Absolute Risk Aversion:

Equation 19

$$ARA(W) = -\frac{U''(W)}{U'(W)}$$

Consider an investor with a negative exponential utility function of wealth and risk aversion

Equation 20

$$U(W) = -e^{-\gamma w}$$

Where $\gamma > 0$, that is, the investor is risk averse.

Imagine a fair gamble in which the investor has equal probabilities of gaining x and loosing x. The investor's certainty equivalent in relation to the fair gamble is given by following function:

Equation 21

$$U(CE) = \frac{1}{2}U(W + x) + \frac{1}{2}U(W - x)$$

Or equivalently,

Equation 22

$$e^{-\gamma} * CE = \frac{1}{2}e^{-\gamma(W+x)} + \frac{1}{2}e^{-\gamma(W-x)}$$

Solving for CE in equation 22 gives us the investor's certainty equivalent for the fair gamble

Equation 23

$$CE = W + \frac{\ln(2)}{\gamma} - \frac{1}{\gamma}\ln(e^{-\gamma x} + e^{\gamma x})$$

In order to avoid the gamble, the maximum amount that the investor is willing give up, is

Equation 24

$$W - CE = \frac{1}{\gamma}\ln(e^{-\gamma x} + e^{\gamma x}) - \frac{\ln(2)}{\gamma}$$

From equation 24, it can be shown that a higher value of γ is associated with a higher difference $W - CE$. This is particularly important, as it is in line with the interpretation of γ as a measure of the investor's level of risk aversion (Munk, 2017).

The negative exponential utility function described above, is a powerful tool when looking at populations where wealth is normally distributed. However more sophisticated utility specifications are required when considering distribution of wealth that is characterized by non-normality.

3.5.3 Portfolio turnover

When developing a theoretical trading strategy, several simplifying assumptions are often introduced. As transaction costs can be difficult to predict ex-ante, people often consider a paper portfolio in which transaction costs are omitted. This allows for an easy way of invalidating poor strategies. If a paper portfolio does not perform well before transaction costs, there is no reason to spend additional effort on estimating the performance after transaction costs. If however, the paper portfolio performs well, the estimation of transaction cost becomes important for assessing the strategy. Investors are concerned with returns after transaction costs, hence the implementation cost becomes a key subject for validating the economic benefit of the strategy. The optimal trading strategy is conditioned by the transaction cost setting that the investor incurs. The cost setting is divided into three categories that describes various transaction cost environments as functions of trading size (Pedersen, 2015). The difference in the three environments is based on the convexity of the function of trading costs. Let TC describe the absolute transaction costs incurred from implementing the strategy and let φ describe the size of the trades. Increasing, constant, and decreasing transaction costs are described by following second order derivatives $TC''(\varphi) > 0$, $TC''(\varphi) = 0$, and $TC''(\varphi) < 0$ respectively. Though, the categories differ in nature, they all assume that the absolute transaction costs are increasing with the level of trade.

When implementing a strategy, the primary cost driver is transaction costs, which are incurred when trading securities. DeMiguel, Garlappi, & Uppal (2009) introduced a simple measure that constitutes a proxy for transaction costs related to implementing an investment strategy. For simplicity purposes, the measure does not take into account the different transaction cost environments but focuses on the portion of the portfolio that is rebalanced. The measure is named portfolio turnover (PT) and is calculated using following formula

Equation 25

$$PT = \frac{1}{T-1}\sum_{t=1}^{T-1}\sum_{j=1}^{N}\left(\left|w_{k,j,t+1} - w_{k,j,t^+}\right|\right)$$

PT computes the average sum of value that is traded across the N available assets over $T-1$ rebalances. For any strategy k, the term w_{k,j,t^+} describes the weights of asset j before rebalancing, while $w_{k,j,t+1}$ describes the desired weights after rebalancing.

In short, PT calculates the average percentage of wealth that is traded over all rebalances in the sample period. Even though the measure does not contain any direct economic interpretation, it is reasonable to assume that a high (low) turnover is associated with high (low) transaction costs.

3.5.4 Opportunity cost

In order to test for significance in economic differences between two trading strategies, Simaan (1993) introduced opportunity cost (or optimization premium) as an analytical framework for comparing investment decisions. Opportunity cost describes a measure of the difference in return that an investor experiences from choosing one strategy over another. Consider a benchmark portfolio X and an alternative portfolio $\hat{X}$, and let r and $\hat{r}$ denote their respective returns. In order to obtain indifference in relation to the choice of portfolio, the investor will require a premium θ such that

Equation 26

$$E[U(r+\theta)] = E[U(\hat{r})].$$

It should be noted that θ can obtain both positive and negative values. In periods, where the alternative portfolio outperforms the benchmark portfolio, θ is positive, while it is negative in periods of underperformance. From equation 26, one can define θ as the premium that an investor requires in order to be indifferent between the two strategies.

By assessing the performance from a utility perspective, both the joint distribution of asset returns and the individual investor's utility specification are taken into account. This enables the performance measure to evaluate a strategy for different utility functions. Further, the comparison with a common benchmark, allows the measure to rank several alternative strategies on a relative basis.

4 Data and methodology

This section describes the variables used in the regressions, the data sources, and how the data has been prepared for the analysis. Also, it discusses the statistical robustness checks as well as the statistical procedure for significance testing of portfolio performance.

4.1 Data description and treatment

Several types of data have been gathered to investigate the co-movement of stock and bond returns. Each of these are described below together with an explanation of how the data has been prepared for the analysis.

4.1.1 Stock and bond return data

To examine the benefits of using forward-looking analyst forecasts to explain the time-varying correlation and volatilities of stock and bond returns, the S&P500 total return index is used as the stock index and the Bloomberg Barclay's US treasury index as the bond index. Both time-series are downloaded from the Bloomberg database. The S&P500 index comprises the 500 largest companies on the New York Stock Exchange based on market capitalization and is frequently used as proxy for the overall US stock market. Bond returns are proxied by the Bloomberg Barclay′s US treasury index which measures US dollar-denominated, fixed-rate, nominal debt issued by the US Treasury. It is a total return index from which bond returns have been calculated.

The thesis is based on monthly frequencies, so daily return data is used to obtain estimates for the realized monthly volatilities and correlation. The monthly frequency has been chosen to be able to compare the empirical findings of this study to other research, since most research use monthly observations (Campbell and Ammer, 1993; Andersson, Krylova and Vähämaa, 2008). The daily returns on both the stock and bond index are turned into log returns in the following way:

Equation 27

$$R_{i,t} = \ln{(1 + r_{i,t})}$$

Where, $R_{i,t}$ is the continuously compounded return of the index on day i in month t and $r_{i,t}$ is the discrete return of the index on day i in month t. Next, the monthly realized volatility and correlation is estimated as suggested by Jivraj (2012). Volatility is estimated from daily returns in month t as:

Equation 28

$$\sigma_t = \sqrt{\sum_{i=1}^{N_t} (R_{i,t})^2}$$

And realized correlation at the end of month t as:

Equation 29

$$\rho_{S,B,t} = \frac{\sum_{i=1}^{N_t} R_{i,t}^S * R_{i,t}^B}{\sigma_{S,t} * \sigma_{B,t}}$$

Where N_t is the number of daily returns in month t and $R_{i,t}^S$ and $R_{i,t}^B$ are the daily continuously compounded stock and bond return on day i in month t. These two variables are used as the regressand in the analysis section.

4.1.2 Analyst forecast data

Analyst forecasts of macroeconomic variables are obtained from The Federal Reserve Bank of Philadelphia's Survey of Professional Forecasters (SPF). Every quarter, surveys are sent out to a diverse group of professionals, mainly from the business world and Wall Street, who all forecast using econometric models as a part of their current job. The SPF provide mean values of the forecasts, which are used as explanatory variables, as well as individual forecasts, which are obtained in order to construct dispersion measures of the forecasts. Two different forecasts are used, a forecast for the remainder of the current calendar year and a forecast for next year.

In order to obtain comparable and consistent time-series of all forecast variables, and to remove the seasonality that exist in the reported forecasts, the expectations for the current and next year are weighted together to get a measure of the 12-month ahead expectations, similarly to Andersson et al., (2008). Although the objective of the analysis is to predict next month's volatilities and correlation, several studies have shown that one-year forecast data displays significant predictability for the one-month ahead volatility and correlation (Andersson, Krylova and Vähämaa, 2008; Jivraj, 2012b).

However, the preference for using monthly data in the light of quarterly releases of expectations complicates the construction of a monthly one-year ahead looking variable. SPF release data in the end of the second month of every quarter, which is the reason for the following construction of a 12-month ahead expectation variable:

Equation 30

$$E_{12_{t-1}} = \frac{12 - Month_{t-2}}{12} E^{C}_{Month_{t-2}} + \frac{Month_{t-2}}{12} E^{N}_{Month_{t-2}}$$

Where $E_{12_{t-1}}$ denotes the 12-month ahead expectations of the macroeconomic variable at time $t-1$. Additionally, $E^{C}_{Month_{t-2}}$ and $E^{N}_{Month_{t-2}}$ denote the current year and next year expectation of the variable from the latest available survey, which is the quarter at month $t-2$. For instance, to explain the variation in stock-bond correlation from daily returns in July 2018, the one-year ahead expectation variable from June is used. This variable is constructed as $\frac{(12-5)}{12} E^{C}_{Q2} + \frac{5}{12} E^{N}_{Q2}$, which means a weight of $\frac{7}{12}$ on current year expectation and $\frac{5}{12}$ on next year expectation, where both forecasts are from the second quarter SPF of 2018. The collection of mean-consensus forecast variables will be denoted $\boldsymbol{F}_t$.

Ideally, one should use monthly survey data to construct the variable. However, in the lack thereof, the one month lagged 12-months ahead expectation variable is the second-best alternative.

As mentioned in the literature review, researchers studying stock-bond correlation have found that dispersion in analyst forecasts has explanatory power in explaining the relationship between the two returns. In line with these researchers, a dispersion measure of all SPF analyst forecasts is constructed by taking the cross-sectional mean-absolute-deviation of the implied 12-months ahead expectation variable from equation 30 in each month. It is calculated as:

Equation 31

$$Mean\ absolute\ deviation_{t-1} = \frac{1}{n_t} \sum_{i=1}^{n} \left| E_{12,i,t-2} - \overline{E_{12}} \right|$$

Where $E_{12,i,t-2}$ is forecaster i's 12-months ahead expectation of the macro variable in the quarter related to month $t-2$. The mean-absolute-deviation is chosen since it is a more robust statistic compared to the traditional standard deviation, which is more prone to be impacted by outliers. The collection of dispersion variables will be denoted $\boldsymbol{D}_t$.

All forecasting variables have been selected based on their qualities described in the existing literature rather than a trial-and-error approach. A total of five macro variables from the SPF is used in this paper, that is, the inflation rate, real GDP growth, short- and long rate as well as growth in corporate profits. The short rate variable is a forecast of the three-month treasury rate, whereas the long rate is a forecast of the ten-year treasury bond rate.

A sixth forward-looking variable is included, which is also motivated by the existing literature. It is the implied volatility extracted from option pricing based on the S&P500 index, also called the VIX index. Connolly, Stivers and Sun (2005) argue that this variable reflects both the level and the uncertainty of the expected future stock volatility. The monthly time-series is sourced from the Chicago Board Options Exchange. Even though this variable is not a part of the SPF data, it is treated as a part of the collection of mean-consensus forecasts in the analysis. The reason is, it can be argued that implied volatility is a forward looking variable, taking into account the expectations of the future macro environment, hence is not backward looking.

4.1.3 Historical data

Historical data is also needed for the analysis. It is interesting to investigate whether analyst forecast add additional explanatory power to the model, or whether volatilities and correlation can be predicted simply by the realized historical values of the above-mentioned macro-variables. This dataset comprises historical values of all macro-variables from the SPF and is also sourced from The Federal Reserve Bank of Philadelphia's database. The realized values of the three-month treasury rate and the ten-year treasury bond rate are gathered from Bloomberg.

Data on the risk-free rate is also necessary for the analysis related to portfolio performance. The one-month T-bill return from Ibbotson and Associates is employed as a proxy for the risk-free rate.

4.2 In-sample and out-of-sample analysis

Predictability tests can be performed using the in-sample fit of a model or alternatively they can be based on the out-of-sample fit obtained by performing a sequence of rolling-window regressions. Using the in-sample specification allows the use of the entire dataset to fit the model of interest. In the latter case, the real-life data constraints of a forecaster are accounted for. Here, rolling-window regressions are based only on data available to the forecaster at time t, and not the entire sample, that is, including data beyond time t, as in the in-sample approach.

In practice, the in-sample predictability tests seem to be more effective, as the null hypothesis of no predictability tends to be rejected more often than out-of-sample tests (Inoue and Kilian, 2002). However, if one is interested in mimicking the actual investment opportunities available to the investors, it does not make sense to use the entire sample. On the other hand, out-of-sample predictability tests that are performed by splitting the dataset into smaller samples, involves a loss of information, and therefore also lower power. As a result, out-of-sample tests may fail to discover actual predictability of a variable, due to the informational loss, whereas the in-sample test correctly reports the predictability. This exact fact has recently been illustrated by Kilian and Taylor (2003). Consequently, both in-sample and out-of-sample predictability tests will be employed in this thesis.

4.3 Ordinary least squares

Throughout this paper, we make extensive use of the statistical method of ordinary least square (OLS). Specifically, OLS provides the statistical method for 1) our in-sample analysis, 2) the out-of-sample analysis, and 3) robustness test via White and Durbin Watson d test.

The method of OLS is attributed the German mathematician Carl Fredrik Gauss. Under certain assumptions, the method is very powerful, for which reason it is one of the most applied methods for regression analysis (Gujarati and Porter, 2008). The purpose of the method is to estimate an unknown population regression function (PRF) that describes a relation between presumably explanatory variables and a dependent variable. In order to approximate the unknown specifications for the PRF, a sample regression function (SRF) is constructed from a sample of the population data. The SRF, then provides a "best guess" for the PRF. To illustrate this point, consider a PRF described by

Equation 32

$$Y_i = \beta_0 + \beta_1 X_{1i} + \epsilon_i$$

Where X_{1i} denotes the i'th observation of the explanatory variable, Y_i is the dependent variable, β_0 is the intersection of the PRF with the y-axis, β_1 describes the dependent variable's relation with the explanatory variable, and ϵ_i denotes a stochastic error term with an expected value of zero. That is, $E(\epsilon_i) = 0$.

As the PRF is not directly observable, we estimate the function by constructing an SRF defined by

Equation 33

$$Y_i = \hat{\beta}_0 + \hat{\beta}_1 X_i + \hat{\epsilon}_i$$

$$= \hat{Y}_i + \hat{\epsilon}_i$$

Where the "hat" describes the sample estimate of the population value. If we isolate $\hat{\epsilon}_i$ in equation 33 we get

Equation 34

$$\hat{\epsilon}_i = Y_i - \hat{\beta}_0 - \hat{\beta}_1 X_i$$
$$= Y_i - \hat{Y}_i$$

which reveals, that $\hat{\epsilon}_i$ is simply the difference between the actual and the estimated dependent variable.

With n pairs of observed values for X and Y, we want to identify the SRF that is closest to the actual Y. Following the OLS method, this is done by minimizing the squares of the residuals. Notionally expressed as

Equation 35

$$\min_{\beta_j \in \mathbb{R}} \sum_{i=1}^{n} \hat{\epsilon}_i^2$$
$$= \min_{\beta_j \in \mathbb{R}} \sum_{i=1}^{n} \left(Y_i - \hat{Y}_i\right)^2$$
$$= \min_{\beta_j \in \mathbb{R}} \sum_{i=1}^{n} \left(Y_i - \hat{\beta}_0 - \hat{\beta}_1 X_i\right)^2$$

Minimizing the square of the residuals, as compared to minimizing the arithmetic mean of the residuals, has two important implications. First, it eliminates the problem that negative and positive residuals cancel each other out. Second, it gives more weight to large residuals than small residuals. A third justification for squaring the residuals is that the obtained estimators β_j has some very desirable statistical properties, given that the Gauss-Markov assumptions are met for the sample data (Gujarati and Porter, 2008). These assumptions are elaborated upon in the following section.

4.3.1 OLS assumptions

According to the Gauss-Markov theorem, an OLS regression will provide the minimum variance estimators of a sample, given that six assumptions are met. Specifically, the assumptions are:

1) The regression model is linear and includes an error term. This means that the regression can be written as

Equation 36

$$Y_i = \beta_0 + \beta_1 X_{1i} + \beta_2 X_{2i} + \cdots + \beta_K X_{Ki} + \epsilon_i$$

2) The error term has a population mean of zero. Formally written as $E(\epsilon_i) = 0$.

3) All explanatory variables are uncorrelated with the error term. Notionally represented as

Equation 37

$$\rho_{X_{1i},\epsilon_i} = \rho_{X_{2i},\epsilon_i} = \cdots = \rho_{X_{Ki},\epsilon_i} = 0$$

4) All error terms are uncorrelated:

Equation 38

$$E(\epsilon_i \epsilon_j) = 0 \text{ for } i \neq j$$

5) All error terms have constant variance: $Var(\epsilon_i) = \sigma^2$

6) There is no perfect linear relationship between two or more explanatory variable (Studenmund, 2014).

Commonly, the acronym BLUE (Best Linear Unbiased Estimator) is used to describe an estimator that is obtained from OLS based on a sample that meets the six abovementioned assumptions. As our strategies are obtained through OLS regressions, we test our in-sample data for some of the most relevant BLUE assumptions. Specifically, we have tested our in-sample data for autocorrelation and homoscedasticity using Durbin-Watson d test and White test, respectively.

The BLUE tests, described in this section, are solely conducted on our in-sample data. The reason being that we do not use the estimators from the out-of-sample analysis to derive any conclusions. The out-of-sample analysis is exclusively conducted in order to assess whether the proposed strategies outperform benchmark strategies. Hence, violation of the BLUE assumptions in relation to our out-of-sample data is more or less irrelevant to our findings. As described in section 5.2, we have made six in-sample regressions, where different explanatory variables are included. We focus on the in-sample regression that includes the lagged variable, the mean-consensus forecasts, and the dispersion measures. As a consequence, the BLUE tests will be conducted on this particular in-sample regression specification.

4.3.2 Durbin-Watson d Test

In order to test for autocorrelation, we conduct a Durbin-Watson d Test. In the strongest form, the BLUE assumption, as stated in equation 38, describes a sample where none of the error terms are correlated. While any order of autocorrelation will influence our findings, we choose to focus on first order autocorrelation, which describes a sample where observations in period t has predictive power in relation to observations in period $t + 1$. The first order autocorrelation is notionally described as $\epsilon_t = \rho\epsilon_{t-1} + u_t$, where ρ describes the first-order autocorrelation coefficient and u_t describes a classical error term (Studenmund, 2014).

By investigating, the nature of the error terms over the period, these can be used to measure the presence of autocorrelation in the sample. The Durbin-Watson d Test incorporates a so-called Durbin-Watson d statistic, which is calculated as

Equation 39

$$d = \frac{\sum_{t=2}^{T}(\epsilon_t - \epsilon_{t-1})^2}{\sum_{t=1}^{T}(\epsilon_t)^2}$$

The d statistic is a number between 0 and 4, where 0 describes extreme positive serial correlation, 2 describes absence of autocorrelation and 4 describes extreme negative serial correlation. By comparing the obtained Durbin-Watson d statistic with some upper and lower thresholds, we are able to determine whether our in-sample data is characterized by autocorrelation. Results from the test are presented in the robustness section of the analysis.

4.3.3 White Test

As noted in section 4.3.1 "OLS assumptions", the BLUE assumption number five is related to the scedasticity of the sample. Specifically, it states that all error terms in the sample should be characterized by a constant variance.

When considering financial markets, the earliest formulations did assume homoscedasticity. Around the turn of the 20th century, Louis Bachelier related Brownian Motion to financial markets. In short, the model describes the development of financial markets as a stochastic function of time and volatility. Though the original model posed a promising model for financial markets, in its simplest form, it provides unsatisfactory predictions for the actual financial market development. Specifically, the simplest form of the model assumes a stationary level of volatility and an expected value of zero. Slight modifications of the model can be implemented to allow for different levels of variance as well as non-zero expected value (drift) (Roberts, 2011).

Although the Brownian Motion model can be constructed in a way that incorporates different levels of variance, it does assume that this level is constant over time. Shiller (1989) finds that this assumption is not in line with observed data, which shows varying levels of return in different time periods. The causes for the observed heteroscedasticity in financial market are not fully understood, and researchers have suggested that it is driven by factors such as, fads, fashions, and social movements (Shiller, 1989), general market conditions (Schwert and Seguin, 1990), and self-fulfilling prophesies, as illustrated by Keynes' famous beauty contest analogy, where speculators devote their "intelligences to anticipating what average opinion expects to be" (Keynes, 1936).

Regardless of the nature of the cause, it is generally accepted that financial data tends to suffer from heteroscedasticity. In order to identify heteroscedasticity in data samples, different tests have been introduced. One such test was proposed by Halbert White in 1980 (White, 1980). The test has become one of the most popular heteroskedasticity tests, due to the relatively easy implementation and the property that it does not rely on the normality assumption (Gujarati and Porter, 2008).

The process of the test proceeds as follows. Consider the regression model from equation 36. From sample data, we estimate equation 36 and obtain residuals $\hat{\epsilon}_i$. We then construct an auxiliary regression, where the dependent variable is the square of the obtained residuals, and the dependent variables constitutes the original X variables, the square of the X variables, and all cross products of the X variables. The auxiliary regression is then defined by

Equation 40

$$\hat{\epsilon}_i^2 = \alpha_0 + \sum_{j=1}^{K} \alpha_j X_{ji} + \sum_{j=1}^{K} \sum_{k=j}^{K} \alpha_{K+s} X_{ji} X_{ki}$$

$$i = 1\,, 2\,, \dots\,, n$$
$$j = 1\,, 2\,, \dots\,, K$$
$$k = 1\,, 2\,, \dots\,, K$$
$$s = 1\,, 2\,, \dots\,, K(K+1)/2\,.$$

The number of explanatory variables in the auxiliary regression is given by

Equation 41

$$df = K\left(1 + \frac{(K+1)}{2}\right)$$

and it can be shown that

Equation 42

$$n * R^2 \underset{asy}{\sim} \chi^2_{df}$$

meaning that the product of R^2 of the auxiliary regression and the number of observations n, assumptotically follows the Chi-squared distribution with degrees of freedom, as defined in equation 41 (White, 1980).

If the obtained value from equation 42 exceeds the critical chi-square value with the chosen level of significance, it is concluded that the data is characterized by heteroscedasticity. If it does not exceed the critical chi-square, the conclusion is, that the data is characterized by homoscedasticity. Which is to say that

Equation 43

$$\alpha_j = \alpha_{K+s} = 0$$

4.4 Calculation of evaluation metrics

To evaluate the out-of-sample performance of the proposed strategies, we obtain portfolio standard deviation, Sharpe ratio, certainty equivalent, portfolio turnover, and opportunity cost for each strategy. These measures comprise our performance matrix which is used to compare our proposed strategies' out-of-sample performance to benchmark investment strategies. It should be noted that standard deviation, Sharpe ratio, and certainty equivalent are all calculated on an annual basis and are based on excess returns, calculated as the return of the portfolio minus the risk-free rate for the same month.

For every month t, we calculate the optimal portfolio weights w_t^k, based on the individual strategy k. For every month t we calculate the out-of-sample return as $r_{t+1}^k = w_t^k \boldsymbol{r}_{t+1}$, where $\boldsymbol{r}_{t+1}$ is an excess return vector for the assets in our portfolio (stocks and bonds) over the period t to $t+1$. The out-of-sample analysis provides a series of excess returns for each strategy, which comprises the data used to calculate the performance measures for individual strategies.

4.4.1 Mean return

The mean return is not reported as a performance metric but is used in the calculations of the metrics below. It defines the average monthly excess return for strategy k and is calculated as

Equation 44

$$\mu^k = \frac{1}{T-\tau} \sum_{t=\tau}^{T-1} r_{t+1}^k$$

Where τ describes the length of the estimation window used in the rolling-window regressions. τ is set equal to 60 months. T is the total number of observations.

4.4.2 Standard deviation

The risk of the individual strategy is defined as the standard deviation, calculated by:

Equation 45

$$\sigma^k = \sqrt{\frac{1}{T-\tau-1}\sum_{t=\tau}^{T-1}\left(r_{t+1}^k - \mu^k\right)}$$

4.4.3 Sharpe ratio

As defined in equation 18, the Sharpe ratio describes the excess return over volatility. In order to comply with convention, we represent the Sharpe ratio as an annualized measure. In order to do so, we annualize mean excess return and standard deviation by following formulas: $\mu_A^k = 12 * \mu^k$ and $\sigma_A^k = \sqrt{12} * \sigma^k$, where subscript A denotes annualized measures.

As defined in equation 18, the Sharpe ratio describes the return over the volatility of excess returns. For our abovementioned measure, the Sharpe ratio is defined as

Equation 46

$$SR = \sqrt{12} * \frac{r_{t+1}^k - rf_{t+1}}{\sigma\left(r_{t+1}^k - rf_{t+1}\right)} = \frac{\mu_A^k}{\sigma_A^k}$$

4.4.4 Certainty equivalent

To calculate the certainty equivalent, we use the definition by DeMiguel et al. (2009). The advantage of their approach is, that the number is calculated as a return measure rather than an absolute number presented by Hershey & Schoemaker (2008). Following the definition by DeMiguel et al. (2009), the certainty equivalent is calculated as

Equation 47

$$CEQ^k = \mu_A^k - \frac{\gamma}{2}\left(\sigma_A^k\right)^2$$

where γ describes the risk aversion. For our primary analysis, we assume that $\gamma = 1$. However, in table 12 in section 5.4 "Robustness check", we present results of certainty equivalent with varying levels of risk aversion.

4.4.5 Portfolio turnover

Following the definition of equation 25, the portfolio turnover for each month is calculated as

Equation 48

$$PT_M^k = \sum_{j=1}^{N} (|w_{j,t+1}^k - w_{j,t^+}^k|)$$

where PT_M^k denotes the portfolio turnover for strategy k for month $t+1$. N denotes the total number of assets in the portfolio. For all strategies k, we exclusively consider stocks and bonds as available assets. Hence $N = 2$ for all strategies throughout this paper.

As we are interested in obtaining a single measure of portfolio turnover for the whole investment period, we average the monthly portfolio turnover and obtain

Equation 49

$$PT^k = \frac{1}{T-\tau-1} \sum_{t=\tau}^{T-1} \sum_{j=1}^{N} (|w_{j,t+1}^k - w_{j,t^+}^k|)$$

which describes the average portfolio turnover for each strategy k over a total of 233 rebalances.

4.4.6 Opportunity cost

As a general concept, opportunity cost is defined as the number θ that solves equation 26. For our specific case, opportunity cost for each month is obtained by isolating OC_{t+1}^k in following equation:

Equation 50

$$U(r_{t+1}^{BM} + OC_{t+1}^k) = U\left(r_{t+1}^k\right)$$

where r_{t+1}^{BM} denotes the return on the Benchmark strategy over the period t to $t+1$. In order to get a measure over the entire period, and not for each individual month, we find the average opportunity cost over the entire period $T - \tau$ as

Equation 51

$$\frac{1}{T-\tau}\sum_{t=\tau}^{T-1} U(OC_{t+1}^{k})$$

The abovementioned measures will comprise our evaluation metrics used to assess the performance of each strategy.

4.5 Statistical significance testing of portfolio performance

In the out-of-sample analysis, portfolio performance is tested for statistical significance. More specifically, the difference in portfolio standard deviation, Sharpe ratio, and certainty equivalent between the two benchmark portfolios and all other strategies are tested. A non-parametric bootstrap method is employed for the analysis. It is the stationary bootstrap method put forward by Politis and Romano (1994), which do not make any assumption about the distribution of portfolio returns.

Furthermore, since the procedure resamples in blocks of random length, the new bootstrapped samples preserve the dependency structure of the data. These features make it an appropriate tool for testing time series data, which applies to this thesis.

In accordance with other similar studies, the number of bootstrap resamples B is set to 10,000 with an expected block size equal to 20 (Jivraj, 2012b). This gives 10,000 samples containing $T-\tau=234$ monthly portfolio returns, where τ is the length of the estimation window, which is set equal to five years of data and T is the total number of months in the sample period which is 294.

Next, the portfolio standard deviation, Sharpe ratio, and certainty equivalent is calculated for all 10,000 resamples. The difference in performance measures between the benchmark strategy and strategy k make up the empirical distributions of the difference in portfolio standard deviation, Sharpe ratio, and certainty equivalent, respectively. From these sample distributions, p-values for the difference in performance can be obtained.

