LONG TERM VOLATILITY OF STOCK RETURNS IN TURKEY:
UNCONDITIONAL AND PREDICTIVE VARIANCE ANALYSES

ALP EREN AKYÜZ

BOĞAZİÇİ UNIVERSITY

2021

LONG TERM VOLATILITY OF STOCK RETURNS IN TURKEY: UNCONDITIONAL AND PREDICTIVE VARIANCE ANALYSES

Thesis submitted to the
Institute for Graduate Studies in Social Sciences
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy
in
Management

by
Alp Eren Akyüz

Boğaziçi University
2021

DECLARATION OF ORIGINALITY

I, Alp Eren Akyüz, certify that

- I am the sole author of this thesis and that I have fully acknowledged and documented in my thesis all sources of ideas and words, including digital resources, which have been produced or published by another person or institution;
- this thesis contains no material that has been submitted or accepted for a degree or diploma in any other educational institution;
- this is a true copy of the thesis approved by my advisor and thesis committee at Boğaziçi University, including final revisions required by them.

Signature…………………………………………….....

Date …………………………………………..............

ABSTRACT

Long Term Volatility of Stock Returns in Turkey:
Unconditional and Predictive Variance Analyses

Long term volatility of stock returns plays a major role in determining the weight of stocks in forward looking portfolios. This thesis investigates the long run stock return volatility in Turkey in two parts: i) unconditional statistics derived from rolling windows samples, ii) conditional statistics derived from a predictive variance analysis using a Bayesian Markov Chain Monte Carlo approach. Unconditional variance decreases with the investment horizon and it becomes more likely for stock returns to beat fixed interest returns. The predictive variance analysis also suggests that return volatility decreases with time. The risk-inducing effect of momentum is dominated by the risk-reducing effect of the negative correlation between the error terms of the current and expected return equations. Assuming a time-varying covariance matrix reduces the portion of predictive variance attributable to the identical and independently distributed risks. Having fewer observations increases the predictive variance estimate. Overall results suggest that it is more preferable from an investor’s perspective to make long-term investments in Borsa Istanbul.

ÖZET

Türkiye'deki Hisse Getirilerinin Uzun Vadedeki Oynaklığı: Koşulsuz ve Öngörücü Varyans Analizleri

Hisse getirilerinin uzun vadedeki oynaklığı ileriye dönük portföylerde hisselerin ağırlığının belirlenmesinde önemli bir rol oynar. Bu tez, Türkiye'deki hisse getirilerinin uzun dönem oynaklığını iki parça halinde incelenmiştir: i) kayan pencerelerden elde edilen koşulsuz istatistikler, ve, ii) Bayesian Markov Zinciri Monte Carlo yöntemiyle yapılan öngörücü (prediktif) varyans analizinden elde edilen koşullu istatistikler aracılığıyla. Koşulsuz varyans yatırım ufku uzadıkça azalmakta ve hisse getirilerinin sabit faiz getirilerini aşması daha olası hale gelmektedir. Öngörücü varyans analizi de getiri oynaklığının zaman ile birlikte azaldığını göstermektedir. Şu anki ve beklenen getiri denklemlerinin hata terimleri arasındaki negatif korelasyonun risk-düşürücü etkisi momentumun risk-arttırıcı etkisine baskın gelmektedir. Zamana bağlı değişen bir kovaryans matrisi kullanmak öngörücü varyansın eş ve bağımsız dağılan (i.i.d.) riskler tarafından açıklanan kısmını azaltmaktadır. Daha az sayıda gözleme sahip olmak öngörücü varyansı arttırmaktadır. Genel sonuçlar, bir yatırımcının bakış açısından Borsa İstanbul'da uzun vadeli yatırımlar yapmanın daha tercih edilebilir olduğu yönündedir.

CURRICULUM VITAE

NAME: Alp Eren Akyüz

DEGREES AWARDED

PhD in Management, 2021, Boğaziçi University

MPhil in Economics, 2013, University of Oxford

BA in Business Administration, 2008, Marmara University

AREAS OF SPECIAL INTEREST

Long Term Investments, Foreign Aid, Econometrics, Programming

PROFESSIONAL EXPERIENCE

EU Project Research Assistant, Istanbul Technical University, 2013-2014

Research Intern, Borsa Istanbul, 2014

Casual Research Assistant, Oxford Centre for the Analysis of Resource Rich Economies, 2012

Research Assistant, Boğaziçi University, 2009-2011

Research Assistant, Gebze Technical University, 2008-2009

AWARDS AND HONORS

Republic of Turkey, Ministry of National Education Graduate Scholarship

First Prize in Deloitte Scientific Research Paper Contest: “Ethics and Business Life”

Tekfen Holding Graduate Scholarship

Istanbul Metropolitan Municipality Undergraduate Scholarship

PUBLICATIONS

The determinants of the composition of foreign aid [Master's thesis], 2013, University of Oxford.

The effects of business partners’ ethical conduct on firm reputation as perceived by its stakeholders, 2009, Published by Deloitte Academy Istanbul as a separate volume.

ACKNOWLEDGEMENTS

First and foremost, I would like to thank my supervisor Prof. Vedat Akgiray for his guidance and support throughout my studies at Boğaziçi University, especially for being so patient with me. His mentorship always motivated me to reach one step further and be a better researcher.

Secondly, I am grateful for the many insights provided to me by my thesis progress committee members Assist. Prof. Ali Coşkun and Prof. Burak Saltoğlu. This thesis would be lacking many crucial parts without their valuable critique and encouragement.

I would like to express my gratitude towards the Department of Management at Boğaziçi University for letting me take the first steps of my academic career as a research assistant and endowing me with the work discipline that served me so well since then. I would also like to thank Prof. Dr. Nesrin Okay and Assoc. Prof. Ulaş Akküçük for the invaluable experience I gained as their teaching assistant. I am grateful to Assoc. Prof. Ozan Hatipoğlu for supervising me during my work on foreign aid which I would argue was the most difficult and formative one for me.

I would like to thank the University of Oxford for the once in a lifetime opportunities I enjoyed throughout two wonderful years. The culture of excellence and competitiveness rubbed off on me there is the source of many breakthroughs I have made in both academia and life. I would also like to thank Prof. Anthony Venables for his mentorship on my past research and challenging me to be better.

I am thankful for the time I spent at Boğaziçi University and the many friends I have made throughout the years. There are too many to list here, but those I hold dear no doubt know that they are the ones that make this campus and program so special for me.

Last but not least, my deepest gratitude goes to my family for supporting all kinds of endeavors I pursue in this life. They are my light.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

CHAPTER 1

INTRODUCTION

The goal of this thesis is to examine the behavior of stock returns over the investment horizon to assess the preferability of stocks in the long run. According to conventional wisdom, the variance of stock returns decreases with time and investors with longer horizons should weigh stocks more in their portfolios. This notion is reflected in the investment advice given to young professionals as well as the portfolio choices made by pension and target date funds. These funds assign more weight to stocks the longer the investment horizon, resulting in a shift from stocks to fixed income assets over time known as the glide path. Long-run stock return variance plays an integral role in determining the shape and features of the glide path; therefore, improvements to its estimation would benefit long term investors.

The conventional wisdom is usually based on a number of unconditional statistics measuring the riskiness of stock returns as a function of investment horizon as well as conditional regressions that model the stock return variance based on a set of predictors. These methods can be argued to neglect components that would enable them to model the problem from an investor’s perspective. In practice, investors face additional estimation risk in the form of parameter uncertainties on top of the riskiness measured by these methods. Approaching from an investor’s perspective requires accounting for these risks and uncertainties, which would drive up the variance measure to potentially overturn the results. Therefore, conventional wisdom should be revisited to assess the extent of its validity while addressing the uncertainties faced by investors.

One proposed solution is calculating the predictive variance of stock returns using a predictive system. Bayesian Markov Chain Monte Carlo (MCMC) analysis accounts for estimation risk while allowing investor beliefs to be incorporated into the system. The analysis is based on a vector auto-regression (VAR) model known as the predictive system with temporal statistical relations (Markov Chain) and is calculated by iterating over the same sample thousands of times (Monte Carlo) to arrive at posterior distributions of all unknowns in the model (Bayesian analysis). Integrating over these distributions accounts for parameter uncertainties and yields the final result. This approach has the advantage of incorporating investor beliefs such as the momentum of expected returns, correlation between shocks to the return and expected return equations, or whether the covariance matrix is time varying. The predictive variance can further be decomposed into its components over the investment horizon to identify the opposing forces in effect and their magnitude.

The main reason provided for the conventional wisdom is mean reversion, i.e., the tendency of stock returns to revert back after shocks. For example, a positive shock to the current return is likely to be followed by a negative shock in the future; therefore, making the shocks to the expected return negative. This negative correlation between the error terms of the current and expected return is one of the main drivers of mean reversion and will constitute a key aspect of the predictive system used in the analysis in later sections.

Return variance has an independent and identically distributed (i.i.d.) component introducing the same amount of risk at each period. Stacking these over time results in a linear increase in variance over the investment horizon. Mean reversion counteracts this increase and depending on its strength may result in a decreasing overall variance.

Mean reversion and i.i.d. risks are not the only factors affecting the final value of predictive variance. Stock return is a function of expected return in a predictive system. Expected return is hard to model due to its unobservable nature. Uncertainties pertaining to the current and future values of expected return increase predictive variance over time. Moreover, expected return may have momentum, i.e., shocks to its value take time to disperse, resulting in another increase to the predictive variance.

Expected return varies with time; making it potentially predictable using a set of predictors which may introduce predictor imperfections to the system. This would still be the case even if there were no estimation risk, i.e., parameters values were known, since the true value of the expected return can only be proxied and not observed. Predictive systems used here account for predictor imperfections whilst calculating the predictive variance. It should be noted that the inclusion of predictors is not necessary in a predictive system. There have been cases in the literature where controlling for momentum and negative error term correlations make predictors negligible. It can be argued that statistically significant predictors can be found even after controlling for these effects; however, their inclusion would reduce the uncertainties in the model and decrease the predictive variance. This would only improve the argument for the preferability of stocks in the long run; therefore, their inclusion would be redundant if a case could already be made without them and can be later pursued for modeling accuracy. Barberis (2000) shows that the predictability of stock returns is weak. Accounting for estimation risk therefore becomes a necessity to prevent overallocation to stocks. Investors are still found to follow the conventional wisdom.

Before checking whether the conventional wisdom holds from an investor's perspective, it must be first established whether it holds at all for Borsa Istanbul using its conventional measures. The analysis here starts by calculating the unconditional variance of stock returns using rolling windows samples for both nominal and real returns. These results constitute a useful benchmark to compare the predictive variance results against. The difference between the unconditional and predictive variances can also be examined to calculate the contributions made to the predictive variance by the additional risks faced by investors using all available information on top of the unconditional variance.

Unconditional variance calculations require listing the stock returns for all relevant investment lengths. Distributions of stock returns for each investment horizon can be obtained using these samples. These distributions can be utilized to find the sample probabilities of stock returns beating fixed interest rates. Plotting them against the investment horizon would document the evolution of an investor's chances of surpassing a predetermined target.

The main contribution of this thesis is the adaptation of the Bayesian MCMC framework for the Turkish stock return data for nominal and real returns to calculate the predictive variance. The analysis is performed using the R programming language under various assumptions and investor beliefs to produce a comprehensive list of outcomes that are robust to changes in the initial values and prior distributions. Overall results suggest that accounting for estimation risk indeed yields a higher variance for all investment horizons; however, it is still more preferable from an investor's perspective to hold stocks for longer durations. Consequently, investments with long horizons such as pension and target date funds stand to benefit from weighing stocks more in their portfolios.

A dedicated section addresses the consequences of having a small sample size on the results. Many developing countries have relatively young stock markets, making it hard to analyze the behavior of long run stock returns simply because the long run has not been reached yet. Furthermore, the Bayesian MCMC method is data hungry: it requires a couple of decades to approximate the values in the system. Simply converting the time intervals to shorter periods would not automatically solve this problem, since the long run in that case would reduce to a couple of months rather than years. In the case of Turkey, around 32 years of stock data exists so the analysis can be performed without any hitches; however, it must be assessed whether the low number of observations put an inherent limit on the ability to estimate the predictive variance. The question asked here is: “does the low number of observations introduce additional uncertainties for investors in calculating the predictive variance?”. In order to answer this, the U.S. stock return data is analyzed under two conditions: i) by using all available data; and ii) restricting the sample size to 32 years. Predictive variances calculated using restricted samples result in higher values and they overturn the results for the U.S. markets. Consequently, it can be concluded that the sample size has a bearing on the results and a low number of observations constitute another form of uncertainty regarding the predictive variance for investors.