The focus of this paper is to investigate whether the use of analyst forecasts in determining portfolio weights perform better than simple benchmark strategies. Hence, the hypothesis is, that the difference in performance metric of strategy k and the benchmark is less than or equal to zero.

Or similarly $H_0: \widehat{PM}^{k} - \widehat{PM}^{bm} \leq 0$, where PM^{k} and PM^{bm} is the respective performance measure of strategy k and the benchmark strategy. However, it will be opposite for portfolio standard deviation, as larger volatility is worse.

It is the one-sided p-value that is presented in the findings section. It gives the probability of obtaining a test statistic at least as small as the one observed, assuming that the null hypothesis is true. This means that the null hypothesis of inferior performance of strategy k compared to the benchmark is rejected given a small p-value. A five percent significance level is employed in all statistical tests.

5 Analysis and empirical findings

In this section the analysis and empirical findings are presented. Firstly, stylized facts and analysis of the stock and bond returns are put forward, alongside an analysis of the volatility and pair-wise correlation of these returns. Secondly, an in-sample regression analysis provides insight into the predictive power of the different variables employed. This is followed by an analysis of the out-of-sample benefits of analyst forecast data to predict the co-movement of stock and bond returns. The last section concerns robustness checks of the abovementioned analysis.

5.1 Stylized facts and the co-movement of stock and bond returns

The two asset classes under investigation, stocks and bonds, have very different risk and return profiles. Not only are they different, but the two asset classes have time-varying characteristics such as mean return and volatility. Hence, choosing the optimal allocation to either asset, accounting for these differences and their time-varying co-movement may seem like a daunting task. If one ignored the risk aspect of the asset allocation problem and looked at past realized returns only, then the task seems simple.

Figure 8 displays the development of one hundred dollar invested at 1 July 1994 and kept until 31 December 2018 in the S&P500 index and the Bloomberg Barclay's US treasury bond index, respectively.

Figure 8 - Total return of the S&P500 index and the treasury bond index for the entire sample period July 1994 to December 2018. Indexed to 100 at July 1994

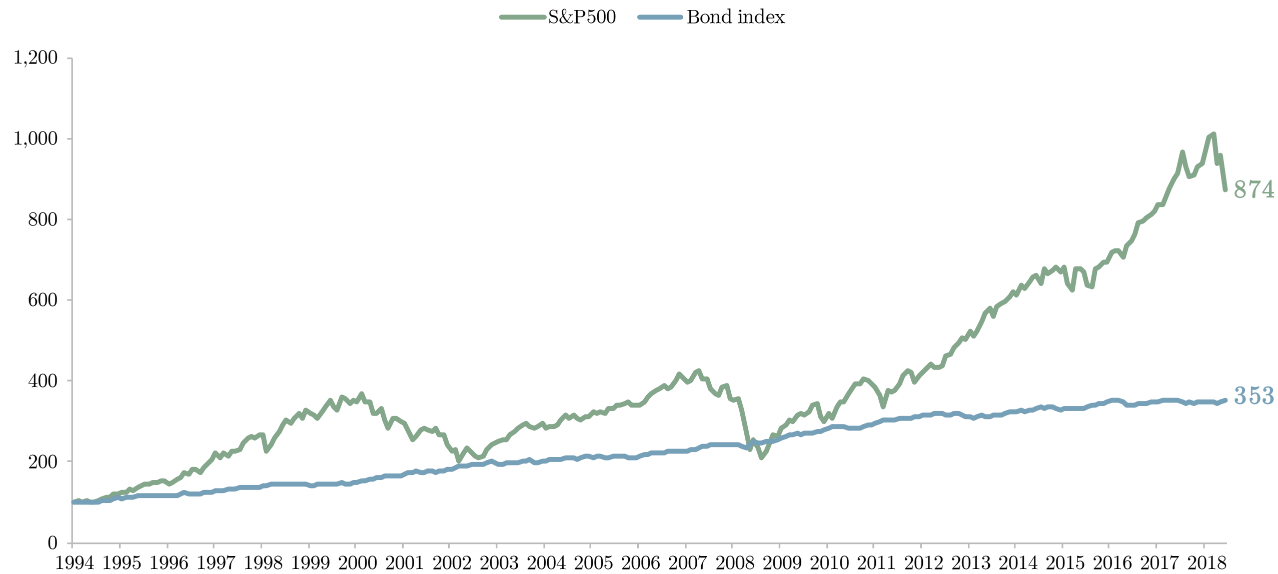

Source: Own contribution

It seems obvious that if expected return is the only decision criteria underlying the asset allocation process, then nobody would be investing in bonds. Over the sample period of approximately 24 years, stocks outperformed bonds by 521 percentage points. Even though stocks grew more than double that of bonds over the same period, this comparison is unfair as it does not account for the inherent risk in the respective asset class as well as the co-movement of returns.

In order to make an informed investment decision, the entire risk-return profile of each asset must be considered. Table 1 conveys such information as it presents descriptive statistics of the stock and bond returns, including the standard deviation of returns.

Table 1 - Descriptive statistics for annualized stock and bond log returns for several sub-periods using daily return data

Stock and bond log returns for sample period July 1994 to December 2018		
	Stocks	Bonds
For entire sample period		
Mean return (%)	9.34	4.99
Standard deviation (%)	18.31	4.44
Correlation	-0.25	
For the period 1995-2002		
Mean return (%)	10.20	8.51
Standard deviation (%)	18.95	4.62
Correlation	-0.07	
For the period 2003-2010		
Mean return (%)	6.62	4.44
Standard deviation (%)	21.24	5.01
Correlation	-0.33	
For the period 2011-2018		
Mean return (%)	11.24	2.32
Standard deviation (%)	14.56	3.59
Correlation	-0.42	

Source: Own contribution

The mean annualized log returns and standard deviation calculated based on daily data, reveal that stocks and bonds indeed have very different risk profiles, even time-varying so. The mean annualized log return of stocks where 9.34% whereas bonds delivered a mean return of 4.99%. This is not surprising considering the vast difference in total return from figure 8, however the table also shows that the annualized standard deviation is very different. Over the entire sample period, stock returns varied on average 18.31% per year, whereas bond returns varied only by 4.44%. The mean annualized risk-free rate over the same period, not reported in the above table, was 2.37%. Hence, over the entire sample period, the ex-post

Sharpe-ratio of stocks was 0.44 while it was 0.84 for bonds. This indicates that when returns are adjusted for risk, it seems like bonds deliver a higher return per unit of risk.

The entire sample has been divided into three sub-samples, each consisting of approximately eight years of data. This provides insight into how the mean return and risk measures change over time. The standard deviation of both asset classes seems to have evolved in the same direction. The variation of returns increased from the first sub-period to second, while the variation of log returns decreased to an even lower level than the initial level in the most recent sub-period 2011-2018. However, the table shows that the development in mean return decoupled over the three sub-periods. Both assets experienced declining mean return from the first to the second period, where after average stock returns increased to the highest level of all sub-periods, while bond returns kept decreasing. From 2011-2018 bonds delivered a mean return of only 2.32% while an investment in stocks yielded 11.24%.

The decoupling of returns over the sub-periods has also resulted in an increasingly negative correlation between the two asset classes. The correlation of-0.07 in 1995-2002 increased to -0.42 in the most recent sub-period, while the correlation over the entire sample period has been -0.25. This means that a portfolio of stocks will receive increasing diversification benefits by adding bonds to the portfolio, as the returns of the two asset classes tend to move in opposite direction.

5.1.1 Realized versus forecasted volatilities and correlation

The realized correlation between the S&P500 stock index and the Bloomberg Barclay's US treasury bond index, calculated as in equation 29, is graphically displayed in figure 9 below.

Figure 9 – One-month ahead correlation: forecast vs. realized correlation

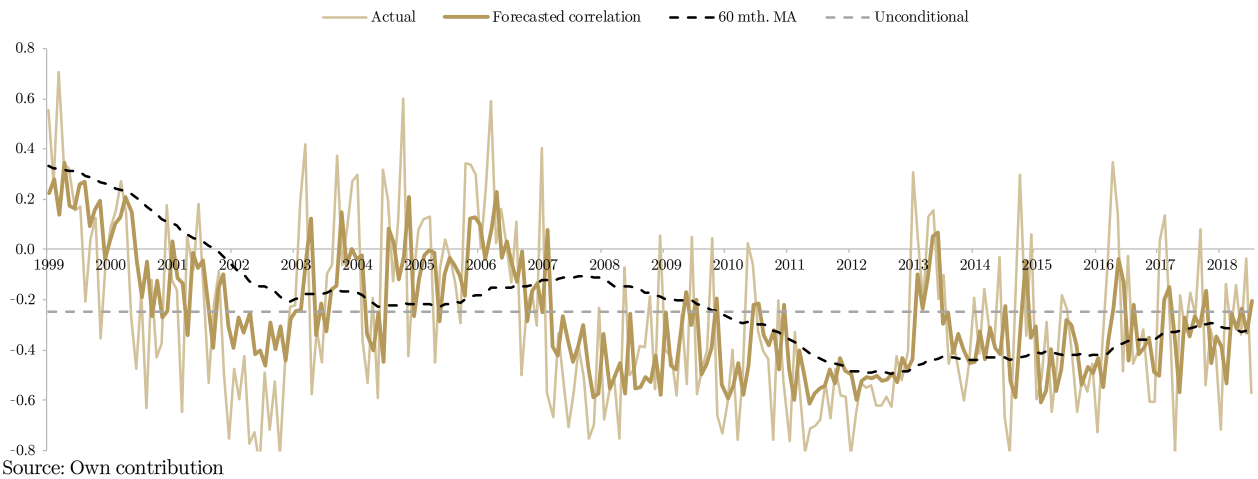

Source: Own contribution

The monthly realized correlation is plotted with the one-month ahead forecasted correlation using the panel C lag-specification from the in-sample analysis in the next part of the analysis section. Additionally, a 60-months' moving-average measure together with the unconditional correlation for the entire sample period is displayed for comparison.

It is very clear from the graph that stock-bond correlation is indeed time-varying. The graph reveals both periods of sustained positive correlation and most recently sustained negative correlation. Interestingly, it seems like the correlation drops and becomes negative around periods of financial distress, such as in the period around year 2000, after the so-called dot-com bubble, and again after the most recent financial crisis starting in 2007. Also, this is in line with the flight-to-quality phenomenon often described in the literature, where in periods of financial distress, investors seek to shift holdings to less risky assets, that is, bonds. Hence, this mechanism bids up bond prices, while at the same time investors offload their stock holdings, ultimately decreasing stock returns. Even today, more than ten years after the financial collapse of 2007, the stock-bond correlation tends to remain negative with only few observations of positive monthly correlation.

The 60-months' moving-average correlation displays the overall trend well. Due to the nature of its construction, it is considerably less volatile compared to the realized correlation and regime shifts in correlation is not detected right away but rather observed with a lag. These attributes show the problems with using historical moving-averages as estimates for next period's correlation.

More interestingly, is the forecasted one-month ahead correlation using the in-sample model from the next section. The bold line follows the realized correlation really well, indicating that the model using analyst forecast data and the dispersion thereof, gives a closer estimate of the realized correlation compared to a model using correlation predictions from historical data. The forecast model does not explain all the variation in the actual correlation, however it does seem to capture the major movements in the co-movement of returns without much delay. This in turn, means that investors will have more precise estimates, when they make their asset allocation decision, hence potentially increasing their portfolio returns.

When forming a portfolio, not only correlation is important, but also how the returns of the two asset classes vary individually. Figure 10 shows the annualized monthly realized standard deviation of returns, as calculated in equation 28.

The green line shows the realized volatility of stock returns whereas the blue line shows the bond volatility. Not only is stock volatility at a constantly higher level than that of bonds, but there is also considerably more dispersion in monthly stock returns, compared to bond returns. This means, that the volatility of stocks explains a larger portion of the time-varying stock-bond correlation than does the bond volatility.

Figure 10 - One-month ahead annualized stock and bond standard deviation of returns: forecast vs. realized correlation

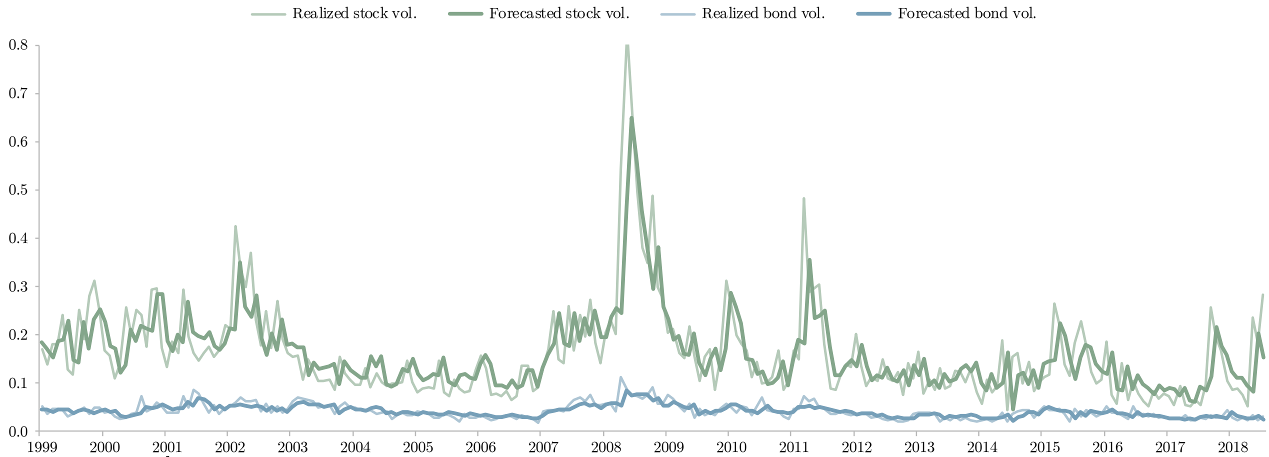

Source: Own contribution

In the graph, the thin lines represent the realized annualized standard deviation of stock and bond returns, and the bold lines show the in-sample forecast model's one-month ahead forecast of stock and bond volatility. The model using analyst forecasts of macroeconomic variables to predict volatility of the two asset classes seems to perform well. For both bond and stock returns, the model seems to capture the major movements, although not to perfection and with some delay.

Table 2 - Descriptive statistics for realized volatilities and correlation

Monthly realized volatilities (annualized) and correlation	σ^S	σ^B	$\rho_{S,B}$
Mean (%)	15.77	4.22	-0.16
Median (%)	13.79	3.97	-0.23
Minimum (%)	4.40	1.89	-0.86
Maximum (%)	82.87	11.25	0.84
Std. dev./Mean	0.60	0.35	-2.74
Correlation with σ^S	1.00	0.61	-0.30
Correlation with σ^B	0.61	1.00	-0.11
Autocorrelation coefficient of order #			
Lag 1	0.73	0.64	0.69
Lag 2	0.58	0.54	0.63
Lag 3	0.47	0.49	0.60
Lag 12	0.17	0.25	0.38

Source: Own contribution

The volatility of returns around time periods with financial distress spikes, especially for stocks. In the 2007 financial crisis, the level of stock volatility reached historically high levels, whereas bond volatility only increased slightly. Hence, stocks are not only riskier than bonds, due to the higher level of volatility but also because stock prices seem to be more susceptible to financial turmoil.

Table 2 summaries the historical annualized monthly volatility and correlation. The mean monthly realized stock volatility has been 15.77%, while bond returns only varied 4.22% on average. Volatility of stock returns varies from 4.40% to 82.87% per month, which confirms the large dispersion in stock returns. On the other hand, the maximum monthly bond volatility is only 11.25%.

The table also shows that the standard deviation of realized monthly correlation is 2.74 times its mean, where it is less than one for stock and bond volatilities. Interestingly, this indicates that the realized correlation fluctuates considerably more than realized volatility of stocks and bonds. Next, the cross-correlations between realized volatilities of each asset class show that realized volatility across stock and bond markets are highly correlated, with a coefficient of 0.61. This phenomenon has been referred to as volatility spill-over between the stock and bond markets (Fleming, Kirby and Ostdiek, 1998). However, the correlation coefficient of the realized volatilities with the correlation itself is negative for both assets. It is most so for stocks with -0.30, indicating that stock-bond correlation decreases when the respective volatilities increase. Again, this can be attributed to the flight-to-quality phenomenon, where investors reallocate more wealth to bonds in order to decrease their portfolio exposure to the higher uncertainty in stock returns. Lastly, from the autocorrelations in the table it can be noted that both volatilities and correlation follow persistent processes which decays slowly.

5.1.2 Analyst forecasts

The forecast data provided by SPF consists of many different analysts' take on how they think a certain macroeconomic variable will evolve in the nearest future. Table 3 provides summary statistics of all five variables.

Panel A shows that the variation in mean consensus of change in corporate profits and real GDP growth is considerably larger than for the remaining three variables. The standard deviation is close to three times the mean for CP and RGDP, whereas it is less than one for the other. Another interesting observation is that consensus mean of CPI, TBILL and TBOND are all highly correlation with coefficients above 0.70. Panel B shows the same statistics but for the dispersion variables instead. Analysts seem to disagree most about the future value of CP and RGDP. Again, the standard deviation, in terms of its mean, is much higher than for the other variables. These observations show that there is great uncertainty among analysts about the level of these variables but also great dispersion over time. Correlation between dispersion measure seems more arbitrary, which also is the case for the correlation between mean consensus and dispersion measures of analyst forecasts in panel C.

Table 3 - Descriptive statistics for analyst forecasts variables from the SPF

Panel A: Mean consensus analyst forecast	F^{CPI}	F^{TBILL}	F^{TBOND}	F^{CP}	F^{RGDP}
Mean	2.30	2.68	4.43	2.12	1.31
Median	2.29	2.24	4.63	1.17	0.08
Minimum	0.40	0.09	2.01	-19.16	-2.65
Maximum	3.44	6.23	7.91	48.20	24.54
Standard deviation	0.53	2.07	1.49	7.77	3.77
Correlations					
F^{CPI}	1.00	0.73	0.70	0.00	-0.07
F^{TBILL}		1.00	0.89	0.04	0.06
F^{TBOND}			1.00	0.08	0.08
F^{CP}				1.00	0.51
F^{RGDP}					1.00
Panel B: Dispersion of analyst forecast	D^{CPI}	D^{TBILL}	D^{TBOND}	D^{CP}	D^{RGDP}
Mean	0.35	0.22	0.24	67.85	37.68
Median	0.32	0.22	0.23	72.67	34.97
Minimum	0.20	0.04	0.13	6.06	14.87
Maximum	0.79	0.43	0.38	144.33	77.89
Standard deviation	0.12	0.08	0.05	36.85	13.56
Correlations					
D^{CPI}	1.00	0.24	0.32	0.40	0.60
D^{TBILL}		1.00	0.54	-0.23	-0.07
D^{TBOND}			1.00	-0.14	0.11
D^{CP}				1.00	0.58
D^{RGDP}					1.00
Panel C: Correlation between measures					
	Consensus mean				
	F^{CPI}	F^{TBILL}	F^{TBOND}	F^{CP}	F^{RGDP}
Dispersion D^{CPI}	-0.36	-0.39	-0.27	-0.08	-0.12
Dispersion D^{TBILL}	0.33	0.36	0.54	0.03	-0.03
Dispersion D^{TBOND}	0.02	0.06	0.24	-0.03	0.00
Dispersion D^{CP}	-0.56	-0.75	-0.75	0.07	-0.02
Dispersion D^{RGDP}	-0.65	-0.71	-0.62	-0.12	0.00

Source: Own contribution

5.2 In-sample predictability of macroeconomic variables

The purpose of this thesis is to investigate the determinants of stock-bond correlation and develop a predictive model that can be used to form efficient portfolios that outperform traditional asset allocation strategies. Firstly, the predictive model must be built before the second part of the analysis can be carried out. The predictive model yields forecasts of the inputs to the variance-covariance matrices needed to calculate mean-variance efficient portfolios. The inputs to the variance-covariance matrix are predictions of next period's standard deviation of stock and bond returns, respectively, as well as the correlation between the two asset classes. In order to build such a model, the determinants of stock-bond correlation must be examined.

The literature has suggested several factors explaining the co-movement of stock and bond returns as mentioned in the literature review. The most convincing factors have been selected to constitute as explanatory variables in the predictive regression model. These variables include inflation (CPI), the short rate (TBIL), the long rate (TBOND), growth in corporate profits (CP), and growth in real gross domestic product (RGDP) as well as implied volatility of the stock market (VIX). An in-sample analysis of the predictive power of these variables may contribute to the understanding of which factors that cause stock and bond returns to co-vary over time.

Additionally, the in-sample analysis will consist of different specifications of the predictive model, in which different combinations of explanatory variables will be investigated. Firstly, an investigation of the predictive power of historical information on the abovementioned macroeconomic variables is performed. Next, analysts' mean consensus forecast of the same variables as well as dispersion in these forecasts are included to determine if analyst forecast actually add any explanatory power. A comparison of the different specifications will provide insight into the question of whether analyst forecast do have predictive power of the co-movement in stock and bond returns.

The discovery of in-sample predictability of some of the macroeconomic variables do not imply that the same factors have predictive power out-of-sample, but it may serve as a guideline to what variables that can be used to forecast stock-bond correlation out-of-sample. As mentioned above, the predictive regressions are needed to provide inputs for the variance-covariance matrix, and in turn to determine the optimal portfolio weights. Hence, predictive regressions for standard deviation of returns as well as for the correlation are performed. Each one of these regressions and their different specifications are presented and analyzed below.

5.2.1 Predictive regressions of stock-bond correlation

In order to predict next month's realized stock-bond correlation, different specifications of the following predictive regression are run.

Equation 52

$$\rho_{t+1} = \hat{\alpha} + \hat{\beta}\rho_t + \hat{\varphi}\boldsymbol{H}_t + \hat{\delta}\boldsymbol{F}_t + \hat{\gamma}\boldsymbol{D}_t + \epsilon_{t+1}$$

Where ρ_{t+1} is the realized correlation in month t+1, determined as in equation 29. ρ_t is the lagged realized correlation, $\boldsymbol{H}_t$ is a collection of historical values of the explanatory variables available at time t, $\boldsymbol{F}_t$ is a collection of analyst mean consensus forecasts of the macroeconomic variables and lastly $\boldsymbol{D}_t$ is a collection of the dispersion measures of the analyst forecasts.

As a base line model, the first specification of the predictive regression includes only past realized historical values of the macroeconomic variables. It is interesting to see how well last month's realized values of inflation, short rate, long rate, change in corporate profits, and real GDP growth predict next month's realized correlation. To test such a specification the following regression will be run.

Equation 53

$$\rho_{t+1} = \hat{\alpha} + \hat{\beta}\rho_t + \hat{\varphi}\boldsymbol{H}_t + \epsilon_{t+1}$$

The results of this regression are presented in panel A of table 4, which includes coefficient estimates, t-statistics and adjusted R^2. In the first specification of panel A, two of the five explanatory variables seem to significantly explain the stock-bond correlation. The realized level of the long rate from last month H^{TBOND}, is highly statistically significant with a t-statistic above five, making it significant at the one percent level. This may not come as a surprise, as intuitively, it makes sense that the long rate has a great impact on bond pricing, hence also on the co-movement of the returns of the two asset classes. Further, the level of long rate tends to follow a quite persistent process without extreme changes to the variable. This observation is in line with the statistically significant findings and hence suggests that it may be possible to forecast next month's correlation using the historical value of last period's long rate.

The second variable that seems to explain stock-bond correlation is the growth rate in the real gross domestic product H^{RGDP}. The coefficient estimate of 0.1039 is also statistically significant at the one percent level with a t-statistic of 2.64.

Stock-bond correlation

Table 4 reports the 1-month horizon in-sample predictive regressions of stock-bond return correlation onto a constant and historical values ($\boldsymbol{H}_t$) as well as standardized measures of the mean consensus ($\boldsymbol{F}_t$) and dispersion ($\boldsymbol{D}_t$) in analyst forecasts of the inflation rate (CPI), the short rate (TBILL), long rate (TBOND), change in corporate profits (CP), change in real GDP (RGDP), and implied volatility (VIX). The table reports coefficient estimates with t-statistics in brackets and the adjusted R^2 of the regression. The critical value of the t-statistic at which the null hypothesis of $\beta = 0$ is rejected at a significance level of 5% is $|t| > 1.96$
**** denotes significance at the 1% level ** denotes significance at the 5% level * denotes significance at the 10% level*

Cons. (*t-stat*)	Lag (*t-stat*)	H^{CPI} (*t-stat*)	H^{TBILL} (*t-stat*)	H^{TBOND} (*t-stat*)	H^{CP} (*t-stat*)	H^{RGDP} (*t-stat*)	F^{VIX} (*t-stat*)	F^{CPI} (*t-stat*)	F^{TBILL} (*t-stat*)	F^{TBOND} (*t-stat*)	F^{CP} (*t-stat*)	F^{RGDP} (*t-stat*)	D^{CPI} (*t-stat*)	D^{TBILL} (*t-stat*)	D^{TBOND} (*t-stat*)	D^{CP} (*t-stat*)	D^{RGDP} (*t-stat*)	Adj. R^2 (%)
Panel A: Historical data only																		
-0.8044***		-0.0814	0.0243	0.1307***	0.0030	0.1039***												40.64
(-11.18)		(-1.07)	(1.25)	(5.03)	(1.01)	(2.64)												
-0.3848***	0.4841***	-0.0512	0.0179	0.0563**	0.0019	0.0555												53.60
(-4.88)	(9.03)	(-0.76)	(1.04)	(2.31)	(0.70)	(1.58)												
Panel B: Historical data as well as analyst mean forecasts and dispersion measures																		
-0.2145		-0.2216***	-0.0972	0.1793***	0.0068**	-0.0466	-0.0199***	-0.0390	0.1053	-0.0215	0.0009	0.0121**	0.0163	-1.0300**	1.1584**	-0.0019*	-0.0010	50.99
(-0.92)		(-3.02)	(-1.27)	(3.45)	(2.23)	(-1.04)	(-5.60)	(-0.66)	(1.25)	(-0.34)	(0.32)	(2.05)	(0.06)	(-2.51)	(2.38)	(-1.92)	(-0.43)	
-0.0580	0.3585***	-0.1426**	-0.0599	0.0669	0.0040	-0.0254	-0.0143***	-0.0373	0.0681	0.0246	-0.0006	0.0115**	0.0517	-0.5451	0.6189	-0.0018**	0.0000	56.47
(-0.26)	(5.99)	(-2.03)	(-0.83)	(1.28)	(1.38)	(-0.60)	(-4.11)	(-0.67)	(0.85)	(0.41)	(-0.23)	(2.08)	(0.22)	(-1.38)	(1.32)	(-2.00)	(0.01)	
Panel C: Analyst mean forecasts and dispersion measures only																		
-0.1854							-0.0197***	-0.0684	-0.0021	0.1561***	-0.0006	0.0158***	0.2129	-0.7732**	0.7490	-0.0029***	-0.0006	48.18
(-0.82)							(-6.08)	(-1.21)	(-0.09)	(4.35)	(-0.23)	(2.70)	(0.88)	(-2.19)	(1.59)	(-3.04)	(-0.23)	
-0.0574	0.4033***						-0.0124***	-0.0462	-0.0003	0.0843**	-0.0015	0.0122**	0.1372	-0.3406	0.3712	-0.0021**	0.0002	56.33
(-0.27)	(7.32)						(-3.95)	(-0.89)	(-0.01)	(2.45)	(-0.56)	(2.26)	(0.62)	(-1.03)	(0.85)	(-2.41)	(0.09)	

Source: Own contribution

This means that the real growth rate in last month's GDP has some predictive power of next month's realized correlation between stocks and bonds. It is commonly known that the economy acts in a cyclical manner, that is, following a persistent trend for some time before experiencing a large correction and then go back to following a persistent trend. With such a pattern it does make sense that last month's growth rate in real GDP is a good estimate of next period's rate, hence the impact of real GDP growth on stock-bond correlation from last period, is likely to be similar in the next period. This in turn means that the realized value of real GDP growth is a good predictor for next month's co-movement of stock and bond returns. While the remaining three regressors remain statistically insignificant, this specification of equation 52 using historical values only, explains 40.64% of the variation in stock-bond correlation.

However, in the second specification from panel A in table 4, when controlling for the lagged value of the realized correlation, adjusted R^2 increases notably to 53.60%. In this setting, last month's realized value of correlation explains most of the variation in next month's correlation, thereby making it the best predictor. The variable is highly significant with a test statistic of 9.03. Interestingly, when controlling for the lagged value, the variable H^{RGDP} is no longer significant in explaining the stock-bond correlation. The explanatory power of this variable from the no-lag specification seems to be caught by the lagged variable. The long rate H^{TBOND} remains statistically significant, however only at the five percent level, with a t-statistic of 2.31. Again, inflation, the short rate, and change in corporate profits remain insignificant.

Both specifications in panel A have positive coefficient estimates for all variables except inflation. This means that a higher short or long rate, or positive change in real GDP or corporate profits increases stock-bond correlation, whereas higher inflation decreases the correlation.

The predictive regression using only historical data on the macro variables performs well. However, it is interesting to see if the performance of the model can increase even further if analyst forecasts are used. Intuitively, analyst forecasts should contain information that is not embedded in past realized data. Embedded within the analysts' forecasts is implied information about their expectation of the future economic regime, macroenvironment, and investment opportunities related to the forecasted variable. Therefore, in order to see if analyst forecasts add any explanatory power over and above historical data, the entire collection of consensus-mean forecasts and dispersion measures are added to equation 53. Hence, the regression being run is equal to equation 52. The performance of this regression is displayed in panel B of table 4.

When including the collection of mean-consensus forecasts and the collection of dispersion measures, the predictive regression has a total of 17 explanatory variables, including the lagged variable.

Starting with the first specification of panel B, the regression excluding the lagged term, three of the five historical variables are statistically significant in explaining the correlation. The first significant variable is inflation H^{CPI} with a t-statistic of -3.02 and a negative coefficient estimate. Interestingly, this variable becomes significant in explaining stock-bond correlation but only so, when the mean-consensus forecasts and dispersions measures are added. The variable was not significant in any of the specifications in panel A.

On the other hand, as in the historical values only regressions, the long rate H^{TBOND} is significant at the one percent level. Again, the coefficient estimate is positive, meaning that a higher level of the long rate is associated with closer co-movement of stock and bond returns. The last of the significant historical variables is change in corporate profits H^{CP}. It is significant at the five percent level with a test statistic of 2.23. Similar to the inflation variable, the addition of analyst forecasts and the dispersion thereof somehow reveals the explanatory power in the realized value of change in corporate profits. The coefficient is positive, suggesting that increasing corporate profits will tend to tie stock and bond returns together and make them move in the same direction. Lastly, unlike in the panel A regressions H^{RGDP} is no longer significant, which indicates that its explanatory power is captured by some of the other variables. Again, the short rate H^{TBILL} is not significant.

The newly included collection of analyst forecasts consists of five standardized measures of mean consensus forecasts variables plus an implied stock market volatility variable. Of the six variables only two seem to explain correlation at a significant level. The first variable is the implied uncertainty of the stock market F^{VIX}. With a t-statistic of -5.60 it is highly significant, even at the one percent level, and is also the variable explaining the most variation in correlation of all the variables in the no-lag specification. The coefficient estimate of the F^{VIX} variable is -0.0199 and a negative value indicates, that increased stock uncertainty is related to a decoupling of stock and bond returns. This observation is in line with the "flight-to-quality" phenomenon, where increasing stock market uncertainty forces investors to liquidate their stock holdings, thereby decreasing stock returns, and instead allocating the capital to a safer asset class, that is bonds, hence increasing bond returns. The VIX variable seem to capture this effect.

The other significant mean-consensus variable is the forecast of growth in real GDP. The coefficient is positive at a level of 0.0211 and significant at the five percent level. It could be the case that the forecasted variable of real GDP growth F^{RGDP} captures all the explanatory power of the same historical variable H^{RGDP}, as the forecast variable contains the same information as the historical variable but also contains additional information implied in the analysts' forecasts. This claim may be supported by the apparent sign change of the coefficient estimate of H^{RGDP} from panel A to B. In panel A, H^{RGDP} had a positive statistically significant coefficient estimate, but in panel B, it is suddenly negative and insignificant. This could be explained by the forecast variable F^{RGDP} making H^{RGDP} redundant. The estimated coefficients on the remaining four forecast variables are all statistically indistinguishable from zero.

The last five regressors comprise the collection of dispersion measures of the five analyst forecast variables. Dispersion in the analysts' forecast of the short rate D^{TBILL} and long rate D^{TBOND} both seem to be statistically significant at the five percent level, with test statistics of -2.51 and 2.38, respectively. The coefficient estimate of short rate D^{TBILL} is -1.0300 suggesting that increasing dispersion in the forecast of the short rate tends to decrease the correlation between stock and bond returns. On the contrary, the estimated coefficient on the long rate is 1.1584, meaning that when analysts start disagreeing more on the future level of the long rate, correlation between stock and bonds increase. The dispersion in analyst forecasts of the future inflation rate and future growth rate in real GDP do not seem to have explanatory power, whereas dispersion in forecasts of change in corporate profits is significant at the ten percent level, with a t-statistic of -1.92.

Using both historical, mean consensus forecasts as well as dispersion measures of forecasts increases the total amount of explained variation in correlation to 50.99%. This is a considerable increase in adjusted R^2 of almost ten percentage points compared to the first specification of the model in panel A.

As in panel A, the same regression as above is run but this time including the lagged correlation variable. Again, the lagged variable is highly significant at the five percent level, with a positive coefficient estimate and a t-statistic of 5.99. The lagged variable seems to capture much of the explanatory power from the other variables, when comparing with the first specification of the predictive regression in panel B. It seems like including the lagged term makes both the historical variables H^{TBOND} and H^{CP}, and dispersion variables D^{TBILL} and D^{TBOND} redundant in predicting stock-bond correlation. It is also the variable that explains most of the variation in realized correlation of all the variables in the regression.

However, the historical variable of inflation H^{CPI} still seems to have significantly explanatory power at the five percent level. The regressor has a negative estimated coefficient of -0.1426, meaning that increasing realized inflation has a decoupling effect on stock and bond returns. Again, the level of implied stock volatility F^{VIX} is highly significant with a t-statistic of -4.11, hence being significant at the one percent level. Additionally, the mean consensus forecast of the growth in real GDP as well as forecast of the long rate F^{TBOND} remain significant when including the lagged term, even at a five percent level. Lastly it can be noted that dispersion of analyst forecasts of change in corporate profits is the only dispersion measure that has explanatory power. It even becomes more significant compared to the other specification in panel B, with a test statistic of -2.00, that is, significant at the five percent level. Hence, increasing disagreement between analysts on the future change in corporate profits, seems to result in decoupling stock and bond returns.

Including the lagged correlation variable in panel B increases adjusted R^2 with approximately five and half percentage point to a R^2 of 56.47%, compared to the first specification. Again, it seems like the past realized correlation is the best predictor for next period's correlation.

Now that predictive regressions using historical data and analyst forecasts have showed that both mean consensus and dispersion measures of analyst forecasts seem to have predictive power of next period's realized stock-bond correlation, it will be interesting to see if the collection of analyst forecasts and dispersion measures alone have explanatory power. Hence, to test such a specification the following regression will be run.

Equation 54

$$\rho_{t+1} = \hat{\alpha} + \hat{\beta}\rho_t + \hat{\delta}\boldsymbol{F}_t + \hat{\gamma}\boldsymbol{D}_t + \epsilon_{t+1}$$

The regression results are outlined in panel C of table 4 above. Correspondingly, as in panel A and B of table 4, panel C consists of two specifications of equation 54. One with and one without the lagged correlation variable.

The first specification without the lagged term shows good performance with three of the forecast variables being statistically significant. The first one being the implied stock market volatility variable F^{VIX}, which is highly significant with a test statistic of -6.08. As argued in the analysis above, the negative coefficient estimate indicates the presence of the "flight-to-quality" phenomenon. Interestingly, by including only analyst forecast data, the mean consensus forecast of the long rate F^{TBOND} becomes highly significant at the one percent level. The variable was not even close to being significant in the panel B regressions. Furthermore, the positive coefficient of 0.5161 suggests that increasing expectations about the future long rate tends to be coupled with increasing realized correlation between stock and bond returns in the following month. This is in line with the findings of the historical variable H^{TBOND}'s impact on next month's realized correlation.

The last mean consensus forecast variable that is statistically significant is the forecast of future real GDP growth F^{RGDP}. It is significant at the one percent level and shows high predictive power as in the panel B regressions. The last three forecast variables do not show any in-sample predictability.

Out of the five dispersion measures only two show significant predictability. As observed in the no-lag regression of panel B, the dispersion in analyst forecasts of the short rate D^{TBILL} seems to have predictive power at the five percent level with a t-statistic of -2.19. The other statistically significant dispersion measure is D^{CP}. This variable also showed predictive power in the panel B regressions, but more significantly so in this specification with only analyst forecast data. The D^{CP} variable is significant at the one percent level, corresponding to a test statistic of -3.04, and the negative coefficient estimate of -0.0029 suggests that increasing disagreement between forecasters about the future change in corporate profits is often followed by a decoupling of stock and bond returns. This may be coupled with the "flight-to-quality" argument, as increasing uncertainty about corporate profits can be interpreted as uncertainty about the future performance of stocks, which in turn is uncertainty about the stock market volatility. Hence, when there is great uncertainty about how stocks will perform in the near future, investors might be more prone to change portfolio weights from stocks to bonds, thereby decoupling the two asset returns.

The adjusted R^2 for the first specification of table 4 in panel C is 48.18%. A comparison of the performance of this regression to that of the no-lag predictive regression only using historical data, reveals significantly better in-sample predictive power of analyst forecasts compared to historical variables. The R^2 increases with more than 7.5 percentage points, using analyst forecasts compared to historical data. However, the predictive regression incorporating both historical and analyst forecast variables shows the best performance with a R^2 of 50.99%, outperforming the first specification in panel C by 2.81 percentage points.

Lastly, the second specification of the predictive regression in panel C, which includes the lagged correlation variable performs well. The regression shows similarity to the no-lag regression in terms of significance of explanatory variables. As in both panel A and B, the lagged variable shows high statistical significance with a t-statistic of 7.32. This variable explains by far the most variation in realized correlation. Similar to the no-lag regression from panel C, the variables F^{VIX}, F^{TBOND}, F^{RGDP}, and D^{CP} all display significant predictive power. However, the forecast variable of the long rate F^{TBOND} and the growth rate in real GDP F^{RGDP}, as well as the dispersion in forecasts of corporate profits D^{CP} are all less significant and only so at the five percent level. Additionally, it seems like the lagged correlation variable captures the explanatory power of the dispersion in forecasts of the short rate D^{TBILL}.

The overall performance of this specification is really good with an adjusted R^2 of 56.33%. This is better than the 53.60% of the historical variables only regression and is almost as good as the historical and forecast variable regression in panel B of 56.47%. Even though the adjusted R^2 of the panel C regression is 14 basis points lower than the panel B regression, one can argue that the former model is better due to its higher degree of simplicity as it has five less explanatory variables. Hence, the results of table 4 shows that in-sample, mean consensus forecast variables as well as dispersion measures do have predictive power of realized stock-bond correlation. Also, a predictive regression model, using only analyst forecast variables and dispersion measures, outperforms models that use only historical or historical and analyst forecast variables in terms of simplicity. Lastly, past realized correlation seems to be the best in-sample predictor for next period's realized correlation.

5.2.2 Predictive regressions of stock volatility

The second input needed for the variance-covariance matrix is the standard deviation of stock returns. Thus, in the same way as to predict realized correlation, different specifications of the following predictive regression of volatility is run.

Equation 55

$$\sigma_{t+1}^{S} = \hat{\alpha} + \hat{\beta}\sigma_{t}^{S} + \hat{\varphi}\boldsymbol{H}_t + \hat{\delta}\boldsymbol{F}_t + \hat{\gamma}\boldsymbol{D}_t + \epsilon_{t+1}$$

Where σ_{t+1}^{S} is the realized standard deviation of daily stock returns in month t+1, determined as in equation 28. Also, σ_{t}^{S} is the lagged realized standard deviation, $\boldsymbol{H}_t$ is a collection of historical values of the explanatory variables available at time t, $\boldsymbol{F}_t$ is a collection of analyst mean consensus forecasts of the macroeconomic variables and lastly $\boldsymbol{D}_t$ is a collection of the dispersion measures of the analyst forecasts.

In order to follow the same methodology as used in the in-sample predictive regression for realized correlation, a base line model using only historical variables is run. Again, it is interesting to see how well realized values of macroeconomic factors explain the time variance in stock returns. To test such a specification the following regression will be run.

Equation 56

$$\sigma_{t+1}^{S} = \hat{\alpha} + \hat{\beta}\sigma_{t}^{S} + \hat{\varphi}\boldsymbol{H}_t + \epsilon_{t+1}$$

The results of this regression are presented in panel A of table 5, which includes coefficient estimates, t-statistics, and adjusted R^2.

In the first specification of panel A all but one of the historical variables are statistically significant at the five percent level. Only last month's realized value of corporate profits H^{CP} is insignificant. On the other hand, realized inflation H^{CPI} displays significant predictive power of standard deviation of stock returns with a coefficient estimate of -0.0128 and t-statistic of -2.37. The interpretation is that in highly inflationary periods, the volatility of stock returns seems to be low.

The same is true for the H^{RGDP} variable, where high growth in real GDP is associated with a decreasing volatility of stock returns.

Lastly, it seems like the realized short rate and long rate have opposite effects on stock volatility. The historical short rate H^{TBIL} has a negative estimated coefficient, which means that increases in the short rate is followed by decreasing stock volatility, whereas increasing long rates H^{TBOND} cause a surge in the standard deviation of daily stock returns. Intuitively, if long rates are high, this may be an expression of high uncertainty about the future, as investors demand a high return to invest their money in a high duration product. This uncertainty may translate into increased stock return uncertainty and thereby volatility.

The no-lag specification of the predictive regression in panel A explains only 25.65% of the variation in realized stock volatility, but almost all of the included variables seem to have some predictive power, although the low R^2 indicates that some factors are missing.

In the second specification in panel A, the realized standard deviation from last month is included as an explanatory variable. This increases the adjusted R^2 considerably to explain 57.53% of the variation. Hence, simply including last period's realized value of stock volatility, alone increase the performance of the regression by 31.88 percentage points. This reveals a high degree of autocorrelation in stock returns, as the last month's realized value has great explanatory power of next month's realized volatility. This is also visible from the large t-statistic 14.74 of the lagged variable.

Interestingly, including the lagged variable seem to capture most of the explanatory power in the collection of historical variables. Both the inflation rate H^{CPI}, the short rate H^{TBIL}, and the long rate H^{TBOND} no longer display predictive power at the five percent level. However, the growth in the real GDP variable H^{RGDP} is still statistically significant with a test statistic of -4.37. Also, the corporate profits variable H^{CP} remains insignificant at the five percent level. Hence, using historical data only, the lagged value of the standard deviation of stock returns is, by far, the best predictor of stock volatility explaining more than half of the explained variation by itself.

Considering how well the in-sample regression performs using historical data including the lagged variable, it is interesting to investigate whether analyst forecast data in terms of mean consensus estimates and dispersion measures can add some explanatory power to the model.

Stock volatility

Table 5 reports the 1-month horizon in-sample predictive regressions of stock return standard deviation onto a constant and historical values ($\boldsymbol{H}_t$) as well as standardized measures of the mean consensus ($\boldsymbol{F}_t$) and dispersion ($\boldsymbol{D}_t$) in analyst forecasts of the inflation rate (CPI), the short rate (TBILL), long rate (TBOND), change in corporate profits (CP), change in real GDP (RGDP), and implied volatility (VIX). The table reports coefficient estimates with t-statistics in brackets and the adjusted R^2 of the regression. The critical value of the t-statistic at which the null hypothesis of $\beta = 0$ is rejected at a significance level of 5% is $|t| > 1.96$
**** denotes significance at the 1% level   ** denotes significance at the 5% level   * denotes significance at the 10% level*

Cons. (*t-stat*)	Lag (*t-stat*)	H^{CPI} (*t-stat*)	H^{TBILL} (*t-stat*)	H^{TBOND} (*t-stat*)	H^{CP} (*t-stat*)	H^{RGDP} (*t-stat*)	F^{VIX} (*t-stat*)	F^{CPI} (*t-stat*)	F^{TBILL} (*t-stat*)	F^{TBOND} (*t-stat*)	F^{CP} (*t-stat*)	F^{RGDP} (*t-stat*)	D^{CPI} (*t-stat*)	D^{TBILL} (*t-stat*)	D^{TBOND} (*t-stat*)	D^{CP} (*t-stat*)	D^{RGDP} (*t-stat*)	Adj. R^2 (%)
Panel A: Historical data only																		
0.0460***		-0.0128**	-0.0026**	0.0058***	-0.0002	-0.0247***												25.65
(8.97)		(-2.37)	(-1.87)	(3.15)	(-1.01)	(-8.82)												
0.0141***	0.6573***	0.0039	-0.0011	0.0025*	-0.0003*	-0.0102***												57.53
(3.19)	(14.74)	(0.91)	(-1.00)	(1.77)	(-1.72)	(-4.37)												
Panel B: Historical data as well as analyst mean forecasts and dispersion measures																		
0.0175		0.0047	0.0068	-0.0021	-0.0007***	-0.0101***	0.0021***	0.0012	-0.0081	0.0024	0.0000	0.0002	0.0186	0.0695***	-0.1040***	-0.0001	0.0001	56.33
(1.26)		(1.06)	(1.47)	(-0.69)	(-3.71)	(-3.78)	(9.70)	(0.33)	(-1.60)	(0.63)	(0.08)	(0.56)	(1.24)	(2.82)	(-3.56)	(-1.43)	(0.65)	
0.0273**	0.4297***	0.0052	0.0075*	-0.0018	-0.0005***	-0.0085***	0.0007*	-0.0013	-0.0093*	0.0030	-0.0001	0.0002	0.0196	0.0635***	-0.0964***	-0.0001*	0.0001	59.13
(2.00)	(4.47)	(1.22)	(1.69)	(-0.60)	(-3.01)	(-3.23)	(1.92)	(-0.37)	(-1.90)	(0.82)	(-0.48)	(0.58)	(1.35)	(2.66)	(-3.40)	(-1.77)	(0.57)	
Panel C: Analyst mean forecasts and dispersion measures only																		
0.0011							0.0022***	0.0082**	0.0006	-0.0042*	0.0000	0.0001	0.0242	0.0695***	-0.1053***	-0.0001	0.0001	51.76
(0.08)							(11.26)	(2.37)	(0.40)	(-1.91)	(0.24)	(0.24)	(1.63)	(3.21)	(-3.64)	(-1.46)	(0.97)	
0.0175	0.5296***						0.0005	0.0033	-0.0002	-0.0019	-0.0001	0.0001	0.0254*	0.0560***	-0.0927***	-0.0001*	0.0001	56.35
(1.29)	(5.53)						(1.25)	(0.95)	(-0.14)	(-0.88)	(-0.43)	(0.36)	(1.80)	(2.70)	(-3.36)	(-1.95)	(0.84)	

Source: Own contribution

Hence, the predictive regression from equation 55, which includes the collections of both historical and forecast variables are run, and results are presented in panel B of table 5.

In the first specification without the lagged variable, only two of the five historical variables seem to have predictive power compared to four out of five in the base line regression from panel A. The first is the realized corporate profits H^{CP}, which is significant at the one percent level with a test statistic of -3.71. The coefficient estimate is -0.007 and since it is negative, it suggests that a positive change in corporate profits often is followed by a reduced volatility in stock returns. The second is the realized change in real GDP H^{RGDP}, which also shows statistical significance in the panel A regressions. This one is also significant at the one percent level. The last three variables H^{CPI}, H^{TBILL}, and H^{TBOND}, which all were significant in the similar specification in panel A, no longer show predictive power at a five percent level.

Among the collection of mean consensus forecasts $\boldsymbol{F_t}$ only the implied stock market volatility variable F^{VIX} shows explanatory power. With a t-statistic of 9.70 the coefficient estimate of 0.0021 is highly significant, even at the one percent level. The interpretation of this is straight forward. When the implied volatility from option prices on the S&P500 index increase, which means that the investors assess the current and near future stock market volatility to have increased, then this is most naturally followed by an actual increase in the realized stock volatility. All the remaining five consensus mean forecast variables display no statistical significance.

Adding dispersion in analyst forecasts to the in-sample regression seems to add explanatory power. The variable describing dispersion in analyst forecasts of the short rate D^{TBILL} is statistically significant at the one percent level with a test statistic of 2.82. The positive coefficient estimate of 0.0695 indicates that when there is a higher degree of misalignment in analysts' forecast of the short rate, then next period's realized stock volatility tend to rise. Dispersion in forecasts of the long rate D^{TBOND} has a different impact on the stock volatility. The negative coefficient estimate of -0.1040, which is significant at the one percent level, suggests that when analysts tend to disagree on the future level of the long rate, the standard deviation of stock returns decreases in the following month. The last three dispersion measures show no explanatory power.

The overall performance of the first specification of panel B is good, posing an adjusted R^2 of 56.33%. This is a considerable increase from the no-lag configuration in panel A, which only explained 25.65% of the variation in realized standard deviation of stock returns. Hence, it does look like the consensus mean and dispersion measures of analyst forecast add power to the predictive in-sample regression.

In the second specification of panel B, the lagged variable is added to the regression line. Again, the lagged variable shows high statistical significance with a test statistic of 4.47 and a high coefficient estimate of 0.4297. The explanatory power of the remaining variables is very much in line with the no-lag configuration. The variables H^{CP}, H^{RGDP}, D^{TBILL}, and D^{TBOND} all show statistical significance at the one percent level like in the other regression. However, the implied volatility variable F^{VIX} is no longer significant, as the lagged variable seem to capture all the explanatory power of F^{VIX}. By including its own lag, adjusted R^2 increases by 2.8 percentage points to 59.13%. However, this is only 1.6 percentage points better than the historical variables only regression from panel A. Also, a total of 12 variables in the panel B regression are insignificant, so even though it seems that analyst forecast data add explanatory power to the predictive regression, maybe there is an even simpler model explaining the same amount of variation in stock volatility.

In that line of thinking, a configuration of the predictive in-sample regression is investigated, which excludes historical variables, thereby only using analyst forecast data to predict realized stock volatility. To test such a specification the following regression will be run.

Equation 57

$$\sigma_{t+1}^{S} = \hat{\alpha} + \hat{\beta}\sigma_{t}^{S} + \hat{\delta}\boldsymbol{F}_t + \hat{\gamma}\boldsymbol{D}_t + \epsilon_{t+1}$$

Regression results are presented in panel C of table 5. The first regression line in panel C, the no-lag specification, shows statistical significance of two of the six mean consensus forecast variables. Like in the no-lag regression from panel B, the implied volatility variable F^{VIX} is statistically significant at the one percent level. More interestingly, when historical variables are excluded from the predictive regression, the mean consensus of inflation F^{CPI} suddenly shows up significant, with at test statistic of 2.37. The coefficient estimate is positive, indicating that when analysts expect an increasing inflation rate, stock returns tend to vary more in the following month.

The last two statistically significant variables are dispersion measures. Similar to the no-lag regression in panel B, the dispersion in forecasts of the short and long rate, D^{TBILL} and D^{TBOND} respectively, are significant at the one percent level. The remaining mean consensus and dispersion measures show no explanatory power.

The no-lag, analyst forecast data only regression performs well overall with an adjusted R^2 of 51.76%. This is significantly better than the first specification in panel A with historical variables only as it explains more than twice the amount of variation in realized standard deviation of stock returns.

Lastly, the analyst forecast variables only regression is also performed including the lagged variable. As in the two other configurations with the lagged variable, the variable shows up highly significant with a coefficient estimate of 0.5296 that is significant at the one percent level. Interestingly, the lag seems to capture much of the predictive power from variables that were significant in the no-lag specification. As a result, the implied volatility variable F^{VIX} as well as mean consensus forecast of inflation F^{CPI} no longer show up significant in the regression results. However, the dispersion in analyst forecasts of the short and long rate remain significant at the one percent level. These are the only two variables together with the lag that have statistically predictive power.

The adjusted R^2 of this regression is 56.35%, which is better than the no-lag configuration. However, surprisingly it is slightly worse than the second specification of historical variables only regression in panel A and also worse than the regression including all variables. However, their performance is quite similar in terms of adjusted R^2.

It is notable that the lagged variable seems to be the best in-sample predictor of next month's realized standard deviation of stock returns across all three configurations. However, it cannot be concluded that a predictive regression using only analyst forecast variables performs better than models including historical variables. It does seem like analyst forecasts add explanatory power to the regressions over and above historical data.

5.2.3 Predictive regressions of bond volatility

The last input needed for the variance-covariance matrix is the standard deviation of bond returns. Thus, in the same way as to predict realized stock volatility, different specifications of the following predictive regression of bond volatility are run.

Equation 58

$$\sigma_{t+1}^{B} = \hat{\alpha} + \hat{\beta}\sigma_{t}^{B} + \hat{\varphi}\boldsymbol{H}_{t} + \hat{\delta}\boldsymbol{F}_{t} + \hat{\gamma}\boldsymbol{D}_{t} + \epsilon_{t+1}$$

Where σ_{t+1}^{B} is the realized standard deviation of daily bond returns in month t+1, determined as in equation 28. Also, σ_{t}^{B} is the lagged realized standard deviation, $\boldsymbol{H}_{t}$ is a collection of historical values of the explanatory variables available at time t, $\boldsymbol{F}_{t}$ is a collection of analyst mean consensus forecasts of the macroeconomic variables and lastly $\boldsymbol{D}_{t}$ is a collection of the dispersion measures of the analyst forecasts.

The first configuration of equation 58 to be tested includes only the collection of historical variables $\boldsymbol{H}_{t}$. This predictive regression, which is listed below, allows for investigation of the predictive power of realized macroeconomic variables in predicting bond volatility. The following regression is run.

Equation 59

$$\sigma_{t+1}^{B} = \hat{\alpha} + \hat{\beta}\sigma_{t}^{B} + \hat{\varphi}\boldsymbol{H}_{t} + \epsilon_{t+1}$$

The results from this regression are presented in panel A of table 6. In the no-lag specification from panel A, three of the five variables display significant predictive power. It is last month's realized values of the short rate H^{TBILL}, the long rate H^{TBIND}, as well as growth in real GDP H^{RGDP}. All three variables are statistically significant at the one percent level, which also was the case in the predictive regression of stock volatility. The last two variables remain insignificant and the in-sample regression delivers an overall weak performance with an adjusted R^2 of 23.33%.

Including the lagged realized bond volatility increases the performance. The variable shows up as highly significant with a t-statistic of 10.50 and a large coefficient estimate of 0.5382. The lagged variable seems to capture some of the predictive power from the other variables as H^{TBILL} no longer is significant and H^{TBOND} only remains significant at the five percent level. However, the past growth rate in real GDP is still statistically significant at the one percent level.

Like the stock volatility regression, the negative coefficient estimate suggests that a growing economy most often is followed by a reduction in the volatility of bond returns. The adjusted R^2 increases to 44.42% when the lag is included. This is not quite as good as in the regressions of stock volatility and stock-bond correlation, with adjusted R^2 of 57.53% and 53.60% respectively.

It is interesting to see if the relatively bad performance of the in-sample regressions from panel A can be fixed by including analyst forecast data in terms of the collection of mean consensus and dispersion measures. This is equivalent to the regression line presented in equation 58. Looking at the no-lag specification in panel B of table 6, one can see, that the realized value of the short rate H^{TBILL} is the only historical variable that is statistically significant.

However, several of the forecast variables seem to have predictive power. The implied volatility of the stock market F^{VIX} seems to explain some of the variation in bond returns. This observation can be described by the volatility spillover phenomenon between stock and bond markets, which was described in the stylized facts section of the analysis. Also, the mean consensus forecast of growth in real GDP F^{RGDP} seems to have explanatory power, as it is significant at the five percent level.