The remainder of this thesis is organized as follows: chapter two reviews the relevant literature, chapter three handles the unconditional part of the analysis, chapter four carries out the predictive variance calculations, and chapter five concludes. A dedicated chapter explains the Bayesian MCMC algorithm in detail.

CHAPTER 2

LITERATURE REVIEW

The literature on the long-run volatility of stock returns is limited for developing financial markets mainly due to their relatively short history. Long-run analyses typically require decades of data that are yet to be accumulated for many emerging markets. In the case of Turkey, there is more than three decades of stock returns data. This allows for an adequate analysis; however, a previous body of work to draw from or compare the results against is still lacking. For these reasons, the predictive variance side of the review focuses on the methodology and results obtained for other countries to highlight the major features of this approach.

In order to paint a comprehensive picture, three bodies of relevant literature are reviewed. First, the arguments for and against the conventional wisdom, i.e., stock riskiness reducing over time are briefly highlighted to provide a general framework. Second, the literature on stock return volatility and its predictors in Turkey is examined. These studies provide a list of popular predictors as well as tests of predictability, random walk, momentum and mean reversion in the stock market. Third and mainly, the literature on predictive variance and Bayesian MCMC methodology is summarized.

2.1 Unconditional variance

Unconditional variance is the principal estimate used in supporting many of the conventional wisdom arguments. A particularly popular one is the time diversification puzzle which is based on the observation that the return variance of k-period long investments is lower than that of *k* consecutive one-period investments.

The time diversification puzzle gave rise to a contentious debate that garnered the attention of many research papers with no definitive conclusion. Proponents of time diversification suggest that investors are better off buying high-risk high-return stocks and argue that risks will be lowered by the passage of time, leaving only high returns. Detractors of this idea point out that despite the reduction in risks, the amount exposed increases at a faster pace, making the investors worse off. These arguments are revisited in a later section where sample probabilities of beating fixed interest rates are examined.

The time diversification argument is mostly based on observations made through model-free statistics without an underlying theory; therefore, it is hard to both defend and refute. The analysis here does not aim to assess the validity of time diversification: the focus is not on whether to buy more or less risky stocks in the long run. The aim here is to analyze whether stocks of any kind should be held for short or long periods. Time diversification literature is still relevant since the debate around it has produced a series of arguments and counterarguments that also apply to the analysis of long run stock return variance. The origins and summaries of these arguments are provided here as a framework to employ when commenting on the results.

Unconditional variance is essentially a sample variance calculated without an underlying model; therefore, it consists a model-free benchmark to compare conditional models against. The samples are constructed from the available dataset using the rolling windows method. In order to construct these samples, relevant investment horizons are listed. The rolling windows are constructed to be the same length as the horizon they correspond to. The windows are not consecutive, they overlap. Each window starts one day after the previous one. The last window is constructed such that its final day is also the final day of the dataset. Consequently, fewer observations will be available for longer horizons. Considering that the sample variance formula punishes lower observations, it should be kept in mind that longer horizons will be at a disadvantage in terms of variance calculations.

The main explanation provided for the time diversification puzzle is mean reversion: the proposed inherent tendency of stock returns to return to their expected value after deviations. The mean need not be constant, it can change over time depending on whether positive or negative returns dominate. If mean reversion exists and positive returns dominate enough in the long run, investors enjoy lowered risks and positive gains in the long run. Therefore, a strand of the literature focuses on investigating whether stock returns exhibit mean-reverting behavior as a test of time diversification.

Some detractors of time diversification argue that stock returns follow a random walk: they are independent and identically distributed (i.i.d.) over time. The price of a stock with i.i.d. returns would follow a meandering path across time. In the presence of random walk, long-term stock returns and their variance can end up anywhere on the board. Moreover, variance would increase with the investment duration. Time diversification cannot exist with full random walk. The random walk hypothesis is one of the key propositions of the efficient markets hypothesis. Therefore, tests of random walk can be utilized as tests of time diversification.

A few research papers can be examined to capture the essence of these arguments. A recent work supporting the conventional wisdom for the U.S. stocks is Siegel (2008). Using a large sample with more than two centuries of observations, this study concludes that longer investments in stocks have lower variances. After 15 years, stocks yield more than bonds and bills. If the period is long enough, the equity premium would guarantee a less risky return. The dataset compiled by Siegel (2008) is also used by the Bayesian MCMC analyses on the U.S. data. Reichenstein and Dorsett (1995) conduct two tests based on two opposing assumptions: random walk and mean-reversion. In either case, they find some support for the time diversification argument. Fisher and Statman (1999) show that investors can wait for an opportune moment to sell for a profit, provided that mean reversion exists. Short run investments also mean more transactions, the costs of which often eat up the earnings, hence even with good tactical allocations it is hard to turn a profit as shown by Tokat and Stockton (2006).

Bodie (1995) argues that higher option premiums suggest higher perceived risk on the part of investors hence the time diversification argument does not hold up. Bodie, Merton, and Samuelson (1992) and Samuelson (1994) draw from the expected utility theory and argue that wealth is not only valued at the end of the investment horizon but at each point in time. Pension and target date funds may be viewed as exceptions, since funds usually become available at the terminal date.

Kritzman (1994) suggests that over long horizons, above-average returns offset below-average returns. The probability of losing money reduces; however, the amount in danger increases faster than this reduction. The investors would be worse off in the end.

In summary, the main arguments for the preferability of stocks are rooted in observations of unconditional variance, lacking the predictive intuition introduced by later approaches. The main counterargument to the conventional wisdom is the faster growth in the amount exposed despite the decrease in volatility. These arguments will be revisited during the unconditional portion of the analysis.

2.2 Analyses of stock returns and variance in Turkey

The literature on stock return variance in Turkey can be reviewed under two categories: i) papers modeling volatility as a function of predictors, ii) papers modeling stock returns as a function of predictors. Each will briefly be reviewed to arrive at a list of popular predictors and also highlight the features of the dataset such as predictability, momentum, mean-reversion and random walk.

Volatility regressions focus on how the riskiness of stocks change conditional on predictor values. Sönmez (2007) shows that inflation has a positive effect. Umutlu, Akdeniz and Altay-Salih (2010) show that financial liberalization have negative effects. Li, Nguyen, Pham and Wei (2011) show that large foreign ownership have negative effects, respectively. Akbaş (2013) shows that stock return and interest rate are cointegrated. Kasman, Vardar and Tunç (2011) show that interest rate and exchange rate are the major determinants of bank stock return volatility. In a slightly tangential paper, Kasman and Torun (2007) show that there is long memory in daily stock returns and volatility. Kandır (2008) finds that exchange rate, interest rate and market return determine the portfolio returns of non-financial stocks. Rjoub, Türsoy and Günsel (2009) find that inflation, interest rate, risk premium and money supply are significant but weak predictors of stock return. Sum (2012) shows that economic policy uncertainty negatively affects stock returns. Toparlı, Çatık, and Balcılar (2019) use a time-varying model and find that interest rate and exchange rate are the main drivers of stock returns.