Bond volatility

Table 6 reports the 1-month horizon in-sample predictive regressions of bond return standard deviation onto a constant and historical values ($\boldsymbol{H_t}$) as well as standardized measures of the mean consensus ($\boldsymbol{F_t}$) and dispersion ($\boldsymbol{D_t}$) in analyst forecasts of the inflation rate (CPI), the short rate (TBILL), long rate (TBOND), change in corporate profits (CP), change in real GDP (RGDP), and implied volatility (VIX). The table reports coefficient estimates with t-statistics in brackets and the adjusted R^2 of the regression. The critical value of the t-statistic at which the null hypothesis of $\beta = 0$ is rejected at a significance level of 5% is $|t| > 1.96$
**** denotes significance at the 1% level   ** denotes significance at the 5% level   * denotes significance at the 10% level*

Cons. (*t-stat*)	Lag (*t-stat*)	H^{CPI} (*t-stat*)	H^{TBILL} (*t-stat*)	H^{TBOND} (*t-stat*)	H^{CP} (*t-stat*)	H^{RGDP} (*t-stat*)	F^{VIX} (*t-stat*)	F^{CPI} (*t-stat*)	F^{TBILL} (*t-stat*)	F^{TBOND} (*t-stat*)	F^{CP} (*t-stat*)	F^{RGDP} (*t-stat*)	D^{CPI} (*t-stat*)	D^{TBILL} (*t-stat*)	D^{TBOND} (*t-stat*)	D^{CP} (*t-stat*)	D^{RGDP} (*t-stat*)	Adj. R^2 (%)
Panel A: Historical data only																		
0.0089***		-0.0012	-0.0012***	0.0020***	0.0001	-0.0030***												23.33
(10.97)		(-1.38)	(-5.41)	(6.73)	(1.53)	(-6.64)												
0.0047***	0.5382***	-0.0011	-0.0004*	0.0007**	0.0000	-0.0015***												44.42
(5.84)	(10.50)	(-1.52)	(-1.75)	(2.52)	(1.07)	(-3.71)												
Panel B: Historical data as well as analyst mean forecasts and dispersion measures																		
0.0085***		0.0007	-0.0023***	0.0005	0.0000	-0.0007	0.0002***	0.0000	0.0015*	-0.0001	0.0000	0.0001**	0.0079***	0.0125***	-0.0059	0.0000***	0.0000*	47.08
(3.53)		(0.97)	(-2.91)	(0.85)	(-0.07)	(-1.55)	(5.06)	(-0.05)	(1.76)	(-0.09)	(-1.31)	(2.29)	(3.03)	(2.95)	(-1.18)	(-3.08)	(-1.94)	
0.0073***	0.2946***	0.0002	-0.0011	-0.0002	0.0000	-0.0007	0.0001***	-0.0001	0.0005	0.0005	0.0000	0.0001**	0.0059**	0.0107***	-0.0085*	0.0000**	0.0000	50.51
(3.12)	(4.50)	(0.28)	(-1.35)	(-0.31)	(-0.05)	(-1.51)	(2.91)	(-0.11)	(0.57)	(0.76)	(-1.30)	(2.03)	(2.34)	(2.60)	(-1.74)	(-2.53)	(-1.51)	
Panel C: Analyst mean forecasts and dispersion measures only																		
0.0084***							0.0002***	0.0002	-0.0008***	0.0000	0.0000*	0.0001**	0.0089***	0.0205***	-0.0112**	0.0000***	0.0000*	45.53
(3.65)							(6.41)	(0.41)	(-3.33)	(0.11)	(-1.71)	(2.15)	(3.59)	(5.69)	(-2.34)	(-3.52)	(-1.75)	
0.0064***	0.3226***						0.0001***	0.0002	-0.0006**	0.0001	0.0000	0.0001*	0.0065***	0.0146***	-0.0113**	0.0000**	0.0000	50.42
(2.85)	(5.37)						(3.89)	(0.44)	(-2.56)	(0.22)	(-1.46)	(1.70)	(2.71)	(4.04)	(-2.46)	(-2.52)	(-1.42)	

Source: Own contribution

Notable is that three out of five dispersion measures seem to predict next period's realized bond volatility well. Dispersion in analyst forecasts of the inflation rate and the short rate are both significant at the one percent level, and both has a sizeable positive coefficient estimate of 0.0079 and 0.0125, respectively. This indicates that when disagreement between forecasters increase in terms of expectations about the future inflation rate and the short rate, then next periods bond returns tend to vary more. The dispersion measure D^{CP} is also statistically significant at the one percent level, however its coefficient estimate is so close to zero that it does not have any economic importance.

It is clear from the performance of this specification, that analyst forecast data do add explanatory power over and beyond that of historical variables. The adjusted R^2 is 47.08%, hence the collection of mean consensus forecasts and the dispersion measures more than doubles the explained variation in bond returns compared to the no-lag specification in panel A.

Including the lagged variable together with historic and forecast data leaves all the historical variables insignificant. The lagged variable has high explanatory power with a test statistic of 4.50 and an estimated coefficient of 0.2946. Hence, it seems like the lag captures all the explanatory power in historical information. All the mean consensus forecast and dispersion variables, that were significant in the no-log regression remain so in the second specification of the predictive regression in panel B. However, D^{CPI} and D^{CP} becomes less significant although still at a five percent level. Also, even though F^{VIX}, F^{RGDP}, and D^{CP} all are statistically significant, the size of the coefficient estimates is not significant in an economical sense, as they contribute very little to the explained variation in bond returns. The overall performance of this configuration increases to 50.51% when the lagged variable is included in the regression.

As all the historical variables become insignificant in the second specification of the regression in panel B, it is interesting to see whether analyst forecast data alone can explain the realized standard deviation of bond returns. To test so, the following regression will be performed.

Equation 60

$$\sigma_{t+1}^{B} = \hat{\alpha} + \hat{\beta}\sigma_{t}^{B} + \hat{\delta}\boldsymbol{F}_{t} + \hat{\gamma}\boldsymbol{D}_{t} + \epsilon_{t+1}$$

In panel C of table 6, seven of the eleven analyst forecast variables display statistical significance. Only the three mean consensus forecast variables, F^{CPI}, F^{TBOND}, and F^{CP} as well as the dispersion measure D^{RGDP} remain insignificant. Hence, most of the forecast variables seem to have good predictability of future bond return volatility. The adjusted R^2 of the regression is 45.53%, which is considerably better than the historical model, but slightly worse than the historical and forecast model.

The last specification of the predictive in-sample regression of standard deviation of bond returns include forecast variables only and its lag. This regression performs very similar to the no-lag specification mentioned above. Noteworthy is that the lagged variable is highly significant with a t-statistic of 5.37 and a large positive coefficient estimate. Also, all the statistically significant variables from the no-lag regression remain so, except from consensus mean forecast of the growth in real GDP F^{RGDP}.

Adding the lagged variable adds another 4.89 percentage points to the adjusted R^2 so that the predictive regression explains a total of 50.42%. This is considerably more explanatory power than the similar specification from panel A, which only explained 44.42%. Interesting to note, the forecast variables only regression explains only nine basis points less of the variation in bond volatility compared to the longer in-sample regression from panel B. Hence, as in the predictive regression of stock-bond correlation, it can be argued that the regression from panel C is the best as it performs almost as good as the one from panel B but contains five less explanatory variables making it considerably simpler.

5.2.4 Performance of the predictive regressions

All three in-sample regressions predicting stock volatility, bond volatility, and correlation, respectively, perform well. In two of the three cases it is argued that the panel C regressions, that is, the in-sample regressions only using the collection of mean consensus forecasts and dispersion measures, deliver the best performance in terms of simplicity and adjusted R^2. Even though a lot of the mean consensus forecast variables and dispersion measures are statistically significant, the lagged variable seems to be the single best in-sample predictor of next month's realized value of correlation, stock and bond volatility, respectively. This is clear from figure 11 below, which shows the aggregate contribution of mean consensus and dispersion variables, as well as the lagged variable to adjusted R^2.

In the regression of stock volatility, the lagged variable explains 59% of the variation alone. The dispersion in analyst forecasts $\boldsymbol{D_t}$ also have a large degree of explanatory power on an aggregate basis, as it explains 34%. Surprisingly mean consensus variables only contribute with 6% of adjusted R^2.

A similar pattern applies to the predictive regression of correlation. However, the degree of predictive power of mean consensus and dispersion measures is reversed. Here mean consensus explains a third of the variation and dispersion only 10%. Again, the lagged variable explains more than half of the variation in stock-bond correlation.

The results paint a different picture of the variables' contribution to R^2 in the bond volatility regression. Here, the lagged variable and the collection of mean consensus and dispersion measures explain almost one third each.

Figure 11 - Contribution to adjusted R^2 for panel C predictive regressions of volatilities and correlation of stock and bond returns

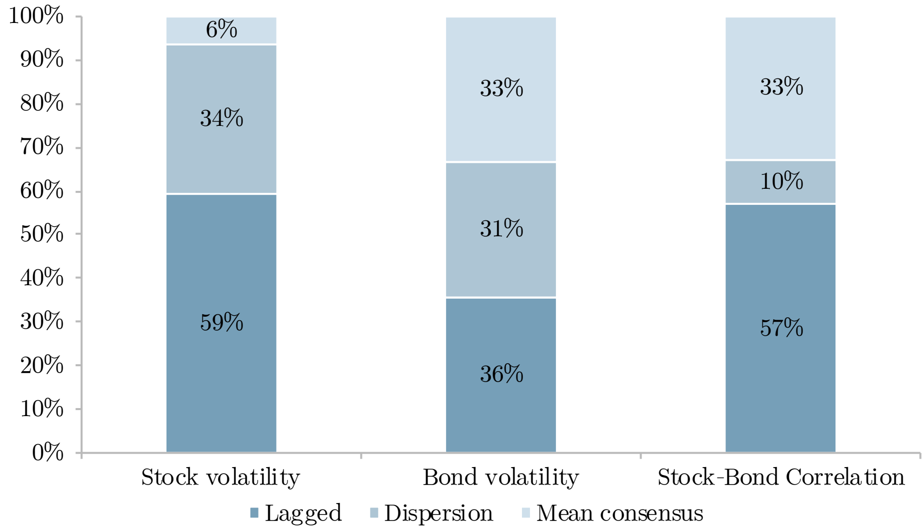

Source: Own contribution

Additionally, it seems like the proposed predictive models is better at forecasting stock-bond correlation and stock volatility than bond volatility. The in-sample regressions of bond volatility consistently underperform in terms of R^2, compared to the models of stock volatility and correlation. On the other hand, the collection of forecasts of economic variables seems to explain a higher proportion of the explained variation in bond volatility compared to stock volatility and correlation. The analyst forecast variables explain a total of 64% of variation in bond volatility whereas it is only 41% and 43% for the other two regressions.

5.3 Out-of-sample predictability analysis

In order to evaluate the out-of-sample performance of the different models presented in the previous section, several performance metrics are calculated. A description of how these metrics are calculated is presented in the data and methodology section. The measures provide standardized metrics to compare the performance of strategies that are based on analyst forecasts as well as historical data to predict volatilities and correlation used for creating the optimal portfolio.

In order to measure the performance, portfolio weights for each strategy needs to be calculated. For every month t, the portfolio weights w^S and w^B, which denote the weights in stocks and bonds, are obtained by solving equation 3. The standard deviation and correlation of stock and bond returns for the next period, ultimately the forecasted variance-covariance matrix needed for this minimization problem, is estimated using a rolling-window procedure. The length of the estimation window is five years of monthly data, that is, 60 months. It is denoted by τ and $\tau < T$, where T is the total number of months in the sample period. In empirical finance, it is quite common to rely on rolling sample periods of five years of monthly data to estimate time-varying variances and covariances (Andersen *et al.*, 2006). The data in the estimation window, which includes macro variables and realized volatilities or correlation, is then used in the regression model to forecast next month's correlation or stock or bond volatility. This in turn is used for determining the portfolio weights for that respective month for all the different strategies.

After the portfolio weights for strategy k in month t are determined, the monthly out-of-sample portfolio return over the period $[t, t+1]$ is calculated as:

Equation 61

$$r_{t+1}^{k} = w_t^{S,k} * r_{t+1}^{S} + w_t^{B,k} * r_{t+1}^{B}$$

Where r_{t+1} is the stock or bond return over month t to t+1 and $w_t^{S,k}$ and $w_t^{B,k}$ denote the weight in stocks and bonds of strategy k at month t.

The portfolios are rebalanced every month by rolling over the estimation window by one month, thereby including the next month and dropping the first observation. This procedure is continued until there is no more data left. It yields a time series of out-of-sample portfolio returns of length $T - \tau$, that is 234 portfolio returns, based on the given strategy k. Next, all the performance metrics are calculated from the vector of portfolio returns. The portfolio standard deviation, Sharpe ratio, and Certainty Equivalent are tested for statistical significance using the non-parametric stationary bootstrap-method described in the data and methodology section. The portfolio turnover and opportunity cost measures are calculated in order to complement the results of statistical significance of portfolio performance with measures of economic significance.

5.3.1 Hypothesis testing for difference in portfolio performance

Not only is different performance metrics for all strategies provided in order to compare their performance, but the statistical significance of the performance of the various portfolios is also tested. More specifically, it is the statistical difference in portfolio standard deviation, Sharpe ratio, and certainty equivalent of the benchmark strategies that is tested. In other words, the following hypothesis will be put to the test. Portfolio volatility, Sharpe ratio, and certainty equivalent of a particular strategy k is worse than that of the benchmark strategy bm. This gives three null hypotheses of the difference in performance.

Equation 62

$$H_0: \hat{\sigma}^{bm} - \hat{\sigma}^{k} \leq 0$$

Equation 63

$$H_0: \widehat{SR}^{k} - \widehat{SR}^{bm} \leq 0$$

Equation 64

$$H_0: \widehat{CEQ}^{k} - \widehat{CEQ}^{bm} \leq 0$$

The alternative hypothesis H_a states that strategy k does outperform the benchmark strategy bm in terms of significantly lower portfolio standard deviation, higher Sharpe ratio, and higher certainty equivalent. Note that the null hypothesis of the difference in volatility in equation 62 is different from the other two null hypotheses. The reason is that a low portfolio standard deviation is considered good, whereas the higher the Sharpe ratio or certainty equivalent the better.

The stationary bootstrap of Politis and Romano (1994) provides an empirical distribution of the difference in all three performance measure from the 10,000 resamples, which then can be used to calculate the p-values.

The test is a one-sided t-test as the null hypothesis states that strategy k performs worse than the benchmark strategy. A two-sided test has not been chosen, since the purpose of this paper is to create asset allocation strategies that perform better than simple benchmark strategies. It is reasonable to think that carefully selecting macro variables and gathering forward-looking analyst data to form portfolio weights should perform at least as well as simply putting 50% in stock and 50% in bonds or alternatively using a moving average to forecast the variance-covariance matrix.

The p-value resulting from the one-sided test can be interpreted as the probability of obtaining a test statistic at least as small as the one observed given that the null hypothesis is true. Hence, there is sufficient evidence against the null hypothesis, if a strategy provides a p-value below five percent, which implies that the result is statistically significant. The next section presents all the strategies considered as well as the relating performance metrics and their statistical significance.

5.3.2 Portfolio strategies to be considered

To measure the performance of portfolios, formed using analyst forecasts, a comparison of the performance of several other portfolio strategies is conducted. These include portfolio strategies based on historical data, based on analyst forecast data, as well as benchmark strategies and other portfolio strategies suggested in the literature. All strategies and their abbreviations are listed in table 7 below, and each strategy is further elaborated upon in the next section.

Table 7 – Presents all portfolio strategies tested in this section and their abbreviations. All strategies below are minimum-variance portfolios expect the benchmark EW

No.	Strategy k	Abbreviation
	Benchmark strategies	
1	Equally-weighted portfolio	EW
2	60-months′ moving-average of realized variance-covariance matrix	MinVarMA
	Historical data only	
3	Based on collection of hisotrical variables	MinVarHt
4	Based on collection of hisotrical variables incl. the lag	MinVarHtLag
	Analyst forecast data only	
5	Collection of mean consensus variables	MinVarFt
6	Collection of mean consensus variables incl. the lag	MinVarFtLag
7	Collection of dispersion variables	MinVarDt
8	Collection of dispersion variables incl. the lag	MinVarDtLag
9	Collection of mean consensus & dispersion variables	MinVarFtDt
10	Collection of mean consensus & dispersion variables incl. the lag	MinVarFtDtLag
	Historical and analyst forecast data	
11	Collection of historical, mean consensus & dispersion variables	MinVarHtFtDt
12	Collection of historical, mean consensus & dispersion variables incl. the lag	MinVarHtFtDtLag
	Other strategies	
13	Regression on lagged volatilities and correlation	MinVarPrevVolCorrEst
14	Last month′s realized variance-covariance matrix	MinVarPrevVolCorr
15	Last month′s realized variances and correlation set to zero	MinVarZeroCorr

Source: Own contribution

As described above, the estimation window is set to 60 months, and given that the sample period runs from July 1994 to December 2018, the first estimation window will be July 1994 to June 1999. This provides portfolio returns for the 234 months' period July 1999 to December 2018. The performance of each group of strategies during this period is presented separately in the following section.

5.3.3 Benchmark strategies

Two simple benchmark strategies have been selected as the anchor point of the comparison of the potential out-of-sample benefits of forecasting variance-covariances matrices. The first is the naive 1/N portfolio also called the equally-weighted portfolio. In the case of only two asset classes, as in this thesis, 50% of wealth is invested in stocks and 50% in bonds at the beginning of the sample period. Weights are rebalanced every month to keep the 50-50 split of allocated capital between stocks and bonds.

In a recent paper, DeMiguel, Garlappi, and Uppal (2009) show that the equally-weighted portfolio performs well compared to other strategies based on mean-variance analysis. They find that the gain from optimal diversification is more than offset by the increase in estimation error. Hence, even though the 1/N portfolio is simplistic, it seems like a good baseline as to compare the performance of portfolios formed by analyst forecast data.

The performance of the equally-weighted strategy in terms of portfolio standard deviation, Sharpe ratio, certainty equivalent, and portfolio turnover as well as opportunity cost is presented in table 8 below.

The equally-weighted portfolio has the highest standard deviation of portfolio returns of all strategies considered. However, this does not come as a surprise since the naive 1/N portfolio does not try to optimize the benefits achieved by diversification. Over the period July 1999 to December 2018, the EW strategy had an annualized standard deviation of 7.71%. The Sharpe ratio was disappointingly low, showing an annualized Sharpe ratio 0.41, while the strategy only delivered an average annual excess return of 4.82%. This is by far the lowest Sharpe ratio compared to the other strategies. The main reason for the poor performance is the relatively high level of risk, as the return is much in line with the other portfolios, although not reported in the table.

Likewise, the certainty equivalent of 2.85%, which states than an investor with a risk aversion constant of one only needs 2.85% in risk-free return to abandon the "higher risk-higher reward" strategy, is the lowest among all strategies. Additionally, it is relative cheap to carry out the 1/N asset allocation strategy compared to other strategies as 2.06% of wealth on average is traded every month for rebalancing the portfolio.

Lastly, the opportunity cost compared to the other benchmark strategy, MinVarMA, is negative 0.13%. This means, that the EW strategy, on average, delivers 0.13% less per month, in annualized terms.

These results are not in line with the findings of DeMiguel, Garlappi, and Uppal (2009) as it does not seem like the equally-weighted strategy forms a desirable portfolio to invest in based on the presented performance metrics.

The second benchmark strategy is called the minimum variance moving-average strategy or MinVarMA. This is another attempt to construct a simplistic portfolio strategy in order to evaluate the more complex strategies in this thesis. The strategy forecasts next period's variance-covariance matrix by taking the 60 months' moving-average of the volatilities and correlation respectively. The procedure of estimating parameters simply by using a historical average is probably the simplest way to account for time-varying volatilities and covariances. The estimated variance-covariance matrix is then used to find the minimum variance portfolio weights. Hence, this provides a good benchmark for all strategies that account for time-varying investment opportunities as well as minimizes the portfolio standard deviation.

Table 8 – This table reports the performance of the first 10 strategies considered. Portfolio standard deviation, Sharpe ratio, and certainty equivalent measures are annualized, whereas portfolio turnover shows the average monthly percentage of wealth traded at the rebalancing date, and opportunity cost displays the annualized average monthly opportunity cost. All measures are displayed in percent except for the Sharpe ratio. A risk aversion of one is employed for the CEQ and OC measure. The numbers in parentheses are p-values for the one-sided test of difference in performance. The upper p-value compares EW and strategy k and the lower compares MinVarMA and strategy k.

Type	Strategy k	σ^k	SR^k	CEQ^k	PT^k	OC^k_{EW}	$OC^k_{MinVarMA}$
Benchmark	*EW*	7.71	0.41	2.85	2.06		-0.13
	MinVarMA	3.32	0.99	3.22	1.13	0.13	
		(0.0000)	*(0.0159)*	*(0.4194)*			
Historical data only	*MinVarHt*	3.43	0.98	3.30	4.66	0.26	0.13
		(0.0000)	*(0.0065)*	*(0.3945)*			
		(0.8930)	*(0.5275)*	*(0.3480)*			
	MinVarHtLag	3.45	0.95	3.22	6.27	0.05	-0.08
		(0.0000)	*(0.0090)*	*(0.4122)*			
		(0.9620)	*(0.7235)*	*(0.5020)*			
Analyst forecast data only	*MinVarFt*	3.46	1.00	3.38	4.73	0.29	0.16
		(0.0000)	*(0.0087)*	*(0.3789)*			
		(0.9744)	*(0.4226)*	*(0.1447)*			
	MinVarFtLag	3.48	0.94	3.22	6.03	0.13	0.00
		(0.0000)	*(0.0122)*	*(0.4157)*			
		(0.9907)	*(0.8064)*	*(0.5136)*			
	MinVarDt	3.41	0.95	3.19	5.47	0.09	-0.03
		(0.0000)	*(0.0075)*	*(0.4200)*			
		(0.9057)	*(0.7076)*	*(0.5641)*			
	MinVarDtLag	3.39	0.90	2.99	6.14	-0.11	-0.24
		(0.0000)	*(0.0157)*	*(0.4700)*			
		(0.9436)	*(0.9340)*	*(0.8860)*			
	MinVarFtDt	3.54	0.95	3.31	8.43	0.22	0.09
		(0.0000)	*(0.0059)*	*(0.3901)*			
		(0.9794)	*(0.7077)*	*(0.3807)*			
	MinVarFtDtLag	3.54	0.98	3.42	9.80	0.33	0.20
		(0.0000)	*(0.0044)*	*(0.3648)*			
		(0.9894)	*(0.5505)*	*(0.2260)*			

Source: Own contribution

The MinVarMA strategy provides the second lowest portfolio standard deviation of 3.32% and the difference from the EW strategy is statistically significant as the p-value is well below the one percent level. However, it is not surprising that the volatility of portfolio returns is lower since the MinVarMA strategy tries to minimize this metric. The strategy also delivers a high return per unit of risk, that is, Sharpe ratio. The Sharpe ratio is 0.99, which is in top five of all portfolios and the difference to the equally-weighted portfolio is significant at the five percent level.

The certainty equivalent return is only slightly higher than the other benchmark strategy and it does not seem to be statistically different from the EW portfolio's CEQ return. An interesting result to note is the very small portfolio turnover of the moving-average benchmark portfolio. Following the MinVarMA strategy would on average demand rebalancing of 1.13% of the wealth per month. This means that not only does the EW portfolio deliver higher risk, lower risk-adjusted return, and lower certain equivalent return, but also cost more money to carry out in terms of transaction costs.

Worth noting is the performance of a pure investment in stocks and bonds over the same period, which posted a Sharpe ratio of 0.22 for stocks and 0.88 for investing in bonds only, however not shown in the table. The EW benchmark strategy performed worse than simply putting all the money in bonds whereas the MinVarMA strategy performed better than both pure asset class investments. Such a result highlights the importance of being able to forecast the volatilities and correlation between the two asset classes. Choosing portfolios weights arbitrarily and thereby ignoring the time-varying volatilities and correlations, such as the equally-weighted portfolio, the strategy may perform worse than simply investing in one of the assets. Also, not reported in the table is the Sharpe ratio of a perfect foresight strategy. Such as portfolio would have yielded a Sharpe ratio of 1.15 which again underlines the importance of good forecasts of the co-movement in stock and bond returns.

All the performance measures collectively suggest that the MinVarMA strategy is the best of the two benchmark portfolios. It beats the EW strategy in all facets with a higher return and lower volatility. It will be interesting to see if any of the strategies can outperform the EW strategy but even more interesting to see if they can outperform the MinVarMA benchmark portfolio. The performance of each of the strategies and a comparison to both benchmark portfolios will be provided in the next section.

5.3.4 Portfolios based on the historical data only

The first set of strategies to be analyzed is based on predictive regressions from the in-sample analysis that uses historical values of a number of macroeconomic variables which have shown to affect stock-bond correlation in the literature. These two strategies correspond to the predictive regressions performed in panel A of table 4, 5, and 6.

The first strategy is called MinVarHt and it estimates the sample variance-covariance matrix based on a regression on the historical values of inflation, the short rate, long rate, growth in corporate profits as well as growth in real GDP. The portfolio standard deviation of 3.43% is significantly better than the EW strategy but slightly worse than MinVarMA. Likewise, the Sharpe ratio of 0.98 is statistically significant compared to the equally-weighted portfolio but indistinguishable from MinVarMA with a p-value of 52.75%.

However, in terms of the certainty equivalent return, the MinVarHt performs better than both the benchmark portfolios. An investor, with a risk aversion constant of one, would forego the risky strategy if she could have at least 3.3% in risk-free return, which is higher than both benchmark portfolios. However, the difference in CEQ of the EW and MinVarMA is not statistically significant. Also, the strategy will be more expensive to implement in terms of transaction costs, as the average portfolio turnover is 4.66%, which is economically higher than the benchmark strategies, but relative average compared to the remaining portfolios.

The opportunity cost shows the economic significance of the difference in performance, as the measure shows how much the benchmark strategy should have yielded over and above strategy k's return, in order for an investor with risk aversion one to be indifferent in terms of utility. Since MinVarHt has a positive opportunity cost of 0.20% in relation to the equally-weighted portfolio and 0.08% in relation to the moving-average benchmark portfolio, the strategy seems to perform significantly better than the EW strategy and almost in line with the MinVarMA strategy.

The second specification of the predictive regression from panel A in the abovementioned tables, which includes the lagged variable, is also investigated as investment strategy. The regression from the previous in-sample analysis showed that including the lagged variable increases performance significantly in terms of adjusted R^2. Hence, it is interesting so see whether the lagged specification of the regression also performs better out-of-sample, that is, in forecasting the variance-covariance matrix and ultimately delivering a higher risk-adjusted portfolio return.

Interestingly, the MinVarHtLag strategy seem to perform worse than its no-lag counterpart. The standard deviation of portfolio returns is slightly higher than the MinVarHt and MinVarMA benchmark portfolio. This is also reflected in the Sharpe ratio, where the lagged specification only delivers a Sharpe ratio of 0.95. It is still highly significant compared to the EW strategy but there is no support against the null hypothesis compared to the MinVarMA portfolio. The certainty equivalent of 3.22% is in line with the MinVarMA strategy. The difference compared to the equally-weighted strategy is not significant, as it gives a p-value of 50.2%.

The results do neither show increased portfolio performance nor economic benefits of including the lagged variable to determine the optimal stock-bond portfolio compared to the MinVarHt strategy. This is a bit surprising given the demonstrated in-sample predictive power of the lagged variable. Also, the results from table 8 show that MinVarHtLag performs in line with, or even slightly worse than, the simple moving-average benchmark strategy. However, it manages to outperform the naive 1/N portfolio, and statistically so, in terms of portfolio volatility and Sharpe ratio.

5.3.5 Portfolios based on analyst forecast data only

The next group of strategies are all based on estimating the variance-covariance matrix by using forward-looking analyst forecast data on macroeconomic variables. Differently from the historical data-based strategies, portfolios formed on analyst forecasts should intuitively perform well out-of-sample. The reason is that the variables ought to have good predictive power, since analyst forecasts on a given variable is based on both the current and forecasted macroeconomic regime. Hence, embedded within these variables is implied information about the future state of the macroeconomic environment.

The out-of-sample prediction performance of analyst forecasts is evaluated on six different strategies. Two of them are based on the predictive regressions from panel C in the in-sample-analysis, where both mean consensus and dispersion measures are included in the regression on volatilities and correlation. The remaining four strategies have not been tested for their in-sample prediction performance, and therefore require a short introduction.

The first strategy, MinVarFt, is a minimum variance portfolio constructed based on a variance-covariance matrix estimated by a rolling-window regression that only includes the vector of mean consensus forecast variables $\boldsymbol{F}_t$. Following the method from the in-sample analysis, specifications both including and excluding the lagged variable are presented. Hence, the stock and bond volatilities and correlation are regressed on a constant, its lag, and the collection of standardized measure of the mean consensus in analysts' forecasts as follows in equation 65.

Equation 65

$$\rho_{t+1} \; or \; \sigma_{t+1} = \hat{\alpha} + \hat{\beta} lag + \hat{\varphi} \boldsymbol{F}_t + \epsilon_{t+1}$$

The MinVarFt strategy uses equation 65 to estimate the variance-covariance matrix, however, not including the lag. The strategy displays good performance in table 8 with a Sharpe ratio above one. This is higher than any of the other strategies analyzed so far. The difference in performance is also highly significant when compared to the naive EW strategy, but not statistically different from the better performing benchmark strategy MinVarMA. The probability, given the null hypothesis is true, of obtaining a test-statistic at least of the same magnitude as the observed value is 42.26%, that is, not enough evidence to reject the null hypothesis in equation 63.

The strategy has a standard deviation in line with the other strategies, except EW, and the certainty equivalent is 18 basis points higher than the MinVarMA benchmark, although the difference is not significant with a p-value of 14.47%. Even though the first three performance metrics show better but insignificant results, the strategy seems to have economically significant outperformance given the opportunity cost with respect to EW and MinVarMA is 29 and 16 basis points. Therefore, it seems like analyst forecasts do provide out-of-sample benefits in allocating capital between stocks and bonds.

Next, the strategy MinVarFtLag, which includes the lagged variable, is tested. This corresponds to the entire regression line in equation 65. Given the high degree of explanatory power of the lagged variable presented in the in-sample analysis, it is surprising that including the lag in predicting covariance, results in worse portfolio performance. Again, it supports the claim made earlier, that forward-looking analyst forecasts, do seem to have out-of-sample-predictability. The MinVarFtLag portfolio delivered a higher standard deviation of returns compared to the MinVarMA benchmark and the similar no-lag specification MinVarFt. Also, its Sharpe ratio was lower at 0.94, worse than the aforementioned portfolio.

Both strategies using mean consensus variables only, have a portfolio turnover around 5-6%, which is close to average across all portfolios considered, but higher than the benchmarks. MinVarFtLag seems only to outperform the equally-weighted portfolio as neither the Sharpe ratio nor certainty equivalent show any evidence for better performance compared to the simple moving-average benchmark strategy.

Similar to including mean consensus variables only, the same analysis is conducted with dispersion measures of analyst forecasts only, $\boldsymbol{D}_t$. Again, the same method is used, where both a lagged and non-lag specification is tested. The regression line that is used to estimate the variance-covariance matrix is.

Equation 66

$$\rho_{t+1} \text{ } or \text{ } \sigma_{t+1} = \hat{\alpha} + \hat{\beta} lag + \hat{\varphi} \boldsymbol{D}_t + \epsilon_{t+1}$$

The first specification without the lag, MinVarDt, performs considerably better than the EW benchmark strategy in terms of portfolio standard deviation and Sharpe ratio, where the difference in both measures is statistically significant at the one percent level. However, the strategy does not perform better than the simple moving-average benchmark.

The strategy displays a Shape ratio of 0.95 over the period and a certainty equivalent return of 3.19%. In fact, only using the collection of dispersion measures to forecast variance-covariance, underperforms in terms of all performance metrics compared to MinVarMA. However, the results are not economically significant, which can be deduced from the low opportunity cost to MinVarMA of -0.03%.

Comparing the two similar strategies of using consensus mean or dispersion measures, it is found that the mean consensus variables of the analyst forecast in isolation have better out-of-sample performance compared to the dispersion variables.

When the lagged variable is added to the regression, as in equation 66, the performance declines even further. It is the same observation as for the strategies that only consider historical or consensus mean variables. The Sharpe ratio of MinVarDtLag declines from 0.95 to 0.90, which in economic terms seems like a large decrease in performance. The strategy is still outperforming the equally-weighted portfolio and significantly so with a p-value of 1.57% for the Sharpe ratio. However, when comparing to the other benchmark portfolio, MinVarDtLag does not seem as an attractive strategy given its underperformance across all performance metrics. The opportunity cost to the MinVarMA strategy is -0.24%. This magnitude must be considered economically significant, as an investor, on average loses 0.24% per month, in annualized terms, when following the MinVarDtLag strategy compared to the MinVarMA benchmark portfolio.

Comparing the performance across all strategies considered, this portfolio yields the third lowest certainty equivalent return and the second lowest Sharpe ratio. Hence, it seems that dispersion of analyst forecasts in isolation have bad out-of-sample forecast performance and investors will be better off following the simpler moving-average benchmark strategy.

The last two analyst forecast-based strategies correspond to the two specifications of the predictive regression from panel C in the in-sample analysis. Here the collection of mean consensus variables and the collection of dispersion measures are both used in the rolling-window regression to predict the variance-covariance matrix and in turn to determine the portfolio weights. The in-sample analysis showed great performance of this model with adjusted R^2 above 56% for the regressions of stock volatility and correlation, and above 50% for the bond volatility regression. Hence, it is interesting to see if the combination of the two analyst forecast variable vectors $\boldsymbol{F}_t$ and $\boldsymbol{D}_t$ also performs well out-of-sample.

The first of the two strategies is the no-lag specification called MinVarFtDt. Surprisingly, the strategy does not perform well. Alongside its lagged version MinVarFtDtLag, it has the highest standard deviation of portfolio returns except from the equally-weighted strategy. The returns vary 0.22 percentage points more per month, in annualized terms, compared to MinVarMA. This may not seem like a lot, but it makes up a large proportion of the total variation in returns of 3.54%.

The Sharpe ratio of the portfolio is 0.95, which is in line with all the other strategies, however still lower than the moving-average benchmark. Like all the other strategies considered, it outperforms the EW strategy and the difference in performance is statistically significant as the p-value of 0.59% is below the five percent level. The certainty equivalent rate of 3.31% is slightly better than both benchmarks, but the difference is not significant. Even though the lower Sharpe ratio suggests inferior performance of MinVarFtDt, the opportunity costs with respect to the benchmarks are both positive, suggesting that the investor demands a higher return from the benchmark portfolios in order to be indifferent in utility terms.

In the analysis of the preceding strategies, introducing the lagged variable to the rolling-window regression of volatilities and correlation decreased the out-of-sample performance. However, this time adding the lag results in better portfolio performance as the Sharpe ratio increases to 0.98, almost equal to the MinVarMA strategy. This finding is in line with the results of the in-sample analysis, where the lagged variable proved to be the single most significant predictor of stock-bond correlation and volatilities.

The standard deviation of returns remains high with 3.54% per month in annualized terms. Like the Sharpe ratio, the certainty equivalent rate increases to 3.42%. The difference in performance of the CEQ measure between MinVarFtDtLag and the benchmarks is not significant as the p-values are 36.48% and 22.60%. Hence, not enough evidence against the null hypothesis. The opportunity cost with respect to the benchmark strategies are also high, with 0.33% and 0.20%, respectively. This suggests that the performance of MinVarFtDtLag may not be statistically significant but economically significant. Finally, implementing the two strategies, using both mean consensus and dispersion variables, are expensive in terms of transaction costs, as on average 8.43% and 9.80% of wealth is traded at each rebalancing date.

5.3.6 Portfolios based on historical and analyst forecast data

Now that strategies based solely on historical or forecast data have been analyzed, next up is the combination where both types of data is used. These strategies correspond to the predictive regression from panel B in the in-sample analysis. The in-sample analysis showed that these specifications has the best performance in terms of highest adjusted R^2 in predicting both volatilities and correlation of stock-bond returns. Hence, it is interesting to see whether these two strategies also deliver the best out-of-sample prediction performance. The results are presented in table 9.

Table 9 – This table reports the performance of the last 5 strategies considered. Portfolio standard deviation, Sharpe ratio, and certainty equivalent measures are annualized, whereas portfolio turnover shows the average monthly percentage of wealth traded at the rebalancing date, and opportunity cost displays the annualized average monthly opportunity cost. A risk aversion of one is employed for the CEQ and OC measure. The numbers in parentheses are p-values for the one-sided test of difference in performance. The upper p-value compares EW and strategy k and the lower compares MinVarMA and strategy k.

Type	Strategy k	σ^k	SR^k	CEQ^k	PT^k	OC^k_{EW}	$OC^k_{MinVarMA}$
Benchmark	*EW*	7.71	0.41	2.85	2.06		-0.13
	MinVarMA	3.32	0.99	3.22	1.13	0.13	
		(0.0000)	(0.0159)	(0.4194)			
Historical & Analyst forecast data	*MinVarHtFtDt*	3.50	1.04	3.57	9.71	0.48	0.35
		(0.0000)	*(0.0015)*	*(0.3260)*			
		(0.9877)	*(0.2838)*	*(0.1277)*			
	MinVarHtFtDtLag	3.50	1.04	3.59	10.29	0.50	0.37
		(0.0000)	*(0.0017)*	*(0.3237)*			
		(0.9910)	*(0.2391)*	*(0.0893)*			
Other strategies	*MinVarPrevVolCorrEs*	3.32	0.95	3.09	6.53	-0.01	-0.13
		(0.0000)	*(0.0192)*	*(0.4479)*			
		(0.5477)	*(0.8181)*	*(0.8539)*			
	MinVarPrevVolCorr	3.54	0.99	3.45	13.15	0.37	0.24
		(0.0000)	*(0.0022)*	*(0.3580)*			
		(0.9871)	*(0.4843)*	*(0.2456)*			
	MinVarZeroCorr	3.26	0.90	2.88	0.97	-0.22	-0.35
		(0.0000)	*(0.0517)*	*(0.4970)*			
		(0.2140)	*(0.9283)*	*(0.9853)*			

Source: Own contribution

The first strategy that includes both historical and forecast data is the no-lag specification abbreviated MinVarHtFtDt. At first sight the performance does not look promising as the monthly standard deviation in annualized terms is 0.18 basis points higher than the best performing benchmark strategy MinVarMA with a level of 3.50%. However, the strategy compensates by delivering a higher average annualized return of 5.30% which in turn yields a Sharpe ratio of 1.04. This is 0.05 higher than the moving-average benchmark strategy. The outperformance is not statistically significant given a p-value of 28.38%.

The strategy also performs well according to the certainty equivalent return measure. An investor, with a risk aversion constant of one, will demand a risk-free return of 3.57% in order to abandon the MinVarHtFtDt strategy. Even though the difference in CEQ between the two benchmarks are 0.72% and 0.35% respectively, the difference remains insignificant in statistical terms.

In economic terms the strategy seems to provide significant outperformance. The opportunity cost is positive with respect to both benchmark portfolios and a magnitude of 0.48% and 0.35%. Interesting to note is the high portfolio turnover. Hence, it seems like implementing MinVarHtFtDt provides a better return, but it is costly to do so. On average 9.99% of wealth is traded at the rebalancing date, which will mean incurring high transaction costs.

In the in-sample analysis, the abovemention specification increased performance when the lag was included. Therefore, the MinVarHtFtDtLag strategy is tested for out-of-sample performance.

The inclusion of the lag does not change the standard deviation of portfolio returns, as the measures remains at 3.50%. The Sharpe ratio increases slightly with 0.007, which cannot be seen in the table. Hence, this strategy presents the highest Sharpe ratio, that is, the highest risk-adjusted return of all the strategies considered. This is quite an interesting finding but again a statistical test of difference in Sharpe ratios do not present enough evidence against the null hypothesis in equation 63, as the chance of observing a similar or smaller test statistic is only 23.91%.

The strategy also presents the highest certainty equivalent return of all portfolios of 3.59%. A p-value of 8.93% indicates that the difference from the MinVarMA CEQ is not statistically significant at the five percent level.

Lastly, the strategy also scores high on the opportunity cost measure, showing the highest relative OC in relation to the EW and MinVarMA strategy with opportunity costs of 0.50% and 0.37%. This suggests economic significance in difference of performance but again implementing this trading strategy would be very expensive with an average monthly portfolio turnover of 10.16%.

Noteworthy is the that MinVarHtFtDtLag delivers the best performance both in-sample as well as out-of-sample.

5.3.7 Other strategies considered

The in-sample analysis found that the lag explains a lot of the variation in volatilities and correlation of stock and bond returns, respectively. Based on this observation two alternative strategies have been formularized, where the first is based on a regression on the lag and the second is based on the previous month's realized variance-covariance matrix.

The first alternative strategy MinVarPrevVolCorrEst is a portfolio that is formed on a variance-covariance matrix estimated from a linear predictive regression on the past realized values of volatility and correlation over the same estimation window as the other models. The following regressions are run.

Equation 67

$$\rho_{t+1} = \hat{\alpha} + \hat{\beta}\rho_t + \epsilon_{t+1}$$

Equation 68

$$\sigma_{t+1} = \hat{\alpha} + \hat{\beta}\sigma_t + \epsilon_{t+1}$$

The results of these strategies are also presented in table 9. The variation of portfolio returns is significantly less than the EW portfolio and in line with moving-average benchmark strategy with 3.32%. That said, the Sharpe ratio and certainty equivalent return are both below the best benchmark with Sharpe ratio of 0.95 and CEQ of 3.09%. The relatively bad performance is also illustrated by the negative opportunity costs with respect to the two benchmarks.

Interestingly, the lagged variable's high predictive power of volatilities and correlation of stock-bond returns from the in-sample analysis cannot be translated into good out-of-sample predictive power.

The second strategy, based on the insights gained from the in-sample analysis, is the simple strategy of using last month's realized variance-covariance matrix as an estimate for next period. This portfolio performs relatively well, although it has one of the highest standard deviation of returns. The strategy provides a Sharpe ratio in line with the MinVarMA portfolio and a higher certainty equivalent return than the benchmarks, although the difference is not statistically significant.

In contrast to the strategy based on the regression on the lagged variable, MinVarPrevVolCorr seem to provide results that are economically significant compared to the two benchmark portfolios with positive opportunity costs of 0.37% and 0.24%, respectively. However, the strategy also has the highest portfolio turnover, with the average amount of wealth traded around the rebalancing date is 13.15%. Hence, significant transaction costs are assumed to be incurred when implementing this strategy.

The last alternative strategy is MinVarZeroCorr, where correlation is set to zero and volatilities are estimated using a moving-average. The strategy has been proposed by DeMiguel *et al.* (2013) to improve the out-of-sample performance. Interestingly, with a volatility of 3.26%, it gives the lowest standard deviation of all portfolios considered. Even though it is low, it is only statistically significantly different from the equally-weighted portfolio.

The rest of the performance measures display inferior performance with a Sharpe ratio of only 0.90. As the only strategy, the difference in Sharpe ratios to both benchmark strategies are insignificant at the five percent level. Also, it provides a certainty equivalent almost as low as the EW portfolio and the opportunity costs to both benchmarks are negative. Surprisingly, this strategy has the lowest portfolio turnover among all portfolio of only 0.23% on average per month.

5.3.8 Overall out-of-sample predictability performance

All strategies statistically significantly outperformed the equally-weighted benchmark portfolio in terms of lower portfolio volatility and higher Sharpe ratio except for the MinVarZeroCorr strategy. Previous literature has found that the EW asset allocation strategy performs well compared to more complex strategies (DeMiguel, Garlappi and Uppal, 2009). Hence, it is interesting to find that all portfolios build on predictive regressions using historical and forecast data are able to provide significantly better performance out-of-sample.

Another interesting finding is that the only portfolio with an insignificant Sharpe ratio compared to both benchmarks, the MinVarZeroCorr strategy, at the same time provides the lowest standard deviation of portfolio returns of all strategies. Especially, considering that this portfolio is based on a rather arbitrary approach, where only volatilities are estimated, and correlation is simply set to zero. The seeming ability to reduce portfolio risk is on the other hand punished by significantly lower portfolio returns. The MinVarZeroCorr delivers the lowest annualized monthly return of all strategies considered.

When performance is compared to the best performing benchmark portfolio, the simple 60-months' moving-average strategy, the results are not as convincing. None of the 13 strategies managed to provide better out-of-sample performance than the MinVarMA. The difference in portfolio volatility, Sharpe ratio, and certainty equivalent is statistically insignificant, suggesting that the investor might be as well off using a simple historical average for the prediction of volatilities and correlation between stock and bond returns in determining portfolio weights.

Although there is no statistical difference in performance, the portfolio turnover and opportunity cost measures are used to uncover any potential economic significance in performance. The opportunity cost relating to the MinVarMA strategy has shown values of significant magnitude. The two strategies using both historical and analyst forecast data, MinVarHtFtDt and MinVarHtFtDtLag, provided an annualized monthly opportunity cost of 0.35% and 0.37%, respectively, which must be considered economically significant. This average monthly outperformance would amount to an outperformance in total return over the entire sample period July 1999 to December 2018 of 17.85% and 19.16%, although not reported in the table.

The opportunity cost metric is a measure before transaction costs. The portfolio turnover metric tries to cover the aspect of implementing each strategy, as it reports the average percentage of wealth traded at the rebalancing dates. The two tables presenting the results show that the strategies providing the best performance also has the highest portfolio turnover, that is, they are the most expensive strategies to implement in terms of transaction costs. The two strategies using historical and forecast data mentioned above both trade close to ten percent of wealth every month, in order to rebalance the portfolio to the desired weights. This is a considerable amount that needs to be reinvested and it is close to nine percentage points higher than the MinVarMA benchmark portfolio.

Even though the MinVarHtFtDt and MinVarHtFtDtLag portfolios seem to have significant economic outperformance, before transaction costs, the high portfolio turnover is an important issue to consider when choosing to implement these strategies or not. The economic outperformance of these strategies might be offset by the large amount of transaction costs incurred every month, thus making them undesirable. The exact effect is difficult to estimate, and the issue will not be modeled or accounted for since it is considered beyond the scope of this thesis.

Another interesting result from the out-of-sample analysis is the performance of the lag. As mentioned before, the lag seems to have a large explanatory power in predicting volatilities and correlation between stocks and bonds, when the predictive regression is fitted on all available data, that is, in-sample.

However, several strategies presented better out-of-sample portfolio performance when the lag was excluded from the predictive regression. This can be seen from the decrease in Sharpe ratios in some of the analyst forecast only strategies, comparing the lag to the non-lag specifications. This could imply that analyst forecasts, in terms of mean consensus and dispersion measures, do have out-of-sample predictability although the effect is not statistically significant.

The analysis presented above suggests that when the equally-weighted portfolio is considered as the benchmark, almost all of the strategies outperform, that is, present enough evidence against the null hypothesis of strategy k performing worse than the benchmark. However, the statistical significance cannot be established when the moving-average benchmark is used, hence we fail to reject the null hypothesis. That being said, some strategies seem to have economically significant performance, but once controlling for transaction costs the significance becomes uncertain.

5.4 Robustness check

5.4.1 OLS assumptions

With the extensive use of Ordinary Least Squares throughout the paper, tests for the most relevant assumptions, described in section 4.3.1 "OLS assumptions", are required. We have decided to focus on autocorrelation and heteroscedasticity as these characteristics are often present in financial data.

5.4.2 Test for autocorrelation

In order to identify potential heteroscedasticity in our data, we have conducted the Durbin-Watson d-test, described in section 4.3.2. The test is carried out on our in-sample data for stock volatility, bond volatility, and stock-bond correlation. Each regression includes the following variables, VIX, all five mean consensus measures, and all five dispersion measures, corresponding to panel C regressions in the in-sample analysis. The Durbin-Watson d-test is not conducted on the regression formulation that includes the lagged variables. The reason being that this particular specification provides a violation of the requirements for the test (Gujarati and Porter, 2008).

We test our data for autocorrelation by comparing the d-statistic, obtained via equation 39, with some specified thresholds. The values for these thresholds are determined by the number of observations in the sample, the number of explanatory variables, and the significance level. With 294 observations, 11 explanatory variables, and a significance level of five percent, the boundaries $d_L = 1.7245$ and $d_U = 1.8941$ are applicable. Table 10 below describes the results for the Durbin-Watson d-test.

Positive first-order serial correlation is particularly present in economic data (Studenmund, 2014). It is therefore not surprising that our sample is serially correlated for both stock volatility, bond volatility, and stock-bond correlation. The first-order correlation seems to be strongest for stock-bond correlation with a DW d-stat of 1.2111, while it is weaker for stock-volatility and bond-volatility with DW d-stats of 1.4961 and 1.5049, respectively.

The finding shows that the previous month's data do have explanatory power in relation to next month's data. The tendency is commonly known as momentum and is extensively covered in the literature (See Carhart, 1997; Jegadeesh & Titman, 1993). Even though it is not surprising that our data is characterized by autocorrelation, it is nevertheless a violation of the classical Gauss-Markow assumptions. Importantly, autocorrelation does not cause any bias in the obtained estimators. It does, however, tend to bias the standard errors of the estimators downwards. As a consequence, there is a risk of obtaining unjustifiably large t-stats, which increases the chance of making Type I errors.

Table 10 - The Durbin-Watson d test for the data in the in-sample analysis. Included in the table are results for data on stock volatility, bond volatility, and stock-bond correlation. The table reports DW d-stats for individual sets of data with significance level of 5%, and result of the test in brackets. "Positive" describes a positive serial correlation and is indicated when: DW t-stat <1.7245, "Negative" describes negative serial correlation and is indicated when: DW t-stat > 1.8941. Finally, "Inconclusive" is indicated, when autocorrelation cannot be determined from the test and is indicated when: 1.7245 < DW < 1.8941.

	Stock volatility	Bond volatility	Stock-bond correlation
DW d-stat	1.4961	1.5049	1.2111
	(Positive)	(Positive)	(Positive)

Source: Own contribution

5.4.3 Test for homoscedasticity

As noted in section 4.3.3 "White Test", there is substantial evidence showing that financial data is characterized by heteroscedasticity. In order to test for this characteristic, we have carried out the earlier described White test.

We follow the procedure, described in section 4.3.3 on stock volatility, bond volatility, and stock-bond correlation, both with and without a lagged variable. The test is carried out on the sample including mean consensus forecasts and dispersion measures, as well as VIX. This gives 12 explanatory variables when the lagged variable is included and 11 when it is excluded.

Using equation 41 and replacing K with 12 and 11, we get df = 90 when the lagged variable is included and df = 77 when it is excluded. With significance level of five percent, the critical chi-square values are $\chi^2_{90} = 113.15$ and $\chi^2_{77} = 98.48$. The results from the test is presented in table 11 below. For all six auxiliary regressions, the table presents R^2 values and $R^2 * n$, which is the input for the test. Lastly, the table presents the test results for the White test in brackets, where negative indicates that we reject the hypothesis of homoscedasticity, that is, heteroscedasticity is present in the data.

As it appears from table 11, there seems to be heteroscedasticity across all six data samples. The result is not surprising, as previous research similarly documents the presence of heteroscedasticity in financial data.

Table 11 - Results of the White Test for homoscedasticity

Specification		Stock volatility	Bond volatility	Stock-bond correlation
With lag				
	R^2	0.72	0.64	0.40
	$R^2 * n$	210.69	186.66	116.44
	Result	(Negative)	(Negative)	(Negative)
Without lag				
	R^2	0.59	0.46	0.35
	$R^2 * n$	172.74	133.94	104.07
	Result	(Negative)	(Negative)	(Negative)

Source: Own contribution

5.4.4 Comments and consequences of the OLS assumption tests

The robustness tests presented in the previous sections, provide evidence that our sample is characterized by both heteroscedasticity and autocorrelation. This was expected, but it does nevertheless influence our findings. This section is included to discuss the related implications.

Table 10 presents the result of our test for autocorrelation. Although our findings present a violation of the Gauss-Markov assumption, it does not cause any bias of the obtained coefficients. It does however, have a tendency to underestimate the standard errors and correspondingly, overestimate the t-statistics of the coefficient estimates (Studenmund, 2014). The identification of autocorrelation in the in-sample data means, that we have to be cautious about our conclusiveness in relation to identification of statistically significant macroeconomic variables, as presented in table 4, 5, and 6. For all strategies, we have proposed a variation that includes the lagged variable, which is a recommended remedy to eliminate the presence of autocorrelation (Studenmund, 2014). Though we cannot test for autocorrelation after including the lagged variable, we must assume that it reduces or even eliminates the problem.

In section 5.4.3 "Test for homoscedasticity", we find that our in-sample data is characterized by heteroscedasticity. As with the case of autocorrelation, the presence of heteroscedasticity does not cause any bias in the coefficient estimators, but only in the standard errors of the coefficient estimates. Conversely, heteroscedasticity has the effect of inflating the standard error and thereby reducing the t-statistics of the coefficient estimates. This increases the risk of type II errors in our in-sample analysis.

In conclusion, the presence of heteroscedasticity and autocorrelation means that our findings in the in-sample analysis bear the risk of both type I and type II errors, and hence should not be considered definitive. However, the primary purpose of this research is to assess the performance of the proposed strategies, and to this end, the presence of autocorrelation and heteroscedasticity does not influence our findings. The reason being that we do not make any conclusions in this regard, based on the t-statistics of the coefficient estimates.

5.4.5 Different levels of risk aversion

As presented in equation 19 and equation 26, the measure of certainty equivalent is defined using the utility function of the investor. In economic theory, the utility function includes a measure for risk aversion, and as we have decided to exclude the approximation of individual investors level of risk aversion, we have followed previous examples and assumed a given level of risk aversion (DeMiguel, Garlappi and Uppal, 2009; Jivraj, 2012b).

To examine the sensitivity of our assumed level of risk aversion to the certainty equivalent measure, we have calculated the measure with varying levels of risk aversion. Specifically, in addition to the absolute risk aversion (ARA) level of one, which we have used throughout the paper, we have calculated certainty equivalent using equation 47 with following levels of risk aversion 0.25, 0.5, 2, and 5.

As described in section 3.6.2 "Certainty equivalent", the measure of certainty is defined as the risk-free return that an investor will require in order to forgo a risky investment with higher return. Table 12 below illustrates certainty equivalent when considering alternative levels of risk aversion for all strategies. To ease the interpretation of the table, a column for excess returns is included. Additionally, the risk-premium is presented in parenthesis below each certainty equivalent measure. The risk-premium is simply the difference between the excess return of the individual strategy and the certainty equivalent.

To illustrate the findings of our sensitivity analysis of CEQ, consider the EW strategy in table 12 below. The realized excess return of the strategy is 3.15%. The column with certainty equivalent of 0.25 represents the least risk averse investor. That is, the column represents investors that primarily are focused on the return of the strategy and cares relatively little about the associated risk.

This is apparent from the corresponding certainty equivalent of 3.08%, which is only 0.07% lower than the realized excess return of the strategy. To clarify, the interpretation is, that the investor will be indifferent between obtaining 3.08% return with certainty or 3.15% with the risky strategy. In this example, the investor requires a mere 0.07% risk premium in order to accept the associated risk of the strategy, which in this case is 7.71%. If we instead consider the most risk averse investor for the same strategy, she will accept 1.66% risk-free return in order to avoid the risky strategy. The corresponding risk premium is 1.49%, which is substantially higher than the previous 0.07% for the investor with low risk aversion. The comparison of the two types of investors show that high (low) risk aversion is followed by a low (high) certainty equivalent. The finding is expected from equation 47 and it is consistent across all strategies.

Table 12 - Certainty equivalent in percent with different levels of risk aversion. Risk premium in parenthesis is also shown in percent.

				Level of Absolute Risk Aversion				
Name	Strategy	Excess rtn.	σ^k	0.25	0.5	1.0	2.0	5.0
Benchmark	EW	3.15	7.71	3.08	3.00	2.85	2.56	1.66
				(-0.07)	(-0.15)	(-0.3)	(-0.6)	(-1.49)
	MinVarMA	3.28	3.32	3.27	3.25	3.22	3.17	3.00
				(0.01)	(0.03)	(0.06)	(0.11)	(0.28)
With lag	MinVarHtLag	3.20	3.48	3.19	3.17	3.14	3.08	2.90
				(0.01)	(0.03)	(0.06)	(0.12)	(0.30)
	MinVarFtDtLag	3.48	3.54	3.46	3.45	3.42	3.36	3.17
				(0.02)	(0.03)	(0.06)	(0.13)	(0.31)
	MinVarHtFtDtLag	3.54	3.53	3.53	3.51	3.48	3.42	3.23
				(0.02)	(0.03)	(0.06)	(0.12)	(0.31)
	MinVarFtLag	3.28	3.48	3.27	3.25	3.22	3.16	2.98
				(0.02)	(0.03)	(0.06)	(0.12)	(0.30)
	MinVarDtLag	3.04	3.39	3.03	3.02	2.99	2.93	2.76
				(0.01)	(0.03)	(0.06)	(0.11)	(0.29)
Without lag	MinVarHt	3.41	3.43	3.40	3.38	3.34	3.28	3.06
				(0.02)	(0.03)	(0.07)	(0.14)	(0.34)
	MinVarFtDt	3.37	3.54	3.36	3.34	3.31	3.25	3.06
				(0.02)	(0.03)	(0.06)	(0.13)	(0.31)
	MinVarHtFtDt	3.54	3.55	3.53	3.51	3.48	3.42	3.23
				(0.02)	(0.03)	(0.06)	(0.12)	(0.32)
	MinVarFt	3.44	3.46	3.43	3.41	3.38	3.32	3.14
				(0.01)	(0.03)	(0.06)	(0.12)	(0.30)
	MinVarDt	3.25	3.41	3.23	3.22	3.19	3.13	2.96
				(0.01)	(0.03)	(0.06)	(0.12)	(0.29)
Other strategies	MinVarPrevVolCorrEst	3.15	3.32	3.13	3.12	3.09	3.04	2.87
				(0.01)	(0.03)	(0.06)	(0.11)	(0.28)
	MinVarPrevVolCorr	3.52	3.54	3.50	3.49	3.45	3.39	3.20
				(0.02)	(0.03)	(0.06)	(0.13)	(0.31)
	MinVarZeroCorr	2.93	3.26	2.92	2.90	2.88	2.82	2.66
				(0.01)	(0.03)	(0.05)	(0.11)	(0.27)

Source: Own contribution

For all levels of risk aversion, the EW strategy provides the lowest levels of certainty equivalent and the highest levels of risk premia. The primary driver is the substantially higher standard deviation for this strategy.