A related study by Welch and Goyal (2008) examines a comprehensive list of predictors and find that they are not very useful for out-of-sample estimations of stock return performance. Investors are cautioned against relying too much on models with good in-sample results.

Bajo-Rubio, Berke and McMillan (2017) studies volatility spillovers between the Turkish stock market and international stock, exchange rate and commodity markets. Vardar and Aydoğan (2019) find that there are bidirectional volatility spillovers between Bitcoin and stock returns.

Predictability of stock returns plays an important role in reducing stock return variance over time. In its absence, the error terms would exhibit random walk and the variance would be constantly increasing over the investment horizon. Therefore, tests of predictability along with momentum and mean reversion would provide preliminary results signaling long run variance reduction. In addition, such studies would provide a list of possible candidates to be used as predictors here. Examples of such predictors include crude oil prices, interest rates, foreign exchange rates, consumer price index, gold price, money supply, industrial production index, risk premium, MSCI world equity index, economic policy uncertainty, trade balance.

It should be noted that predictors may be imperfect as shown by Pastor and Stambaugh (2009). Fortunately, Bayesian MCMC methods can account for this.

2.3 Predictive variance and Bayesian MCMC

Predictive systems have mainly been proposed to solve estimation and inference problems arising from serial correlation between residuals and predictors, persistence in predictors, and correlation between return residuals and predictor innovations commonly encountered in finite samples that cannot be fully addressed by predictive regressions. The main aim is to arrive at better forecasting models. A byproduct of this research came by in the form of long-run return volatility estimations which yield a statistic that is shortly called as the predictive variance. Predictive systems allow estimating stock return volatilities while accounting for modeling shortcomings and incorporating beliefs a real-life investor would have. Therefore, the resulting predictive variance is thought to reflect an investor's perspective.

Predictive systems are essentially Vector Auto-Regression models that are estimated using the Bayesian MCMC methodology. They model return as a function of expected return. Expected return and predictors are modeled as auto-regressive series. Predictors need not be directly tied to return and expected return through their equations: it is possible for them to be related through the covariance matrix of the system. The covariance matrix can be constant or time-varying, depending on the assumptions.

This section of the literature details this approach and compares the results with a special focus on four articles: Pastor and Stambaugh (2012), Carvalho, Lopes and McCulloch (2018), and Johannes, Korteweg, and Polson (2014).

Pastor and Stambaugh (2012) propose a predictive system for stock returns accounting for imperfect predictors continuing from their previous work in Pastor and Stambaugh (2009). In their 2012 article, they propose a state-space system modeling the co-movements of returns, conditional expected returns, and predictors. Stock returns, r_{t+1}, are written as a function of conditional expected return, μ_t:

$$r_{t+1} = \mu_t + u_{t+1}$$

while the conditional expected return is an auto-regressive process of order one (AR(1)):

$$\mu_{t+1} = \alpha + \beta\mu_t + w_{t+1}$$

where $\beta \in (0,1)$ and $\rho_{uw} < 0$

The key assumption here is that the correlation between the error terms of these two equations is negative, i.e., $\rho_{uw} < 0$. This creates a mean-reversion effect where a shock to the expected future returns (μ_t) are accompanied by an unexpected counter-shock to returns (r_{t+1}). In other words, an unexpected increase in conditional expected return is not thought to be sustainable.

The β term in the AR(1) equation is assumed to be between zero and one for stationarity. This coefficient guarantees momentum in returns, meaning the effects of shocks to the expected return will persist for a duration that depends on how close β is to one.

The assumptions $\beta \in (0,1)$ and $\rho_{uw} < 0$ create two opposing forces, suggesting that shocks to the system will persist for a while and then dissipate, similar to what is observed in real life. The speed of diversion and reversion depends on these parameters.

The negative correlation between the error terms gains a special importance when predictors are imperfect. Imperfections will later be explained in detail; however, for now it should be noted that the main difference between the cases of having perfect versus imperfect predictors is that imperfect predictors require all lagged values to be considered in forming equations. With perfect predictors, it is sufficient to keep all the variables at level with, i.e., featuring only the same period values in the equation rather than lagged ones. If imperfect predictors are to be acknowledged, relations tying the behavior of variables from consecutive periods must be defined. In the predictive systems setup, this means the term ρ_{uw} assumes a critical role in inference. Predictive regressions ignore the negativity of this correlation term, resulting in incorrect assessments of predictor significance and explanatory power. Inference with predictive regressions is invalid if ρ_{uw} is negative. Predictive systems must be used instead.

Another problem arising from predictor imperfections is serially correlated residuals. Serial correlation in residuals means that standard errors are miscalculated, once again causing inference to be flawed. In other words, the significance of predictors cannot be correctly assessed. Stambaugh (1999) furthermore suggest that this serial correlation is not the only issue: predictor persistence and correlations between the residuals and error terms in predictors worsen inference in finite samples.

Persistent predictors are problematic because they lead to spurious regressions as shown by Ferson, Sarkissian, and Simin (2003). Spurious regressions attribute explanatory power to predictors in finite samples when they have none in the population. Pastor and Stambaugh (2012) argue that the term ρ_{uw} can also help solve this problem: spurious predictors will not result in a negative correlation. In fact, an estimate for the term ρ_{uw} can be used to detect whether there are any spurious results.

In summary, imperfect predictors cause the results obtained from predictive regressions to be misleading. Standard errors calculated using the residuals obtained from predictive regressions are incorrect and cannot be used in statistical inference. Testing for predictor imperfections is possible. Pastor and Stambaugh (2012) argues that imperfect predictors will result in the residuals from the predictive regression of r_{t+1} on x_t having a non-zero correlation.

In order to overcome these problems, a Bayesian VAR setting is proposed by Pastor and Stambaugh (2012), dubbed “predictive system”. The features of this system are explained in detail in Section 4.2.