As a consequence, investors will require a relatively low secured return to forgo the high risk of the strategy. The highest levels of certainty equivalent are present for the strategy MinVarHtFtDt. The high certainty equivalent for this strategy is primarily driven by the excess return of 3.54%, which is higher than any other strategy. The strategy MinVarZeroCorr has the lowest levels of risk premia as a consequence of the low standard deviation of 3.25%.

To test the sensitivity of varying levels of risk aversion, in relation to statistical significance of different strategies' certainty equivalent, p-values are presented in table 13 below. To obtain the p-values for this table we make use of the stationary bootstrapping as described in section 4.5.

When we consider p-values for the comparison with the EW benchmark, we observe that all strategies show increasing levels of significance of CEQ with increasing levels of ARA. The finding is not surprising considering the significantly higher standard deviation of EW, which means that the strategy is the most sensitive to ARA.

Interestingly, the opposite is the case for almost all strategies when comparing to the MinVarMA strategy. Only MinVarPrevVolCorrEst and MinVarZeroCorr indicates falling p-values with increasing risk aversion. This tendency is a direct result of the comparison of risk from the different strategies. As only MinVarPrevVolCorrEst and MinVarZeroCorr have smaller standard deviations than MinVarMA, these two strategies are the only ones that shows a decrease in p-values as a result of an increase in risk aversion.

In terms of p-value significance, our sensitivity analysis has a single notable implication, namely that the CEQ for the strategy MinVarHtFtDt compared to MinVarMA becomes insignificant at a 10% significance level for a risk aversion of 5. We observe no other implications in terms of p-values, and hence we conclude that our analysis of CEQ is robust to varying levels of risk aversion.

Table 13 – Sensitivity analysis of the certainty equivalent measure. CEQ is presented in in percent and p-values in parenthesis. The upper p-value compares EW and strategy k and the lower compares MinVarMA and strategy k.

		Level of Absolute Risk Aversion				
Type	Strategy	0.25	0.5	1.0	2.0	5.0
Benchmark	*EW*	3.08	3.00	2.85	2.56	1.66
	MinVarMA	3.27	3.25	3.22	3.17	3.00
		(0.4584)	*(0.4452)*	*(0.4194)*	(0.3709)	(0.2513)
Historical data only	*MinVarHt*	3.40	3.38	3.34	3.28	3.07
		(0.4356)	*(0.4216)*	*(0.3945)*	*(0.3440)*	*(0.2236)*
		(0.3425)	*(0.3443)*	*(0.3480)*	*(0.3554)*	*(0.3780)*
	MinVarHtLag	3.19	3.17	3.14	3.08	2.90
		(0.4539)	*(0.4397)*	*(0.4122)*	*(0.3608)*	*(0.2370)*
		(0.4952)	*(0.4975)*	*(0.5020)*	*(0.5111)*	*(0.5383)*
Analyst forecast data only	*MinVarFt*	3.43	3.41	3.38	3.32	3.14
		(0.4177)	*(0.4045)*	*(0.3789)*	*(0.3312)*	*(0.2174)*
		(0.1382)	*(0.1403)*	*(0.1447)*	*(0.1538)*	*(0.1827)*
	MinVarFtLag	3.27	3.25	3.22	3.16	2.98
		(0.4560)	*(0.4423)*	*(0.4157)*	*(0.3656)*	*(0.2435)*
		(0.5027)	*(0.5063)*	*(0.5136)*	*(0.5281)*	*(0.5702)*
	MinVarDt	3.23	3.22	3.19	3.13	2.96
		(0.4628)	*(0.4483)*	*(0.4200)*	*(0.3671)*	*(0.2399)*
		(0.5601)	*(0.5614)*	*(0.5641)*	*(0.5695)*	*(0.5857)*
	MinVarDtLag	3.03	3.02	2.99	2.93	2.76
		(0.5132)	*(0.4986)*	*(0.4700)*	*(0.4156)*	*(0.2805)*
		(0.8839)	*(0.8846)*	*(0.8860)*	*(0.8887)*	*(0.8966)*
	MinVarFtDt	3.36	3.34	3.31	3.25	3.06
		(0.4314)	*(0.4173)*	*(0.3901)*	*(0.3397)*	*(0.2206)*
		(0.3731)	*(0.3756)*	*(0.3807)*	*(0.3910)*	*(0.4233)*
	MinVarFtDtLag	3.46	3.45	3.42	3.36	3.17
		(0.4045)	*(0.3909)*	*(0.3648)*	*(0.3165)*	*(0.2041)*
		(0.2202)	*(0.2221)*	*(0.2260)*	*(0.2339)*	*(0.2593)*
Historical & Analyst forecast data	*MinVarHtFtDt*	3.53	3.51	3.48	3.42	3.23
		(0.3647)	*(0.3514)*	*(0.3260)*	(0.2798)	(0.1753)
		(0.1253)	*(0.1261)*	*(0.1277)*	(0.1310)	(0.1417)
	MinVarHtFtDtLag	3.53	3.51	3.48	3.42	3.23
		(0.3616)	*(0.3486)*	*(0.3237)*	(0.2785)	(0.1757)
		(0.0872)	*(0.0879)*	*(0.0893)*	(0.0923)	(0.1018)
Other strategies	*MinVarPrevVolCorrEst*	3.13	3.12	3.09	3.04	2.87
		(0.4885)	*(0.4747)*	*(0.4479)*	(0.3970)	(0.2709)
		(0.8549)	*(0.8546)*	*(0.8539)*	(0.8525)	(0.8480)
	MinVarPrevVolCorr	3.50	3.49	3.45	3.39	3.20
		(0.3968)	*(0.3835)*	*(0.3580)*	(0.3110)	(0.2014)
		(0.2408)	*(0.2424)*	*(0.2456)*	(0.2521)	(0.2724)
	MinVarZeroCorr	2.92	2.90	2.88	2.82	2.66
		(0.5358)	*(0.5227)*	*(0.4970)*	(0.4474)	(0.3189)
		(0.9861)	*(0.9859)*	*(0.9853)*	(0.9842)	(0.9803)

Source: Own contribution

5.4.6 Different length of estimation window

As another robustness check, we have conducted a test in order to investigate the "memory" of financial data. In contrast to our primary analysis, where we apply a 60 months' estimation window, in this robustness check we use 120 months of data. The test will give an indication of whether there is any predictive information in data that is more than 5 years old.

The performance measures for the 120 months' robustness check are calculated in an identical manner as the procedure for the primary analysis, however, with the estimation window being twice the length. An effect from the longer estimation window is that we need 60 additional months to estimate next month's variance-covariance matrix. Consequently, the analyzed period starts in July 2004, where the analyzed period for the primary analysis starts in July 1999. The results from the analysis are presented in table 14 below. In order to compare the results with the primary analysis with a 60 months' estimation window, the latter is presented below each measure in parenthesis.

As the equally-weighted portfolio does not make use of an estimation window, changes for this strategy will not be a result of the change in estimation window but merely an effect of the different period analyzed. For this reason, the changes for this strategy will not be commented upon, and we therefore focus on the remaining 14 strategies.

When comparing the average return and excess return of the two estimation lengths, there is a consistent trend of the short estimation window (60 months) outperforming the long estimation window (120 months). Not a single strategy shows superior performance when using the long estimation window in terms of excess returns, suggesting that investors cannot benefit from longer estimation windows. When comparing standard deviation, the picture is more ambiguous with 9 out of 14 strategies showing lower standard deviation with the long estimation window. The results regarding Sharpe ratios are also ambiguous with 8 out of 14 performing better with the long estimation window. For certainty equivalent, 6 out of 14 strategies perform better with the long strategy, while it is 7 out of 14 for portfolio turnover.

All 14 strategies perform worse with the long estimation window when measuring opportunity cost in relation to the equally-weighted strategy. However, one should be cautious about making conclusions about this result, as it is clearly biased by the better performance of the equally-weighted portfolio in the different analysis period. When comparing opportunity costs in relation to the MinVarMA, 10 out of 14 outperform when the long estimation window is used.

Above comparison of the performance measures with the different lengths of estimation window, indicates that there is not any clear tendency of one estimation window being superior. Most notable are the returns, that are superior for all strategies when using the short estimation window.

Table 14 – Out-of-sample portfolio performance with 120 months estimation window. All measures are reported in percent except the Sharpe ratio.

Type	Name	Avg. rtn.	Excess rtn.	σ^k	SR^k	CEQ^k	PT^k	OC^k_{EW}	$OC^k_{MinVarMA}$
Benchmark	EW	5.71	4.49	7.54	0.60	4.21	1.93		1.33
		(4.82)	(3.15)	(7.71)	(0.41)	(2.85)	(2.06)		(-0.13)
	MinVarMA	4.38	3.17	3.29	0.96	3.11	1.26	-1.33	
		(4.95)	(3.28)	(3.32)	(0.99)	(3.22)	(1.13)	(0.13)	
With lag	MinVarHtLag	4.48	3.26	3.43	0.95	3.20	6.30	-1.23	0.09
		(4.95)	(3.28)	(3.45)	(0.95)	(3.22)	(5.71)	(0.13)	(0.00)
	MinVarFtDtLag	4.48	3.26	3.48	0.94	3.20	7.97	-1.23	0.09
		(5.15)	(3.48)	(3.54)	(0.98)	(3.42)	(9.80)	(0.33)	(0.20)
	MinVarHtFtDtLag	4.69	3.47	3.52	0.99	3.41	8.65	-1.02	0.31
		(5.32)	(3.65)	(3.50)	(1.04)	(3.59)	(10.29)	(0.5%)	(0.37%)
	MinVarFtLag	4.24	3.03	3.46	0.88	2.97	6.64	-1.47	-0.14
		(4.95)	(3.28)	(3.48)	(0.94)	(3.22)	(6.03)	(0.13)	(0.00)
	MinVarDtLag	4.48	3.27	3.30	0.99	3.21	7.85	-1.23	0.10
		(4.71)	(3.04)	(3.39)	(0.90)	(2.99)	(6.14)	(-0.11)	(-0.24)
Without lag	MinVarHt	4.58	3.36	3.51	0.96	3.30	3.18	-1.13	0.19
		(5.03)	(3.36)	(3.43)	(0.98)	(3.30)	(4.04)	(0.20)	(0.08)
	MinVarFtDt	4.66	3.44	3.50	0.98	3.38	5.98	-1.05	0.27
		(5.04)	(3.37)	(3.54)	(0.95)	(3.31)	(8.43)	(0.22)	(0.09)
	MinVarHtFtDt	4.84	3.62	3.57	1.01	3.56	8.45	-0.87	0.45
		(5.30)	(3.63)	(3.50)	(1.04)	(3.57)	(9.71)	(0.48)	(0.35)
	MinVarFt	4.39	3.17	3.54	0.90	3.11	4.03	-1.32	0.01
		(5.11)	(3.44)	(3.46)	(1.00)	(3.38)	(4.73)	(0.29)	(0.16)
	MinVarDt	4.76	3.55	3.42	1.04	3.49	6.09	-0.95	0.38
		(4.92)	(3.25)	(3.41)	(0.95)	(3.19)	(5.47)	(0.09)	(-0.03)
Other strategies	MinVarPrevVolCorrEst	4.33	3.11	3.56	0.87	3.05	7.58	-1.38	-0.06
		(4.82)	(3.15)	(3.32)	(0.95)	(3.09)	(6.53)	(-0.01)	(-0.13)
	MinVarPrevVolCorr	4.91	3.69	3.56	1.04	3.63	14.07	-0.80	0.52
		(5.19)	(3.52)	(3.54)	(0.99)	(3.45)	(13.15)	(0.37)	(0.24)
	MinVarZeroCorr	4.13	2.91	3.09	0.94	2.87	0.90	-1.58	-0.25
		(4.60)	(2.93)	(3.26)	(0.90)	(2.88)	(0.97)	(-0.22)	(-0.35)

Source: Own contribution

As in the primary analysis, we test the performance measures with the long estimation window for statistical difference compared to the two benchmark portfolios. In line with our primary analysis, we use stationary bootstrapping to obtain p-values for the performance measures with the 120 months' estimation window. The results are presented in table 15 below.

When comparing p-values for the difference in standard deviation, the picture is very similar to our primary analysis. Namely that each individual strategy has a significantly lower standard deviation than the EW strategy, while there is no statistical significance when comparing to the MinVarMA strategy. One notable difference, however, is that the MinVarZeroCorr becomes significantly different than the MinVarMA for a significance level of 10%.

Table 15 - Sensitivity analysis for estimation window of 120 months. Performance metrics are in percent except the Sharpe ratio and p-values are presented in parenthesis. Upper p-value compares EW and strategy k and the lower compares MinVarMA and strategy k.

Type	Strategy k	σ^k	SR^k	CEQ^k	PT^k	OC^k_{EW}	$OC^k_{MinVarMA}$
Benchmark	*EW*	7.54	0.60	4.21	1.93		1.33
	MinVarMA	3.29	0.96	3.11	1.26	-1.33	
		(0.0000)	*(0.0803)*	*(0.7502)*			
Historical data only	*MinVarHt*	3.51	0.96	3.30	3.18	-1.13	0.19
		(0.0000)	*(0.0451)*	*(0.7305)*			
		(0.9710)	*(0.4362)*	*(0.1343)*			
	MinVarHtLag	3.43	0.95	3.20	6.30	-1.23	0.09
		(0.0000)	*(0.0550)*	*(0.7459)*			
		(0.9920)	*(0.5247)*	*(0.3034)*			
Analyst forecast data only	*MinVarFt*	3.54	0.90	3.11	4.03	-1.32	0.01
		(0.0000)	*(0.0823)*	*(0.7734)*			
		(0.9908)	*(0.7875)*	*(0.5052)*			
	MinVarFtLag	3.46	0.88	2.97	6.64	-1.47	-0.14
		(0.0000)	*(0.1181)*	*(0.7899)*			
		(0.9873)	*(0.9048)*	*(0.7887)*			
	MinVarDt	3.42	1.04	3.49	6.09	-0.95	0.38
		(0.0000)	*(0.0487)*	*(0.6752)*			
		(0.9955)	*(0.0770)*	*(0.0160)*			
	MinVarDtLag	3.30	0.99	3.21	7.85	-1.23	0.10
		(0.0000)	*(0.0705)*	*(0.7292)*			
		(0.5575)	*(0.3207)*	*(0.2594)*			
	MinVarFtDt	3.50	0.98	3.38	5.98	-1.05	0.27
		(0.0000)	*(0.0411)*	*(0.7076)*			
		(0.9966)	*(0.3137)*	*(0.0656)*			
	MinVarFtDtLag	3.48	0.94	3.20	7.97	-1.23	0.09
		(0.0000)	*(0.0739)*	*(0.7408)*			
		(0.9967)	*(0.6251)*	*(0.3388)*			
Historical & Analyst forecast data	*MinVarHtFtDt*	3.57	1.01	3.56	8.45	-0.87	0.45
		(0.0000)	*(0.0231)*	*(0.6699)*			
		(1.0000)	*(0.2217)*	*(0.0278)*			
	MinVarHtFtDtLag	3.52	0.99	3.41	8.65	-1.02	0.31
		(0.0000)	*(0.0341)*	*(0.7005)*			
		(0.9999)	*(0.3479)*	*(0.1039)*			
Other strategies	*MinVarPrevVolCorrEst*	3.56	0.87	3.05	7.58	-1.38	-0.06
		(0.0001)	*(0.1073)*	*(0.7574)*			
		(0.5340)	*(0.6366)*	*(0.6102)*			
	MinVarPrevVolCorr	3.56	1.04	3.63	14.07	-0.80	0.52
		(0.0000)	*(0.0320)*	*(0.6399)*			
		(1.0000)	*(0.1629)*	*(0.0150)*			
	MinVarZeroCorr	3.09	0.94	2.87	0.90	-1.58	-0.25
		(0.0001)	*(0.1824)*	*(0.7628)*			
		(0.0731)	*(0.6415)*	*(0.8332)*			

Source: Own contribution

The influence on p-values for Sharpe ratios, when comparing the two estimation windows is more explicit. Particularly, when comparing strategies with the EW portfolio, which has a much higher Sharpe ratio for the long estimation window. In the primary analysis, 13 out of 14 alternative strategies performed significantly better with a significance level of five percent. For the long estimation window, it is only the case for 6 out of 14 alternative strategies. As with our primary analysis, we do not find any signs of statistically significant outperformance, when comparing Sharpe ratios with the MinVarMA benchmark.

When looking at certainty equivalent in our primary analysis, we do not find any significant values at a five percent level, and only one p-value below 10%. Namely, CEQ for the MinVarHtFtDtLag strategy, when comparing to the MinVarMA portfolio. However, when we investigate the significance of CEQ for the long estimation window, we find that MinVarDt, MinVarHtFtDt, and MinVarPreVolCorr all show significantly better performance in terms of CEQ compared to the best performing benchmark portfolio MinVarMA.

Generally, there does not seem to be a clear effect from the increase in estimation window from 60 to 120 months. The longer holding period seems to provide lower returns and lower risks. The consequence is an unambiguous effect on the Sharpe ratio. Comparing p-values for the long and short estimation window, we observe slightly more significant values for the long estimation window. We find that changing the length of the estimation window does not have any significant impact on our findings.

5.4.7 Different length of holding period and quarterly data

In our primary analysis we use monthly data and a holding period of one-month, meaning that the portfolio is rebalanced each month. It is reasonable to assume that some investors will tend to rebalance less frequently, and this section is included to test the performance of each strategy, when assuming a holding period of three months and using quarterly data. The table below reports the performance of each strategy.

An interesting observation from the table is, that returns (both absolute and excess) are worse for all strategies when the time between rebalances is increased to three months. A reason could be that individual strategies cannot react to investment opportunities as quickly. If a different allocation is optimal in the end of the first month, two additional months will have to pass before it is possible to rebalance.

With 8 out of 15 strategies experiencing a lower volatility when rebalancing less frequently, the effect on standard deviation seems to be ambiguous. The clearly negative effect on returns and the ambiguous effect on standard deviation results primarily in a negative effect on Sharpe ratio. Only the EW portfolio shows a minor improvement in Sharpe ratio from 0.41 to 0.42 when rebalancing less frequently. The remaining 14 portfolios shows significantly lower Sharpe ratios, with MinVarHtFtDt being most affected with a decrease in Sharpe ratio from 1.04 to 0.69. As a consequence of the consistently worse return, certainty equivalent is also worse for all strategies when compared to monthly rebalances.

Table 16 - Out-of-sample portfolio performance with quarterly data and three-months' holding period. All measures are reported in percent except the Sharpe ratio.

Type	Name	Avg. rtn.	Excess rtn.	σ^k	SR^k	CEQ^k	PT^k	OC^k_{EW}	$OC^k_{MinVarMA}$
Benchmark	EW	4.65	2.99	7.21	0.41	2.73	4.16		0.01
		(4.82)	(3.15)	(7.71)	(0.41)	(2.85)	(2.06)		(-0.13)
	MinVarMA	4.64	2.97	3.62	0.82	2.91	2.53	-0.01	
		(4.95)	(3.28)	(3.32)	(0.99)	(3.22)	(1.13)	(0.13)	
With lag	MinVarHtLag	4.68	3.01	3.51	0.86	2.95	10.19	0.03	0.04
		(4.95)	(3.28)	(3.45)	(0.95)	(3.22)	(5.71)	(0.13)	(0.00)
	MinVarFtDtLag	4.54	2.87	3.40	0.84	2.81	9.47	-0.12	-0.11
		(5.15)	(3.48)	(3.54)	(0.98)	(3.42)	(9.80)	(0.33)	(0.20)
	MinVarHtFtDtLag	4.40	2.73	3.37	0.81	2.67	9.77	-0.25	-0.24
		(5.32)	(3.65)	(3.50)	(1.04)	(3.59)	(10.29)	(0.50)	(0.37)
	MinVarFtLag	4.48	2.81	3.43	0.82	2.75	7.19	-0.18	-0.17
		(4.95)	(3.28)	(3.48)	(0.94)	(3.22)	(6.03)	(0.13)	(0.00)
	MinVarDtLag	4.12	2.45	3.35	0.73	2.39	8.14	-0.54	-0.52
		(4.71)	(3.04)	(3.39)	(0.90)	(2.99)	(6.14)	(-0.11)	(-0.24)
Without lag	MinVarHt	4.75	3.08	3.47	0.89	3.02	7.66	0.10	0.11
		(5.03)	(3.36)	(3.43)	(0.98)	(3.30)	(4.04)	(0.20)	(0.08)
	MinVarFtDt	4.59	2.92	3.45	0.85	2.86	9.25	-0.06	-0.05
		(5.04)	(3.37)	(3.54)	(0.95)	(3.31)	(8.43)	(0.22)	(0.09)
	MinVarHtFtDt	4.48	2.81	4.02	0.70	2.73	8.25	-0.18	-0.17
		(5.30)	(3.63)	(3.50)	(1.04)	(3.57)	(9.71)	(0.48)	(0.35)
	MinVarFt	4.62	2.95	3.52	0.84	2.89	7.13	-0.03	-0.02
		(5.11)	(3.44)	(3.46)	(1.00)	(3.38)	(4.73)	(0.29)	(0.16)
	MinVarDt	4.16	2.49	3.29	0.76	2.44	9.07	-0.49	-0.48
		(4.92)	(3.25)	(3.41)	(0.95)	(3.19)	(5.47)	(0.09)	(-0.03)
Other strategies	MinVarPrevVolCorrEst	4.43	2.76	3.44	0.80	2.70	5.59	-0.22	-0.21
		(4.82)	(3.15)	(3.32)	(0.95)	(3.09)	(6.53)	(-0.01)	(-0.13)
	MinVarPrevVolCorr	4.93	3.26	4.57	0.71	3.16	10.60	0.28	0.29
		(5.19)	(3.52)	(3.54)	(0.99)	(3.45)	(13.15)	(0.37)	(0.24)
	MinVarZeroCorr	4.25	2.58	3.86	0.67	2.51	2.14	-0.40	-0.39
		(4.60)	(2.93)	(3.26)	(0.90)	(2.88)	(0.97)	(-0.22)	(-0.35)

Source: Own contribution

Portfolio turnover is higher for 10 out of 15 strategies, which can be expected as a consequence of less frequent rebalancing. One should be careful about interpreting this difference, as it is a comparison of monthly and quarterly rebalances. Even though the latter is larger for two thirds of the portfolios, total transaction costs are likely to be higher for monthly rebalances.

Looking at opportunity costs in relation to the EW portfolio, there is a clear negative tendency, indicating that the alternative strategies are more negatively affected than is the EW benchmark strategy. Conversely, when considering opportunity costs in relation to the MinVarMA benchmark, 4 out of 14 strategies show higher opportunity costs. This indicates, that more often than not, the MinVarMA is more negatively affected than are the alternative strategies.

Again, in order to test the sensitivity of p-values from the primary analysis, we have calculated p-values for the difference in performance metrics when considering quarterly, as opposed to monthly, rebalancing. Table 17 below shows the findings from this sensitivity analysis.

The difference between standard deviation of the alternative strategies and the EW benchmark shows similar results as the primary analysis, namely that the EW performs significantly worse than the remaining 14 strategies. When considering the difference between alternative strategies and the MinVarMA portfolio, we do observe some differences to the primary analysis. In the original analysis, no alternative strategy presented significantly lower standard deviation. In contrast, when using quarterly data, 9 of 14 strategies have significantly lower standard deviation at a five percent level.

Looking at Sharpe ratios, in the primary analysis 13 out of 14 alternative strategies, significantly outperform the EW benchmark at a significance level of five percent. For quarterly data, the number is 12 out of 14, as the MinVarHtFtDt strategy becomes insignificant with a p-value of 0.0577. When comparing alternative strategies with MinVarMA, one strategy obtains statistically significant p-value at a significance level of five percent for quarterly data.

Looking at CEQ for monthly rebalances, no alternative strategy significantly outperforms the EW benchmark and only one strategy significantly outperforms the MinVarMA strategy at a 10% level. For quarterly rebalances, no alternative strategy outperforms either benchmark.

The sensitivity analysis in relation to quarterly data and rebalances does not indicate any major implications to our primary analysis. Most notable is the influence on standard deviation, however, the negative influence on returns means that there is no significant influence on neither Sharpe ratios nor certainty equivalent.

Table 17 - Sensitivity analysis with quarterly data. Performance measures are in percent except the Sharpe ratio and p-values are presented in parenthesis. The upper p-value compares EW and strategy k and the lower compares MinVarMA and strategy k.

Type	Strategy *k*	σ^k	SR^k	CEQ^k	PT^k	OC^k_{EW}	$OC^k_{MinVarMA}$
Benchmark	*EW*	7.21	0.41	2.73	4.16		0.01
	MinVarMA	3.62	0.82	2.91	2.53	-0.01	
		(0.0000)	*(0.0081)*	*(0.2577)*			
Historical data only	*MinVarHt*	3.47	0.89	3.02	7.66	0.10	0.11
		(0.0000)	*(0.0015)*	*(0.2189)*			
		(0.0129)	*(0.0291)*	*(0.2613)*			
	MinVarHtLag	3.51	0.86	2.95	10.19	0.03	0.04
		(0.0000)	*(0.0024)*	*(0.2299)*			
		(0.0696)	*(0.1378)*	*(0.4150)*			
Analyst forecast data only	*MinVarFt*	3.52	0.84	2.89	7.13	-0.03	-0.02
		(0.0000)	*(0.0032)*	*(0.2468)*			
		(0.0093)	*(0.1545)*	*(0.5415)*			
	MinVarFtLag	3.43	0.82	2.75	7.19	-0.18	-0.17
		(0.0000)	*(0.0055)*	*(0.2719)*			
		(0.0293)	*(0.2044)*	*(0.8532)*			
	MinVarDt	3.29	0.76	2.44	9.07	-0.49	-0.48
		(0.0000)	*(0.0067)*	*(0.3316)*			
		(0.0003)	*(0.7094)*	*(0.9818)*			
	MinVarDtLag	3.35	0.73	2.39	8.14	-0.54	-0.52
		(0.0000)	*(0.0154)*	*(0.3474)*			
		(0.0010)	*(0.8795)*	*(0.9890)*			
	MinVarFtDt	3.45	0.85	2.86	9.25	-0.06	-0.05
		(0.0000)	*(0.0015)*	*(0.2446)*			
		(0.0042)	*(0.1256)*	*(0.5839)*			
	MinVarFtDtLag	3.40	0.84	2.81	9.47	-0.12	-0.11
		(0.0000)	*(0.0021)*	*(0.2582)*			
		(0.0022)	*(0.0959)*	*(0.6944)*			
Historical & Analyst forecast data	*MinVarHtFtDt*	4.02	0.70	2.73	8.25	-0.18	-0.17
		(0.0000)	*(0.0577)*	*(0.3071)*			
		(0.9509)	*(0.9759)*	*(0.7530)*			
	MinVarHtFtDtLag	3.37	0.81	2.67	9.77	-0.25	-0.24
		(0.0000)	*(0.0032)*	*(0.2821)*			
		(0.0033)	*(0.3026)*	*(0.8873)*			
Other strategies	*MinVarPrevVolCorrEst*	3.44	0.80	2.70	5.59	-0.22	-0.21
		(0.0000)	*(0.0085)*	*(0.2889)*			
		(0.0025)	*(0.4912)*	*(0.9868)*			
	MinVarPrevVolCorr	4.57	0.71	3.16	10.60	0.28	0.29
		(0.0000)	*(0.0155)*	*(0.2224)*			
		(0.9856)	*(0.9927)*	*(0.1492)*			
	MinVarZeroCorr	3.86	0.67	2.51	2.14	-0.40	-0.39
		(0.0000)	*(0.0626)*	*(0.3378)*			
		(0.9613)	*(0.9962)*	*(0.9961)*			

Source: Own contribution

6 Conclusion and discussion

6.1 Concluding remarks on empirical results

As described in the introduction, the research question has been attempted answered in several steps. This is done by trying to answer sub-questions that ultimately provide a concluding answer to the research question. The first half of the research question looks into the determinants of future stock-bond correlation and whether forward-looking analyst forecasts of these variables have any explanatory power. After a methodical search of macroeconomic variables in the current literature, which previously have been found to explain stock-bond return co-movement, analyst forecasts as well as historical data on these variables are gathered. Additionally, the forecast data is used to create dispersion variables of the forecasters' mean consensus estimates of each variable, allowing to create a collection of dispersion measures. To help answer the first part of the research question, the following hypothesis was formalized in the introduction.