Pastor and Stambaugh (2012) present arguments on why $\rho_{uw} < 0$ is a plausible assumption. It is shown to be perfectly negative in the case of bonds since ρ_{uw} basically ties asset prices to discount rates which are negatively correlated. In the case of stocks, this negative correlation is less than perfect since cash flow shocks also play a role. Large cash flows can offset negative shocks and yield a positive ρ_{uw}. In short, in order for this assumption to be violated, stock returns must be less volatile than the expected value of dividends. This can be easily shown to not happen in reality since stock prices would follow a well-defined curve rather than the meandering path they are observed to take in reality.

Another finding supporting $\rho_{uw} < 0$ comes from the works of Shiller (1981) and LeRoy and Porter (1981). They observe a negative relation between the volatilities of stock returns and expected returns. Stocks display “excess volatility” when expected returns are near constant. This excess is explained by taking discount rate volatility into account. Discount rate volatility combined with cash flow shocks explains why stock returns are more volatile than expected returns.

While $\rho_{uw} < 0$ means that return history imposes a negative effect on expected returns, β imposes a positive effect. Pastor and Stambaugh (2012) shows that it is possible to calculate a breakeven value for ρ_{uw} where the two effects cancel each other.

Pastor and Stambaugh (2012) also focuses on the volatility derived from predictive systems. They decompose the predictive variance into five components (the direction of their effect is provided in brackets): *i)* i.i.d. uncertainty [+], *ii)* mean reversion [-], *iii)* uncertainty about future expected returns [+], *iv)* uncertainty about current expected return [+], and *v)* estimation risk [+]. Mean reversion is the only one that reduces variance, the rest increase it. The largest contribution to variance is by the uncertainty about future expected returns, μ_{t+1}. As the investment horizon lengthens, the remaining four positive contributors overcome the effect of mean reversion, resulting in predictive variance first decreasing then increasing with the investment horizon. Therefore, the conclusion arrived by Pastor and Stambaugh (2012) is that stocks are more volatile in the long run from an investor's perspective.

Johannes et al. (2014) differ in their approach as they perform a predictive regression without a predictive system and introduce time-varying covariances. They also use monthly data instead of annual returns. Their results based on the implied volatility calculated from their model increases with the investment horizon. Their treatment assumes time variation in volatilities but not the ρ_{uw} parameter.

Carvalho et al. (2018) directly iterate from the results of Pastor and Stambaugh (2012) and Johannes et al. (2014) by introducing time-varying volatilities and covariances to their predictive system. They conclude that with plausible priors, the result of Pastor and Stambaugh (2012) is reversed: predictive variance decreases with the investment horizon. They extend their analysis to cover more countries than the U.S. alone and discover that the results echo except for the German and Japanese stock exchanges.

The predictive variance section of the analysis presented here follows from the contributions made by Carvalho et al. (2018) and adapts the methodology for Borsa Istanbul. It extends the approach by performing the analysis for both real and nominal returns. In addition to adapting the model, a contribution to the literature is made by investigating the consequences of having a small sample size in the Bayesian MCMC analysis. The results suggest that the predictive variance estimates are prone to: i) being higher, and ii) curving upwards earlier with small samples. In other words, predictive variance estimates obtained for developing countries with typically young stock markets are likely biased upwards, making it seem riskier to invest in the long run.

CHAPTER 3

UNCONDITIONAL ANALYSIS

3.1 Data

The data used for this section consists of the daily closing prices of the index summarizing the performance of the stocks in Borsa Istanbul: XUTUM (the index of all stocks). In addition, monthly consumer price index (CPI) data is used to derive the real returns from the nominal ones. The steps followed are described in the following methodology section. Figure 1 displays the increasing pattern of XUTUM over time. Figure 2 shows the real returns which have a mean-reverting profile. Figure 3 plots the CPI deflator.

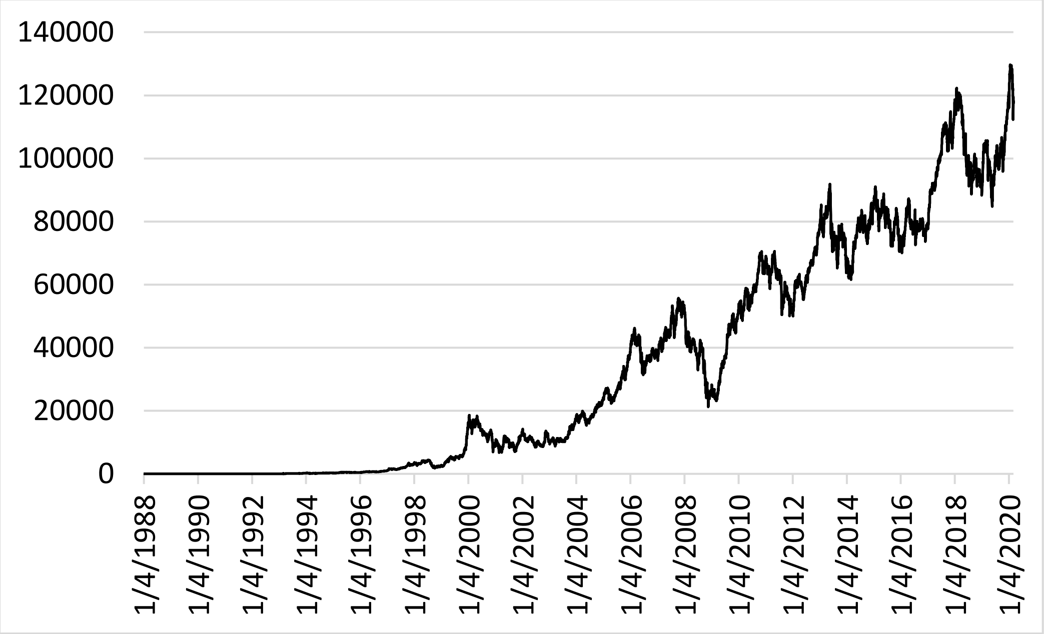

Figure 1. Nominal value of XUTUM

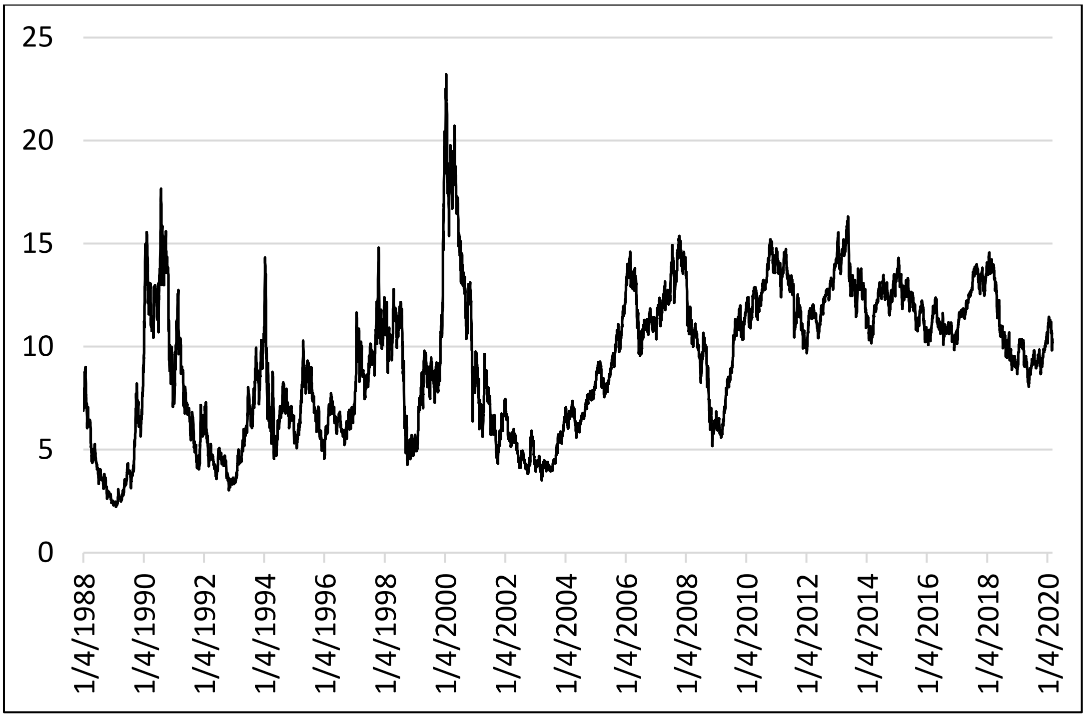

Figure 2. Real values of XUTUM

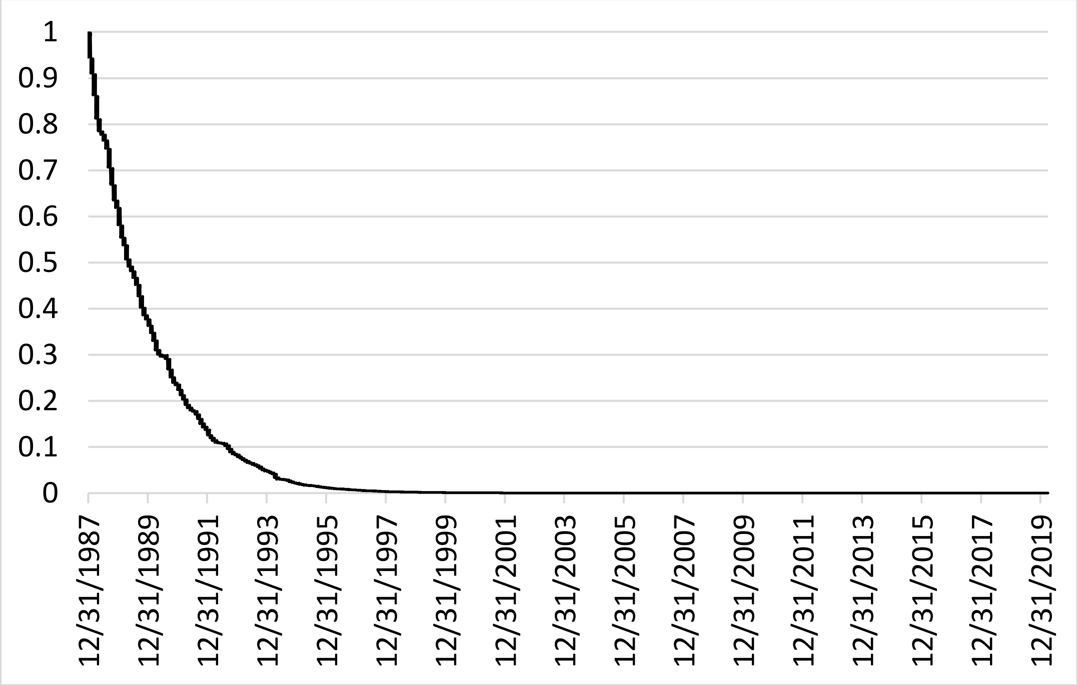

Figure 3. CPI deflator

3.2 Methodology

The first part of the analysis computes an unconditional variance along with a few other statistics to illustrate the behavior of stock returns over different investment horizons. These statistics are all calculated without conditioning them on any other variables or assumptions, hence they are not thought to be a realistic approximation of investors’ estimations. They are mainly computed to summarize the statistical properties of stock returns and risk in detail; as well as to have a benchmark to compare the results from the Bayesian analysis against. Alongside unconditional variance, the distribution of the general mass of returns is also plotted over different investment horizons in moving boxplot graphs. Lastly, comparative probabilities of stock returns beating various fixed income returns are plotted over different investment horizons to emphasize how the actual chances of investors making an excess return on their investment have been in the past.

Our treatment starts by first replacing the missing values through linear interpolation. Since the days missing are usually a few consecutive ones, this would mean that daily returns would appear to be slightly less volatile than they really are, hence a version of the analysis omitting these interpolated values is also performed as a robustness check.

The nominal return data is converted to real returns using a daily deflator obtained by linearly interpolating the monthly consumer price index data. The changes in the consumer price index (CPI) basket are transferred by keeping the ratio between the consecutive days when the change occurred. The final values for the deflator start with 1 on the first day and applied to the nominal returns dataset to arrive at the real return series.

All of the calculations are carried out using a rolling windows method to construct subsamples of stock returns associated with different investment horizons. These subsamples are constructed using the daily returns data. Compared to the annual return data later used in the Bayesian analysis chapter of the thesis, unconditional variance is calculated using a dataset with higher frequency. Therefore, it allows extending the analysis to cover investment horizons shorter than one year.