Hypothesis 1

H_0: The collection of analyst forecast variables, $\boldsymbol{F}_t$ and $\boldsymbol{D}_t$, do not explain variation in volatilities or correlation of stock and bond returns

In order to test hypothesis one, several predictive regression models are build using different combinations of historical data and forecast data to predict next months' volatilities and correlation. The predictive power of these variables are tested in-sample, that is, using data from the entire sample period from July 1994 to December 2018.

The predictive regression specifications, using only forward-looking analyst forecast data, outperformed the models using historical data only in terms of adjusted R^2. This is true for the regressions of both stock and bond volatility as well as stock-bond correlation. When the lagged variable is excluded from the regressions, the analyst forecasts seem to explain considerably more of the variation in volatilities and correlation compared to the historical variables. Hence, it suggests that in-sample, analyst forecast variables seem to add explanatory power over and above that of historical data.

Interestingly, we find that when the collection of mean consensus and dispersion measures of analyst forecasts are combined with historical data, the model provides an even better fit. Although the fit is only slightly better than the analyst forecast specification, it means that historical data possess some predictive power, which is not captured by forward-looking analyst forecasts. This is the case for all three in-sample regressions.

Turning to the statistical significance of the macroeconomic variables, the results show that all mean consensus and dispersion variables are statistically significant in predicting volatilities or stock-bond correlation except for mean consensus of corporate profits, F^{CP}, and the dispersion in forecasts of real GDP growth, D^{RGDP}. Many of the variables even show statistical significance at the one percent level. Another interesting finding is that most of the forecast variables remain significant after including the lag, whereas the opposite is true for the specification with historical data. When the lag is included, the few historical variables that showed statistical significance becomes insignificant. This suggests that only forward-looking analyst forecasts provide in-sample predictability when controlling for last month's realized value.

Lastly, the lag seems to be the single best in-sample predictor of both volatilities and correlation. The lag is not only statistically significant across all three regressions, but the size of the estimate is also economically significant. Comparing the size of the coefficient estimate of the lag with that of the other variables reveals that it is, by far, explaining most of the variation. The high contribution to adjusted R^2 is displayed in figure 11. From the remaining variables, dispersion in analyst forecasts of the short and long rate seem to be the second and third most economically significant variables.

Effectively, the null hypothesis, outlined in hypothesis one, is rejected. This means that implied stock market uncertainty, inflation rate, short and long rates as well as change in corporate profits and change in real GDP seem to be determinants of future stock-bond correlation and that analyst forecasts thereof do have statistically significant predictive power, even over and beyond historical data. This provides the answer to the first part of the research question.

The second part of the research question concerns the possibility of constructing mean-variance efficient portfolios based on the macroeconomic variables from the first part of the analysis, and whether these portfolios can outperform simple benchmark strategies. To help answer this part of the research question, the following hypothesis was formalized in the introduction.

Hypothesis 2

H_0: The performance of minimum-variance portfolios formed using the predictive regressions is worse than or equal to benchmark strategies

The hypothesis above is tested by forming minimum-variance portfolios where variance-covariance matrices, and in turn portfolio weights, are predicted using 60-months' rolling-window regressions. The purpose is to report the out-of-sample predictive power of analyst forecasts in relation to portfolio performance. The performance of each strategy is compared to an equally-weighted portfolio and a 60-months' moving-average benchmark strategy.

A total of 13 strategies are tested against the two benchmarks. The results of the analysis show that all but one strategy statistically significantly outperform the equally-weighted benchmark strategy in terms of Sharpe ratio and portfolio volatility. However, the picture is less clear regarding the certainty equivalent measure, where no strategy reports statistically significant outperformance.

This means that all strategies based on either historical data, forecast data, or both provide enough evidence against the null hypothesis, reported in hypothesis two, when the equally-weighted strategy is considered as the benchmark. Hence, if investors used the presented models to predict volatilities and stock-bond correlation to determine the allocation between stocks and bonds, they would have been better off than simply splitting wealth equally among the two asset classes.

When the 60-months' moving-average benchmark strategy is used instead, the results are less convincing. None of the 13 strategies provide statistically significant outperformance in any of the performance measures. However, the opportunity cost measure indicates that some of the strategies do economically outperform over the sample period. Nevertheless, the analysis reveals that the best performing strategies have considerably larger portfolio turnover than any of the benchmark strategies making them more expensive to implement. Effectively, when considering the MinVarMA as benchmark, all of the strategies fail to reject the null hypothesis presented by hypothesis two.

Similar to the in-sample analysis, the strategies using both historical and forecast data deliver the best out-of-sample portfolio performance. This is not surprising since the in-sample analysis also shows that predictive regressions using historical and forecast variables deliver the highest adjusted R^2. The difference between pure forecast strategies and pure historical data strategies are less clear since some perform better, while some perform worse.

Interestingly, some of the pure forecast strategies deliver better performance when the lag is excluded from the predictive regression. This is in contrast to the in-sample analysis, where the lag was the best predictor of next month's volatilities and correlation, which in turn suggests that analyst forecasts do deliver out-of-sample benefits in predicting co-movement of stock and bond returns.

Ultimately, the paper shows that implied stock market uncertainty, inflation rate, short and long rates as well as change in corporate profits and change in real GDP seem to predict future stock-bond correlation and the analyst forecasts thereof do have statistically significant predictive power, even over and above historical data. Additionally, using this information, investors are able to construct minimum-variance portfolios that significantly outperform the equally-weighted benchmark strategy. However, outperformance cannot be proven statistically significant compared to a 60-months' moving-average benchmark strategy.

6.1.1 Comparing findings to similar research

Our findings have several commonalities with previous research on the same topic but there are also several discrepancies. The inconsistencies might be due to differences in methodological structure of the studies. Li (2002) predicts stock-bond correlation from a predictive regression using macroeconomic variables but assumes volatilities and expected returns to be constant over time, thereby forming mean-variance instead of minimum-variance portfolios. He uses the same 60-months' moving-average portfolio as benchmark and finds that the strategy significantly outperforms the moving-average strategy in terms of certainty equivalent return. On the contrary, we could not establish statistical significance of the outperformance compared to the MinVarMA strategy. The differing results may be explained by the use of minimum-variance portfolios in our thesis.

In a similar study, Jivraj (2012) uses analyst forecasts to predict both volatilities and correlation of stock and bond returns. Interestingly, he fails to detect statistically superior performance compared to the equally-weighted portfolio, at a five percent significance level, although it is significant at the 10% percent level. Almost all strategies outperform the 1/N strategy in this thesis, but the difference may stem from differing sample periods and different sources of analyst forecast data. Jivraj's study also finds that portfolios based on analyst forecasts cannot outperform a moving-average benchmark strategy when employing a five percent significance level. Lastly, he demonstrates that removing the lagged value as an independent variable from the out-of-sample predictive regressions, increases portfolio performance, despite the lag showing best in-sample predictive power.

Lastly, the most recent comparable study uses forward-looking option implied information to forecast stock-bond correlation, while volatilities are estimated from historical data. The results show that minimum-variance portfolios do not outperform the naive 1/N benchmark portfolio, however the study rebalances more frequently with weekly and fortnightly rebalancing (DeMiguel *et al.*, 2013).

Hence the results in this thesis contribute to the existing literature with evidence suggesting that analyst forecasts can be used to outperform an equally-weighted portfolio, while the remaining results are in line with previous research.

6.2 Practical validity of our findings

In this paper, we identify strategies that seem to outperform the naive equally-weighted portfolio in a minimum-variance setting. Although this finding proposes potential academic importance, the practical and economic importance is not as clear. Firstly, the practical validity is conditioned of the premise that the benchmark strategy is actually practiced in the real world. Secondly, while the finding is valid in a minimum-variance setting, it might not be valid for investors that do not necessarily seek to minimize variance. This section will discuss these considerations.

6.2.1 Practical validity of the equally-weighted benchmark strategy

The 1/N rule is not a modern concept. Reportedly the rule dates back, at least, to the 4th century where the Rabbi Isaac bar Aha stated that *"One should always divide his wealth into three parts: a third in land, a third in merchandise, and a third ready to hand"* (Babylonian Talmud: Tractate Baba Mezi'a, folio 42a, as cited in DeMiguel, Garlappi and Uppal, 2009, p. 1).

The 1/N heuristic is evidently simple, and its use is both criticized and advocated by several researchers. The rule has been subject to extensive scrutiny and, in contrast to this thesis, a substantial amount of research find that the strategy performs well (DeMiguel, Garlappi and Uppal, 2009; Duchin and Levy, 2009). The acceptable performance in combination with its simplicity are probably the main reasons that the allocation rule is indeed practiced. For example, Samuelson and Zeckhauser (1988) studied, among other things, decision making in relation to pension plans. They found that the 50-50 split between stocks and bonds, by far, was the most popular allocation of pension contributions with about half of the participants choosing exactly the 50-50 split.

In line with these findings, Benartzi and Thaler (2001) find that investors distribute their wealth using the 1/N heuristic. Furthermore, Huberman and Jiang (2006) document that investors tend to allocate their wealth over relatively few assets and that they tend to allocate funds evenly over these few assets.

Finally, even the renowned pioneer of modern portfolio theory, Harry Markowitz, reportedly used the 1/N rule personally with the following argument: *"My intention was to minimize my future regret. So I split my contributions fifty-fifty between bonds and equities"* (Benartzi and Thaler, 2001, p. 80).

6.2.2 Validity of using minimum-variance portfolios

As presented in section 3.2 "The Markowitz portfolio optimization model", the global minimum-variance portfolio is defined as the portfolio, of all possible portfolios, with the lowest variance. According to modern portfolio theory, the GMVP is efficient from a mean-variance perspective, but not from a Sharpe ratio optimization perspective. The GMVP is characterized by relatively high exposure to low beta assets, and in theory it should be possible to increase portfolio Sharpe ratio by increasing the exposure to assets with higher risk and correspondingly higher return. In other words, because of diversification an investor can obtain a better risk-adjusted return if she accepts a higher level of risk than the GMVP, while the portfolio remains mean-variance efficient (this point is illustrated in figure 4 on page 18). Additionally, the theory stipulates that an asset's market beta should be proportionally and positively related to the return of the asset.

Although the model provides an acceptable approximation of financial markets, the empirical data does not exactly comply with the simplified theoretical model. One well documented discrepancy between theoretical expectations and empirical observations is, that high market-beta stocks are not proportionally rewarded by high returns. Fama and French (1992) for instance, document that the risk-reward anomaly is present in financial data from 1941 to 1990, and that it is consistent across firm sizes. More recently, Ang *et al.* (2006) find that the anomaly is present in data from 1986 to December 2000, and argues that it is explained by idiosyncratic risk. Their findings show that the least risky companies provide a CAPM-Alpha of 0.14, compared to -1.35 for the riskiest companies. Clarke *et al.* (2011) study the performance of minimum-variance portfolios and find that their cumulative excess returns are slightly higher than the cumulative excess return of the market in the period 1968 to 2009. Generally, the literature shows a tendency of academic recognition of minimum variance portfolios as a better choice than initially implied by modern portfolio theory.

The impressive performance of minimum-variance portfolios poses a puzzling violation of the risk-return principles of modern finance, and it has spurred an increasing popularity among investors. Clarke *et al.* (2011) argue that the recent increase in popularity, to a high degree, is driven by the increased appreciation for risk management subsequent to the financial crisis. The popularity of minimum-variance portfolios prompted MSCI to launch a range of global minimum-variance portfolios for various regions in the years 2008 and 2009 (MSCI, 2019). Figure 12 below shows the development of the MSCI USA Minimum Volatility Index and the MSCI USA Index over the last 15 years. Appendix A presents the development of the MSCI EAFE Minimum Volatility Index and the MSCI EAFE Index for benchmark comparison in the same 15-year period. The EAFE index includes all developed markets except North America.

Figure 12 – Development in the MSCI USA Minimum Volatility Index vs. MSCI USA Index. Indices indexed at 100 in May 2005

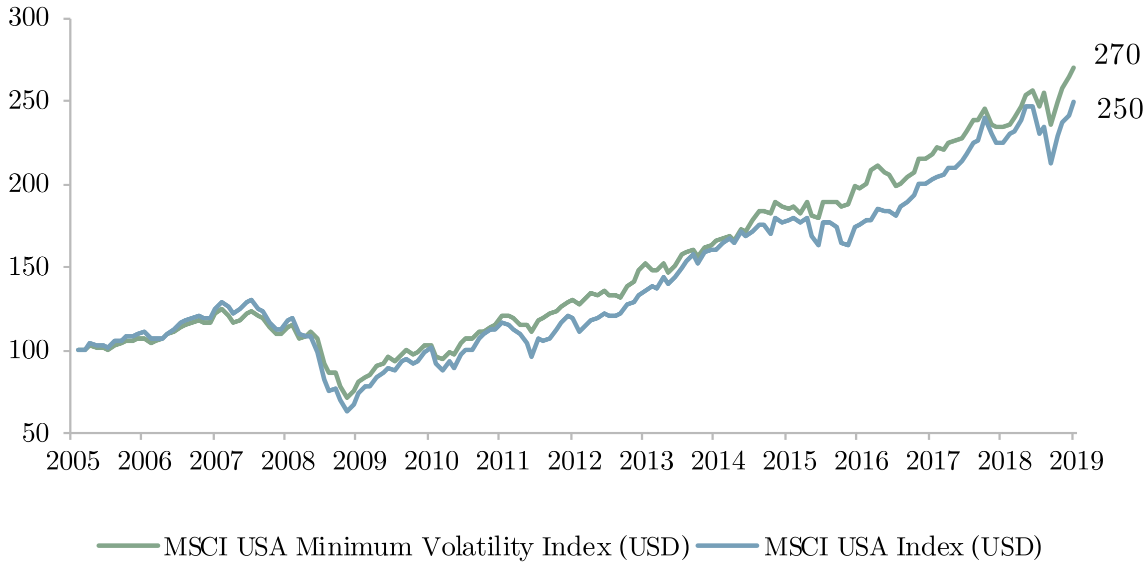

Source: (MSCI, 2019)

In line with the findings of Clarke *et al.* (2011), both minimum volatility indices outperform their respective benchmarks.

Another concern regarding the validity of using minimum-variance portfolios is the rather extreme realized portfolio weights of all MinVar strategies. Over the sample period July 1999 to December 2018 and across all 13 strategies, the average monthly allocation to stocks is 12.68% with the remaining 87.32% invested in bonds. Figure 13 below shows the development in portfolio weights across the sample period for the MinVarMA benchmark strategy and the best performing MinVarHtFtDtLag strategy. The development in portfolio weights for the remaining strategies can be found in appendix B.

It is not surprising that a lot of weight is put on low risk assets, that is, bonds, given we are forming minimum-variance portfolios. However, it might be difficult to convince an investor that such a limited exposure to equities is the best possible strategy, when the time-varying nature of volatilities and stock-bond correlation is taken into consideration.

Even though the results clearly show better performance, investors might be reluctant to adopt such strategies, given that more traditional asset allocation strategies often have far more equity exposure.

Figure 13 – Change in portfolio weights of stocks and bonds for the moving-average benchmark strategy and MinVarHtFtDtLag strategy across the sample period

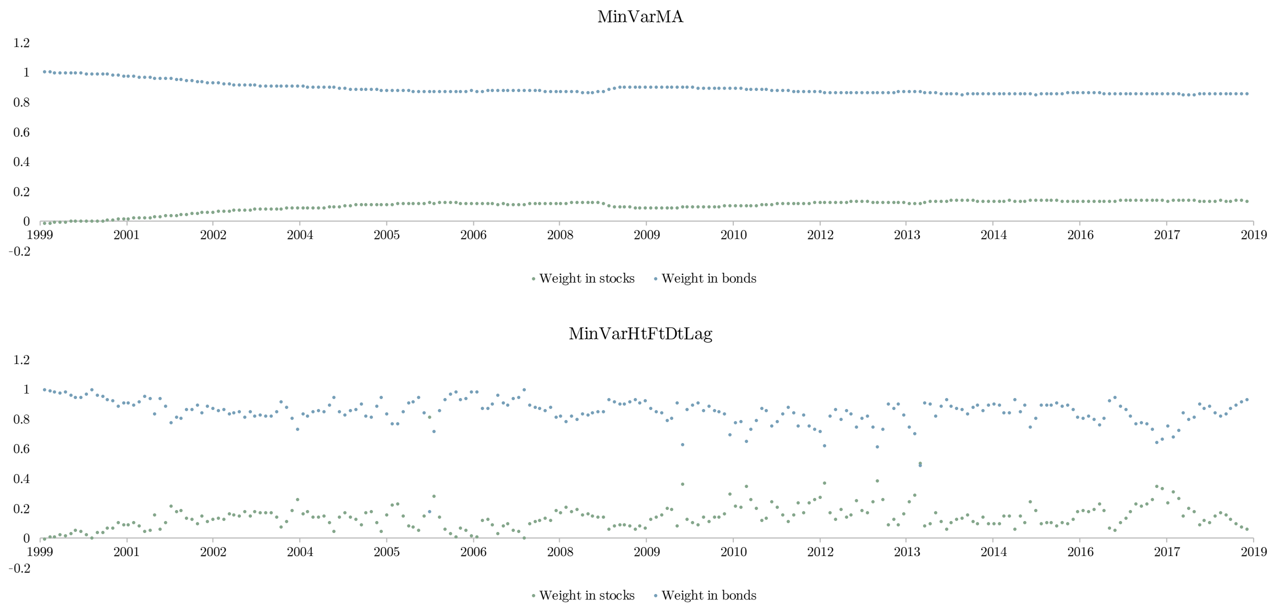

Source: Own contribution

6.2.3 Overall validity of the empirical results

With the discussion above and support from the literature, we argue that the 1/N rule is indeed practiced, which legitimizes the use of the equally-weighted portfolio as a benchmark. Furthermore, the discussion also establishes that global minimum-variance portfolios are a practical concern and not merely an academic bi-product from Markowitz's (1952) research, as investors do use them as investment vehicles. Hence, it does not seem as a far-fetched assumption that some investors are minimum-variance optimizers. All in all, this implies that our results are not merely of academic concern but have practical implications for investors. Using analyst forecasts of macroeconomic variables to make informed asset allocation decisions provide better risk-adjusted returns than the simple and widespread equally-weighted portfolio.

6.3 Asset liability management: Pension funds

The importance of pension funds in the financial markets is undisputed. In a recent report on the pension fund market, it was found that the industry's combined assets under management in the OECD countries have exceeded USD 40 trillion, corresponding to no less than 133.6% of the area's GDP (OECD, 2018). The same report shows that pension funds, on average, allocate 26.1% to equities and 43.9% to bonds. With 70% of pension funds' assets invested in stocks and bonds, stock-bond correlation becomes of practical importance.

The significant exposure to stocks and bonds signifies the importance of asset/liability management (ALM). Pension funds tend to be short net bonds and long net stocks. The former net position is primarily driven by future payments to pensioners as they encounter retirement age. The latter is the long positions in stocks, financed by contributions from individuals or employers (Jivraj, 2012a). From an ALM perspective, the discrepancy in net positions poses a potential risk for underfunding in periods where stock positions are declining dramatically in value. The financial crisis in 2008 and 2009 provides an example of the ALM impact in turbulent financial periods. In 2007, before the crisis, the gap between liabilities and assets for pension funds in OECD countries was 13%. By the end of 2009 the gap had increased to 26%, signifying a substantial worsening of the asset/liability structure of pension funds in the OECD area (Keeley and Love, 2010). The funding gap for pension funds in the crisis years showcases that ALM and mitigation procedures are a practical concern for pension funds.

Jivraj (2012) notes that there is limited focus on shocks in stock-bond correlation in the ALM literature. To solve the problem of decoupling asset and liability positions, Jivraj argues that a solution could be provided by using multi-asset class derivatives. Specifically, he argues that pension funds are too exposed to stock-bond correlation shocks and that this could be mitigated by increasing positions in, the already existing, stock-bond correlation swap contracts.

From our analysis, we observe a pattern that could potentially provide a more direct mitigation procedure in order to avoid the extensive decoupling of assets and liabilities for pension funds in periods with plunging stock prices. To illustrate the potential mitigation procedure, we have graphed the portfolio weights in stocks for the 15 investment strategies during the financial crisis years in figure 14 below. As our analysis only includes two assets and we condition the analysis on a 100% combined exposure, the weight in stocks is simply given by $w^S = 1 - w^B$ where w^S and w^B represent the percentage allocation in stocks and bonds, respectively. In consideration of visual clarity, we only illustrate equity positions in figure 14. Closing prices of the S&P 500 index has been included, on the right-hand axis, to illustrate the crisis period.

From the figure below, we observe a clear tendency of a reduction in equity exposure shortly subsequent to the fall of Leeman Brothers in September 2008. Only one strategy notably defies this tendency, namely the EW portfolio, which has a constant stock exposure of 50%. The average equity exposure from the remaining 14 portfolios shows an average reduction from 14.53% in October 2008 to 8.53% in November 2008. The delay from the decline in the S&P 500 index and the reaction in the stock exposure illustrates that the strategies where not able to predict the upcoming crisis.

Figure 14 - Equity exposure during the financial crisis of all 15 strategies. For comparison, the level of the S&P500 index is graphed on the right-hand axis.

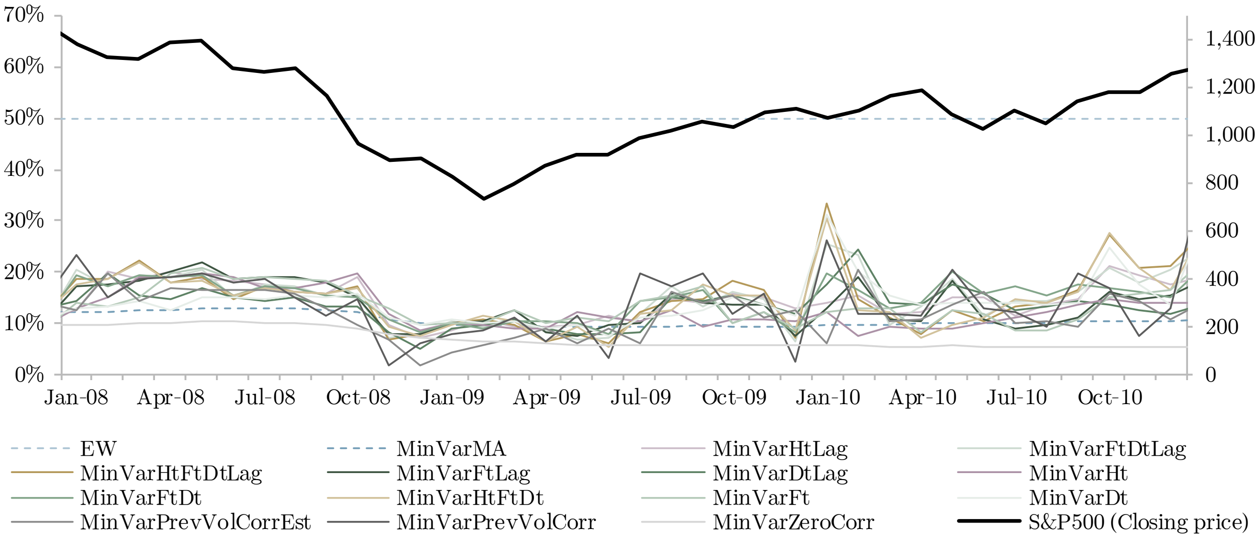

Source: Own contribution

However, the reduction in stock exposure to 8.53% in the beginning of November, shows that the strategies were quick to respond to the changing market environment. The exposure of the 14 portfolios remain low through June 2009, where the worst period of the financial crisis is over and the portfolios again start to increase the exposure to equities.

The analysis above shows that pension funds can use the proposed strategies in relation to asset/liability management. A clear example of this is provided by the comparison of the EW strategy and the MinVarHtFtDtLag strategy during the financial crisis. The latter strategy reacts to the crisis in October 2008, where the equity exposure is reduced from 17.19% to 7.01% the following month. The strategy holds a low equity exposure until July 2009 where it increases from 5.98% to 12.21%. By definition, the EW portfolio retains a constant stock exposure of 50% throughout the period. The fast and appropriate reaction from the MinVarHtFtDtLag strategy means that the portfolio manages to deliver a positive return of 3.27% over the two-year period from the beginning of 2008 to the end of 2009. In the same period the EW strategy loses 9.33%.

The above example illustrates that asset allocation strategies can be used as a mechanism to cope with the decoupling of pension funds' assets and liabilities in volatile market environments.

The MinVarHtFtDtLag strategy provides the best example of a strategy that reacts appropriately to the declining equity prices. Interestingly, this strategy also provided the best in-sample predictability of volatilities and correlation of stock and bond returns. This finding indicates that strategies, which take analyst forecasts into account, can prove to be valuable for pension funds' asset/liability management. However, thorough analysis with focus on the particular subject is required in order to make conclusive statements in this regard.

6.4 Future research

Our findings suggest that analyst forecasts possess valuable information in relation to future correlation of stocks and bonds. Our focus is rather narrow and there are several possible extensions. The first being an investigation of mean-variance portfolio optimization, in contrast to our minimum-variance approach. For this type of study, one is required to forecast future returns and it would be interesting to explore, whether the forecast data we apply, has predictive power in this regard.

Secondly, an extensive amount of literature uses relatively advanced econometric models, such as ARCH, GARCH, and M-GARCH, to estimate covariance between asset classes (see for example Bollerslev, Engle and Wooldridge, 1988; Kroner and Ng, 1998; de Goeij and Marquering, 2004). It would be interesting to investigate the relative performance of portfolios that are constructed using these historical data-based models in contrast to portfolios that are constructed using analyst forecast data.

Thirdly, an investigation of the potential benefits of predictive regressions to pension funds in mitigating their exposure to decoupling asset and liability positions, could be considered. Such a study would be a more academic extension to the discussion provided in the previous section.

Lastly, our analysis focuses on the correlation between stocks and bonds. As a multitude of additional asset classes are available to investors, analyst forecasts' predictive power in relation to the correlation between alternative asset classes could provide additional value for investors.

Bibliography

Andersen, T. G. *et al.* (2006) *Volatility and correlation forecasting*, *Handbook of Economic Forecasting*.

Andersson, M., Krylova, E. and Vähämaa, S. (2008) 'Why does the correlation between stock and bond returns vary over time?', *Applied Financial Economics*, 18(1), pp. 139–151.

Ang, A. *et al.* (2006) 'The cross-section of volatility and expected returns', *Journal of Finance*, 61(1), pp. 259–299.

Baele, L., Bakaert, G. and Inghelbrecht, K. (2010) 'The Determinants of Stock and Bond Return Comovements', *The Review of Financial Studies*, 23(6), pp. 2374–2428.

Benartzi, S. and Thaler, R. H. (2001) 'Naive diversification strategies in defined contribution saving plans', *American Economic Review*, 91(1), pp. 79–98.

Berk, J. and DeMarzo, P. (2017) *Corporate finance*. Fourth11 edn. Pearson.

Bodie, Z., Kane, A. and Marcus, A. (2014) *Investments*. 10th globa. Berkshire: McGraw-Hill Education.

Bollerslev, T. (1986) 'Generalized Autoregressive Conditional Heteroscedacity', *Journal of Econometrics*, 31(3), pp. 307–327.

Bollerslev, T., Engle, R. F. and Wooldridge, J. M. (1988) 'A Capital Asset Pricing Model with Time-Varying Covariances', *Journal of Political Economy*, 96(1), pp. 116–131.

Brinson, G. P., Hood, L. R. and Beebower, G. L. (1986) 'Determinants of Portfolio Performance', *Financial Analysts Journal*, 42(4), pp. 39–44.

Brinson, G. P., Singer, B. and Beebower, G. L. (1991) 'Determinants of Portfolio Performance II: An Update', *Financial Analysts Journal*, 47(3), pp. 40–48.

Campbell, J. Y. and Ammer, J. (1993) 'What Moves the Stock and Bond Markets? A Variance Decomposition for Long-Term Asset Returns', *The Journal of Finance*, 48(1), pp. 3–37.

Cappiello, L., Engle, R. and Sheppard, K. (2006) 'Asymmetric Dynamics in the Correlations of Global Equity and Bond Returns', 4(4), pp. 537–572.

Carhart, M. M. (1997) 'On Persistence in Mutual Fund Performance', LII(1), pp. 57–82.

Clarke, R., De Silva, H. and Thorley, S. (2011) 'Minimum-Variance Portfolio Composition', *The Journal of Portfolio Management*, 37(2), pp. 31–45.

Connolly, R., Stivers, C. and Sun, L. (2005) 'Stock market uncertainty and the stock-bond return relation', *Journal of Financial and Quantitative Analysis*, 40(1), pp. 161–194.

David, A. and Veronesi, P. (2008) *Inflation and Earnings Uncertainty and Volatility Forecasts: A Structural Form Approach*.