Rolling windows samples can be constructed for any investment length. For the sake of brevity, a list of relevant investment lengths is prepared. Then this list is iterated through by taking one investment length at a time and constructing all possible windows that can be associated with it. For each specific investment length h, the rolling window samples consist of a series of returns obtained by starting from the first day and calculating the returns of h-day investment periods by iterating one day forward. This iteration would stop once the end is reached, meaning that the last window would start on the day numbered $T - h$. The number of total return values that can be calculated for an investment length of h is $T - h$. Since T is constant, this means that fewer observations will be available for longer investment lengths. This property results in stock return variance estimates of longer investment lengths being punished, i.e., marked up to compensate for the loss of accuracy since lower sample sizes result in less precise estimates. In comparisons of variances of different investment lengths, this should be kept in mind. The following equations summarize the calculations performed to arrive at the subsamples of returns associated with these investment horizons.

$$h \in H = List\ of\ relevant\ investment\ horizons$$

$$r_t = \frac{p_{t+h} - p_t}{p_t},$$

$$r_{t+1} = \frac{p_{t+1+h} - p_{t+1}}{p_{t+1}}$$

$$\vdots$$

$$r_{T-h} = \frac{p_T - p_{T-h}}{p_{T-h}}$$

Returns obtained from different investment horizons would not be directly comparable without converting them to their equivalent rates corresponding to a base time length. This is usually done by finding the annualized return rates for comparable investment lengths; however, the investment horizons set contains mere days in addition to decades. Annualizing a 30-year return value means discounting it by taking its 30^{th} root whereas annualizing a 1-day return would mean compounding it by taking its 365^{th} power. This unnecessarily distorts the graphs, especially when reporting the extreme values such as the minimum and maximum observations of short-term investments. It is also misleading since it yields a value that could only be obtained had daily returns could be consecutively made at the same rate. This is unlikely to occur, especially for the more extreme values. On the other hand, discounting a long run return simply yields the daily equivalent of a return that actually occurred in the data.

The distortions resulting from annualizing returns obtained from investments lasting less than a year can be remedied by rescaling the graph by taking the logarithm of the values, but the resulting graph would not be immediately intuitive as it would require mentally deconstructing the values in a two-step process. Therefore, instead of using annualized values, daily equivalent returns are calculated, since the shortest investment length chosen is a single day. For each investment horizon h, the return is discounted by taking its h^{th} root.

$$r_{T-h}^{*} = \left(\frac{p_T - p_{T-h}}{p_{T-h}}\right)^{\frac{1}{h}}$$

$$\bar{r}_h = \frac{\sum_{t=1}^{T-h} r_t^{*}}{T-h}$$

$$\sigma_h^2 = \frac{\sum_{t=1}^{T-h} (r_t - \bar{r}_h)^2}{T-h-1}$$

Daily equivalent returns calculated in this manner are comparable to each other since they all correspond to a base investment length of a single day.

Once the associated rolling windows return samples are constructed for each relevant investment length, their sample variances can be calculated. The unconditional variance line can also be plotted across all investment lengths.

The probabilities of beating fixed interest rates can be estimated using the rolling windows samples constructed. Simply counting the times it occurred in the dataset would give us the sample probabilities. It is possible to plot these sample probabilities across investment horizons to visually examine their evolution.

$$SampleProb(r > i) = \frac{\sum_{t=1}^{T-h} \begin{cases} 1 \; if \; r_t^{*} > i \\ 0 \; if \; r_t^{*} \leq i \end{cases}}{T-h}$$

where $i \in I = \{List\ of\ relevant\ fixed\ interest\ rates\}$

Rolling windows by their nature consist of overlapping periods. Although the returns are calculated using the first and the last days of each window and the same data points are never used twice, they cover the same intervals which streamlines the data in the sense that it reduces the variety that would be faced by an investor in the long run. As a result, the unconditional variance is likely to be downward biased. This feature should be kept in mind that when comparing its results with those of predictive variance. Any glide paths constructed using the unconditional variance is likely to weigh stocks more than what would be appropriate for their risk levels.

3.3 Nominal returns findings

Findings obtained using the nominal returns are within expectations.

Figure 4 shows that the variance drops as the investment horizon lengthens while the Sharpe ratio increases. These findings support the conventional wisdom: long-term investments promise better returns with lower risk. Going from daily to weeklong investments decrease the risks nearly by 80%. Variance starts to approximate toward zero as the investment horizon lengthens.

Figure 5 plots a magnified version of Figure 4 to focus on the behavior of the Sharpe ratio since it displays explosive growth after 15 years. The Sharpe ratio climbs rapidly with the investment horizon in an exponential fashion. This is mainly due to the reduction in variance since the mean return is nearly constant within the 35-40% band, as can be seen from Figure 6.

It should be noted that the x-axis of the figures here does not represent time as usual but the relevant investment horizons picked; therefore, it is not to scale.

The decrease in unconditional variance only tells a portion of the picture. Variance is a measure of risk; however, it is a singular value not incapable of differentiating between upside and downside risk. This is important, because depending on which one dominates, the investor will be more likely to suffer losses or enjoy gains. In order to compare the upside and downside risks, a moving boxplot can be constructed to have a better understanding of how the distribution of the actual mass of returns behave as investment period lengthens.

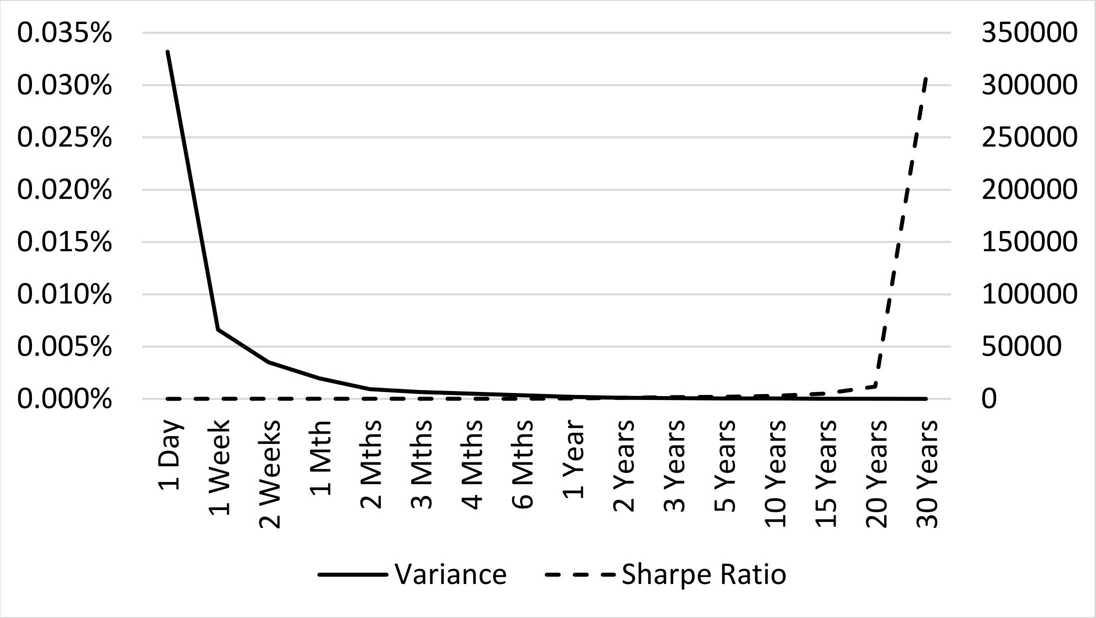

Figure 4. Variance and the Sharpe Ratio (nominal)