DeMiguel, V. *et al.* (2013) 'Improving Portfolio Selection Using Option-Implied Volatility and Skewness', *The Journal of Financial and Quantitative Analysis*, 48(6), pp. 1813–1845.

DeMiguel, V., Garlappi, L. and Uppal, R. (2009) 'Optimal versus naive diversification : How inefficient is the 1 / N portfolio strategy?', *Review of Financial Studies*, 22(5), pp. 1915–1953.

Duchin, R. and Levy, H. (2009) 'Markowitz Versus the Talmudic Portfolio Diversification Strategies', *The Journal of Portfolio Management*, 35(2), pp. 71–74.

Engle, R. (1982) 'Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation Author', *The Econometric Society*, 50(4), pp. 987–1007.

Engle, R. (2002) 'Dynamic conditional correlation: A simple class of multivariate generalized autoregressive conditional heteroskedasticity models', *Journal of Business and Economic Statistics*, 20(3), pp. 339–350.

Fama, E. F. and French, K. R. (1992) 'The Cross-Section of Expected Stock Returns', *The Journal of Finance*, 47(2), pp. 427–465.

Fleming, J., Kirby, C. and Ostdiek, B. (1998) 'Information and volatility linkages in the stock, bond, and money markets', *Journal of Financial Economics*, 49(1), pp. 111–137.

Fleming, J., Kirby, C. and Ostdiek, B. (2003) 'The economic value of volatility timing using "'realized'" volatility', *Journal of Financial Economics*, 67(3), pp. 473–509.

de Goeij, P. and Marquering, W. (2004) 'Modeling the Conditional Covariance between Stock and Bond Returns: A Multivariate GARCH Approach', *Journal of Financial Econometrics*, 2(4), pp. 531–564.

Gujarati, D. N. and Porter, D. (2008) *Basic Econometric*, *Fifth Edition*.

Gulko, L. (2002) 'Decoupling: If the U.S. Treasury repays its debt, what then?', *The Journal of Portfolio Management*, 28(3), pp. 59–66.

Hershey, J. C. and Schoemaker, P. J. H. (2008) 'Probability Versus Certainty Equivalence Methods in Utility Measurement: Are they Equivalent?', *Management Science*, 31(10), pp. 1213–1231.

Huberman, G. and Jiang, W. (2006) 'Offering versus choice in 401(k) plans: Equity exposure and number of funds', *Journal of Finance*, 61(2), pp. 763–801.

Ilmanen, A. (1985) 'Probability versus certainty equivalence methods in utility measurement: Are they equivalent?', *Management Science*, 31(10), pp. 1213–1231.

Inoue, A. and Kilian, L. (2002) *In-sample or out-of-sample tests of predictability: Which one should we use?*

Jegadeesh, N. and Titman, S. (1993) 'American Finance Association Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency', *The Journal of Finance*, 48(1), pp. 65–91.

Jivraj, F. (2012a) *Stock-Bond Correlation: Time variation , Predictability & Hedging*. Imperial College London.

Jivraj, F. (2012b) *Stock-Bond Correlation and Out-of-Sample Portfolio Performance Using Analyst Forecasts*.

Jivraj, F. and Kosowski, R. (2011) *Changing Expectations and the Correlation of Stocks and Bonds*.

Keeley, B. and Love, P. (2010) 'Pensions and the Crisis', in *From Crisis to Recovery: The Causes, Course and Consequences of the Great Recession*. Paris: OECD Publishing, pp. 68–87.

Keynes, J. M. (1936) *The General Theory of Output, Interest and Money*, *The Collected Writings of John Maynard Keynes*.

Kilian, L. and Taylor, M. P. (2003) 'Why is it difficult to beat the random walk forecast of exchange rates?', *Journal of International Economics*, 60(1), pp. 85–107.

Kostakis, A., Panigirtzoglou, N. and Skiadopoulos, G. (2011) 'Market Timing with Option-Implied Distributions: A Forward-Looking Approach', *Management Science*, 57(7), pp. 1231–1249.

Kroner, K. F. and Ng, V. K. (1998) 'Modeling Asymmetric Comovements of Asset Returns', *The Review of Financial Studies*, 11(4), pp. 817–844.

Li, L. (2002) *Macroeconomic factors and the correlation of stock and bond returns*, *Yale ICF Working Paper*.

Lintner, J. (1965) 'The valuation of risk assets and the selection of risky investments in stock portfolios and capital budgets', *The Review of Economics and Statistics*, 47(1), pp. 13–37.

Markowitz, H. (1952) 'Portfolio selection', *The Journal of Finance*, 7(1), pp. 77–91.

Markowitz, H. (1959) 'Portfolio selection: Efficient diversification of investments', *The*

Journal of Finance, 15(3), p. 447.

Mossin, J. (1966) 'Equilibrium in a capital asset market', *The Econometric Society*, 34(4), pp. 768–783.

MSCI (2019) *Minimum Variance Indices*. Available at: https://www.msci.com/end-of-day-history?chart=country&priceLevel=40&scope=C&style=M&currency=15&size=30&indexId=91630 (Accessed: 1 May 2019).

Munk, C. (2017) *Financial Markets and Investments*.

OECD (2018) 'Pension Markets In Focus', *Pension Markets in Focus*.

Pedersen, L. H. (2015) *Efficiently inefficient*. Princeton University Press.

Piazzesi, M., Salomao, J. and Schneider, M. (2011) *Trend and Cycle in Bond Premia, Working Paper*.

Politis, D. and Romano, J. (1994) 'The Stationary Bootstrap', *Journal of the American Statistical Association*, 89(428), pp. 1303–1313.

Pratt, J. W. (1964) 'Risk Aversion in the Small and in the Large', *Econometrica*, 32(1), pp. 122–136.

Roberts, A. J. (2011) *Elementary Calculus of Financial Mathematics*, *Elementary Calculus of Financial Mathematics*.

Samuelson, W. and Zeckhauser, R. (1988) 'Status Quo Bias in Decision Making', *Journal of Risk and Uncertainty*, 1(1), pp. 7–59.

Schwert, W. G. and Seguin, P. J. (1990) 'Heteroskedasticity in Stock Returns', *The Journal of Finance*.

Scruggs, J. T. and Glabadanidis, P. (2003) 'Risk Premia and the Dynamic Covariance between Stock and Bond Returns', *The Journal of Financial and Quantitative Analysis*, 38(2), pp. 295–316.

Sharpe, W. F. (1966) 'Mutual fund performance', *The Journal of Business*, 39(1), pp. 119–138.

Shiller, R. and Beltratti, A. (1992) 'Stock prices and bond yields: Can their comovements be explained in terms of present value models?', *Journal of Monetary Economics*, 30(1), pp. 25–46.

Shiller, R. J. (1989) *Market Volatility*. MIT Press.

Simaan, Y. (1993) 'What is the opportunity cost of mean-variance investment strategies?', *Management Science*, 39(5), pp. 578–587.

Studenmund, A. H. (2014) *Using econometrics: a practical guide*. 6. edition. Pearson Education.

Treynor, J. (1961) *Market value, time, and risk*, *Originally unpublished*.

Treynor, J. L. (1965) 'How to rate management of investment funds', *Havard business review*, 43(1), pp. 63–75.

Viceira, L. M. (2012) 'Bond Risk, Bond Return Volatility, and the Term Structure of Interest Rates', *International Journal of Forecasting*, 28(1), pp. 97–117.

White, H. (1980) 'A heteroskedasticity-consistent covariance matrix estimator and a direct test for heteroskedasticity'. Econometrica, 48(4), pp. 817–838.

William, S. (1964) 'Capital asset prices: A theory of market equilibrium under conditions of risk', *The Journal of Finance*, 19(3), pp. 425–442.

Appendices

Appendix A

Appendix A – Development in MSCI EAFE Minimum Volatility vs. MSCI EAFE

Indices are indexed at 100 in May 2005

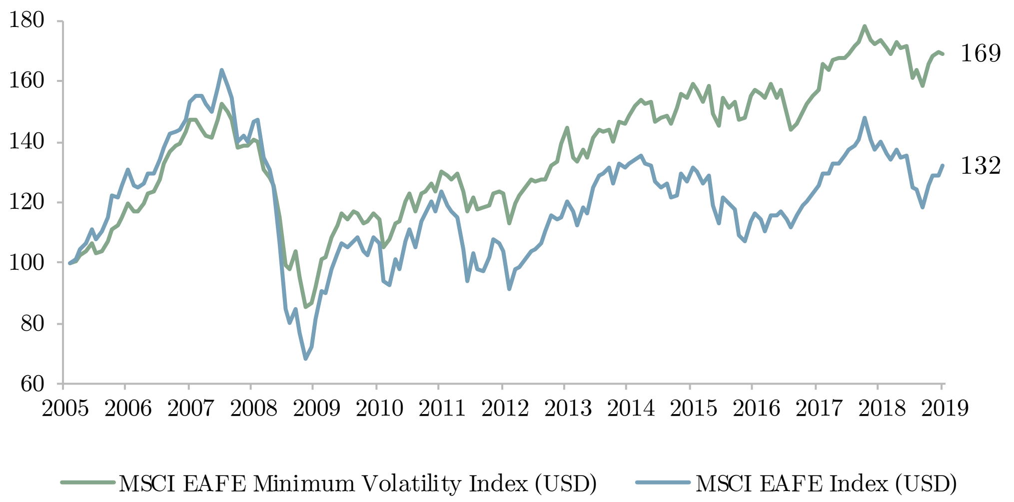

Source: (MSCI, 2019)

Appendix B

Appendix B – Change in portfolio weights of stocks and bonds for all strategies

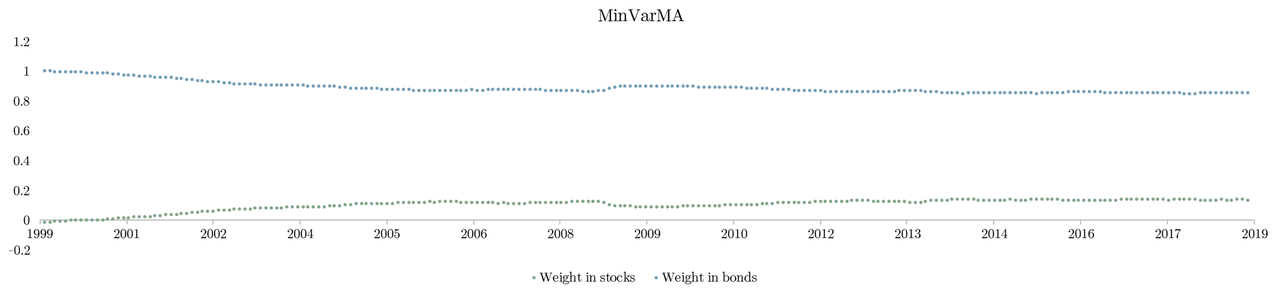

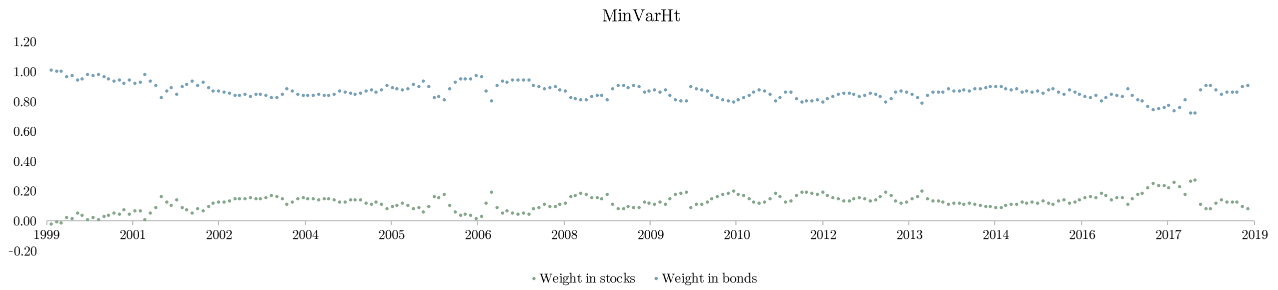

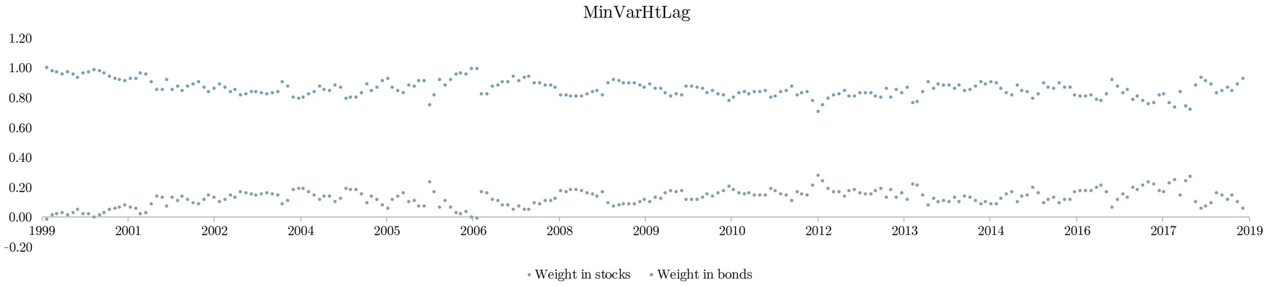
MinVarHtLag
Weight in stocks
Weight in bonds

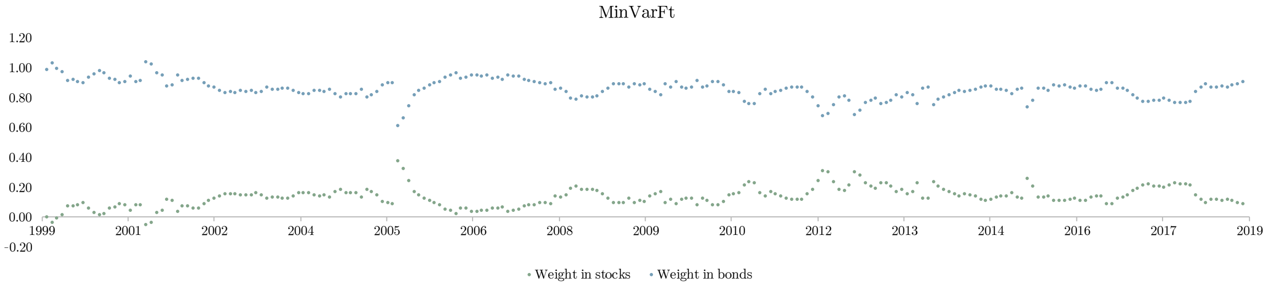
MinVarFt
Weight in stocks
Weight in bonds

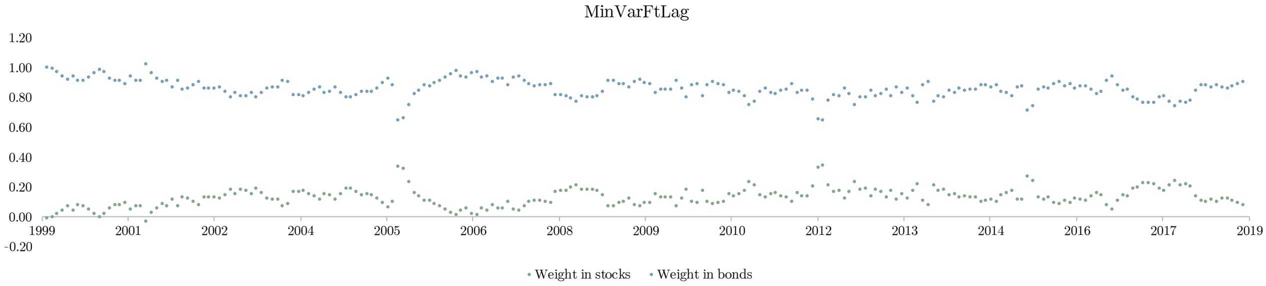
MinVarFtLag
Weight in stocks
Weight in bonds

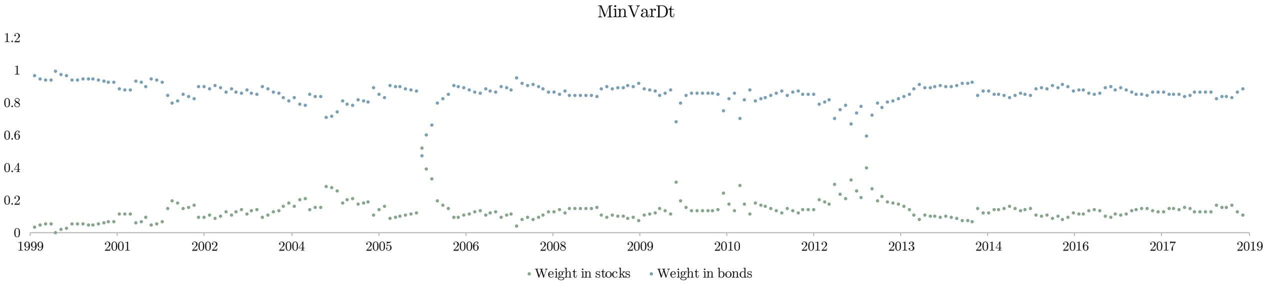
MinVarDt
Weight in stocks
Weight in bonds

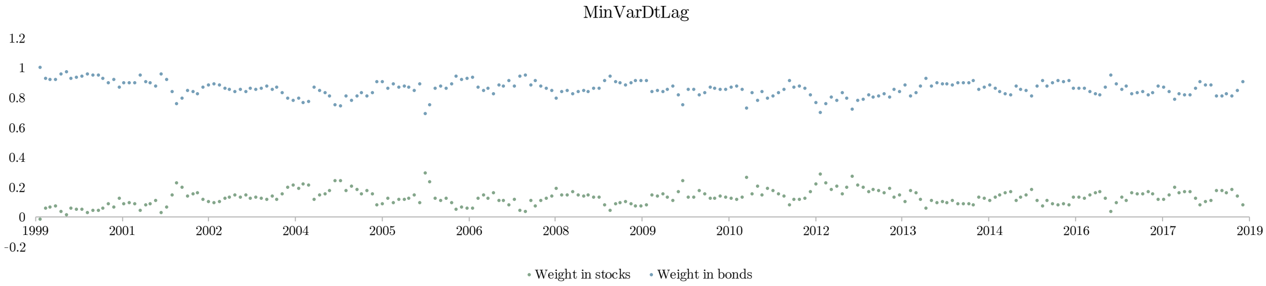
MinVarDtLag
Weight in stocks
Weight in bonds

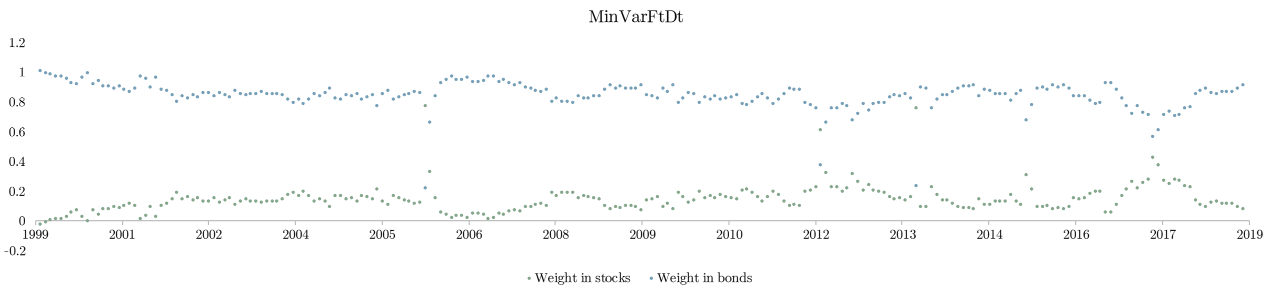
MinVarFtDt
1.2
1
0.8
0.6
0.4
0.2
0
-0.2
1999 2001 2002 2004 2005 2006 2008 2009 2010 2012 2013 2014 2016 2017 2019
Weight in stocks
Weight in bonds

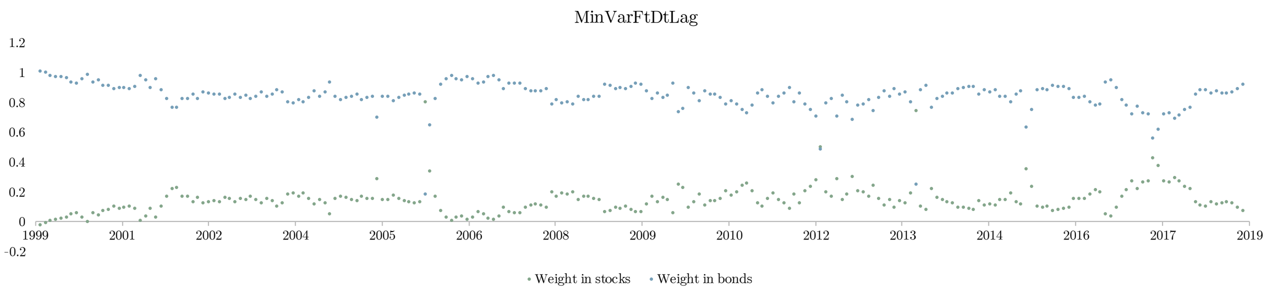
MinVarFtDtLag
1.2
1
0.8
0.6
0.4
0.2
0
-0.2
1999 2001 2002 2004 2005 2006 2008 2009 2010 2012 2013 2014 2016 2017 2019
Weight in stocks
Weight in bonds

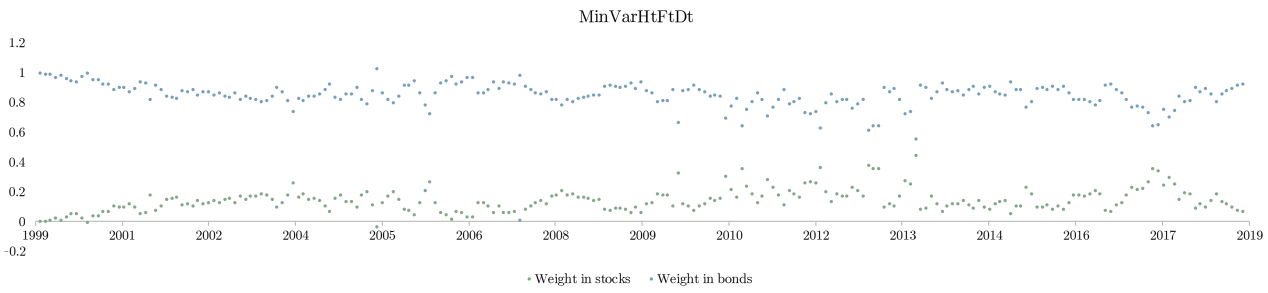
MinVarHtFtDt
1.2
1
0.8
0.6
0.4
0.2
0
-0.2
1999 2001 2002 2004 2005 2006 2008 2009 2010 2012 2013 2014 2016 2017 2019
Weight in stocks
Weight in bonds

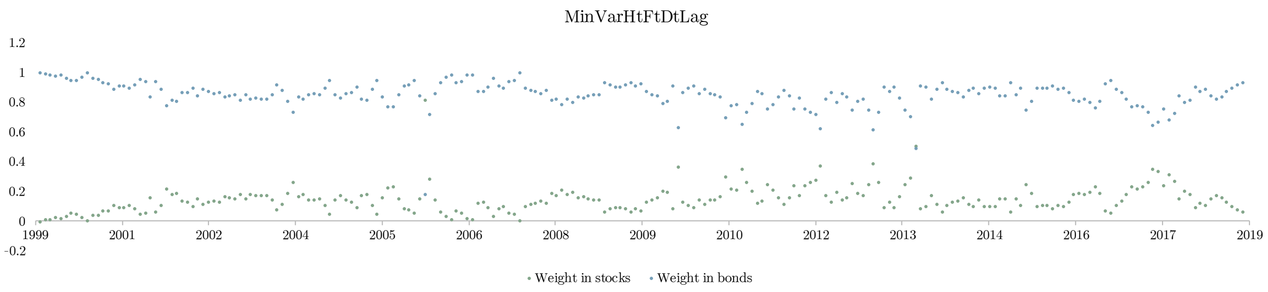
MinVarHtFtDtLag
1.2
1
0.8
0.6
0.4
0.2
0
-0.2
1999 2001 2002 2004 2005 2006 2008 2009 2010 2012 2013 2014 2016 2017 2019
Weight in stocks
Weight in bonds

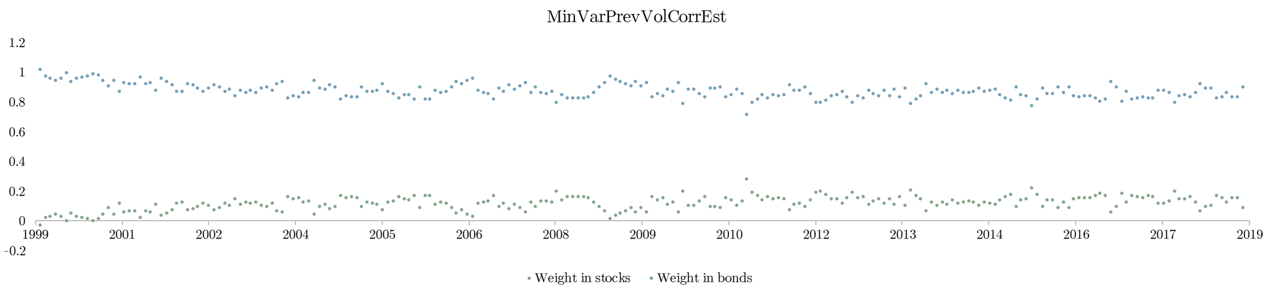
MinVarPrevVolCorrEst
1.2
1
0.8
0.6
0.4
0.2
0
-0.2
1999 2001 2002 2004 2005 2006 2008 2009 2010 2012 2013 2014 2016 2017 2019
Weight in stocks
Weight in bonds

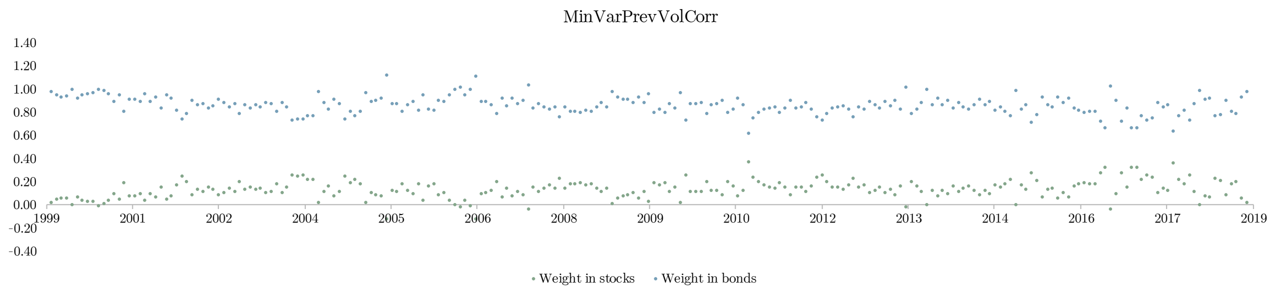
MinVarPrevVolCorr
1.40
1.20
1.00
0.80
0.60
0.40
0.20
0.00
-0.20
-0.40
1999
2001
2002
2004
2005
2006
2008
2009
2010
2012
2013
2014
2016
2017
2019
Weight in stocks
Weight in bonds

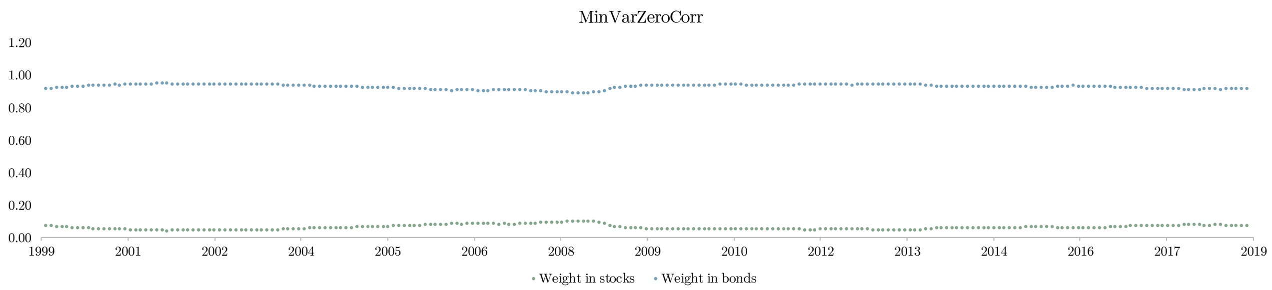
MinVarZeroCorr
1.20
1.00
0.80
0.60
0.40
0.20
0.00
1999
2001
2002
2004
2005
2006
2008
2009
2010
2012
2013
2014
2016
2017
2019
Weight in stocks
Weight in bonds