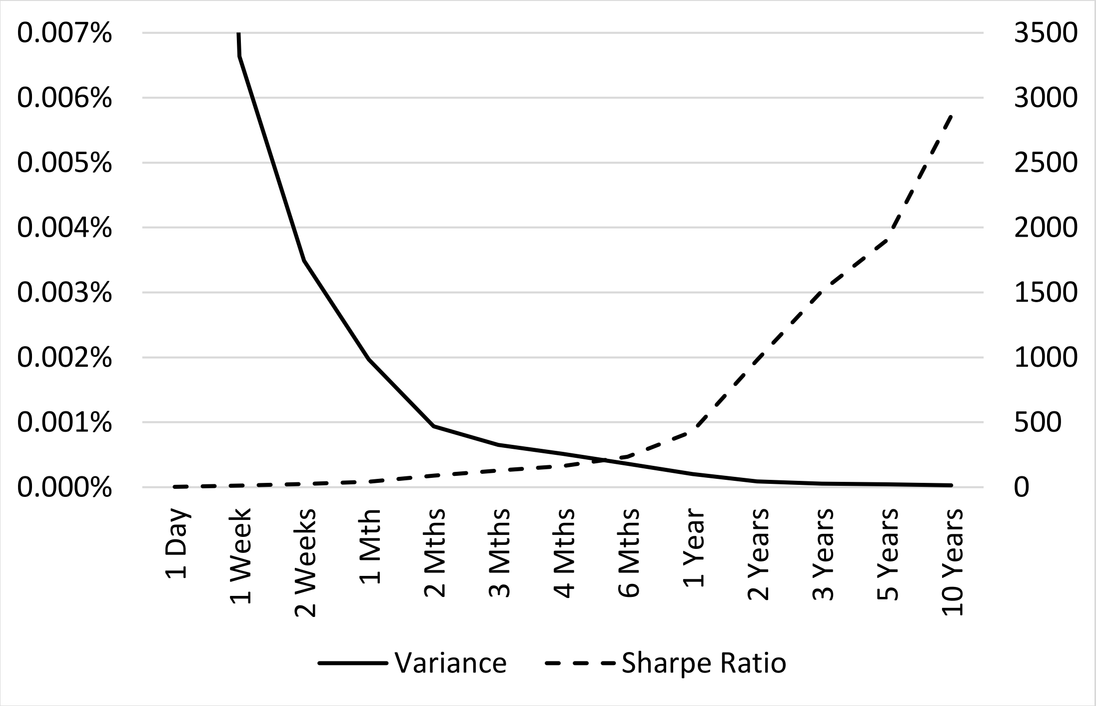

Figure 5. Variance and the Sharpe Ratio (nominal, magnified)

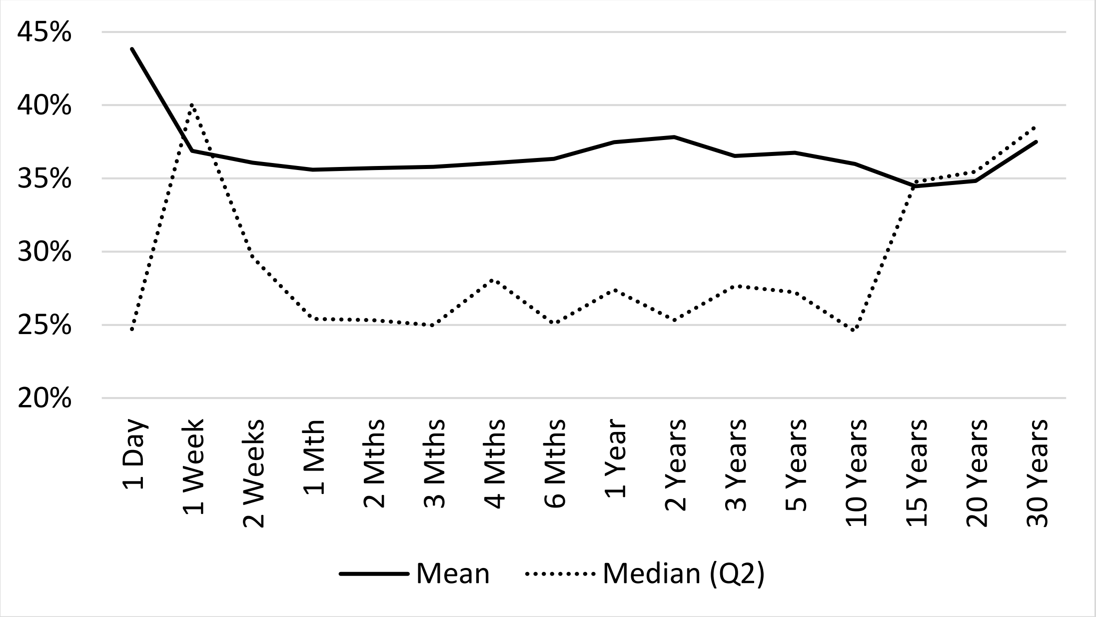

Figure 6. Annualized mean and median returns (nominal)

Each boxplot consists of the mode, mean, median, first and the third quartiles, minimum and maximum values calculated from each subsample. The only difference is that instead of a single stem-and-leaf display, each of these values are plotted as lines depicting their evolution through time. Plotting them in a similar way to unconditional variance yields the moving boxplot graph plotted in Figure 7. This figure highlights a few important properties. First, it confirms what the unconditional variance graphs tell us: the band of return values narrows as investment period lengthens. Second, zooming in on the interquartile range shows that the mean and median return values are fairly stable across investment horizons. This is a convenient result from a portfolio choice perspective, since with comparable mean returns, variance becomes the sole determining factor of an investor’s preferences in regard to such investments. If it were not the case, a risk-return preference-based model would have to be used to figure out the premiums that would be required to justify a certain amount of risk. Since the mean return is nearly identical across all horizons, it can be broadly argued that longer investment horizons are more preferable since the variance is decreasing. Figure 8 contains a magnified version of Figure 7 for easier examination.

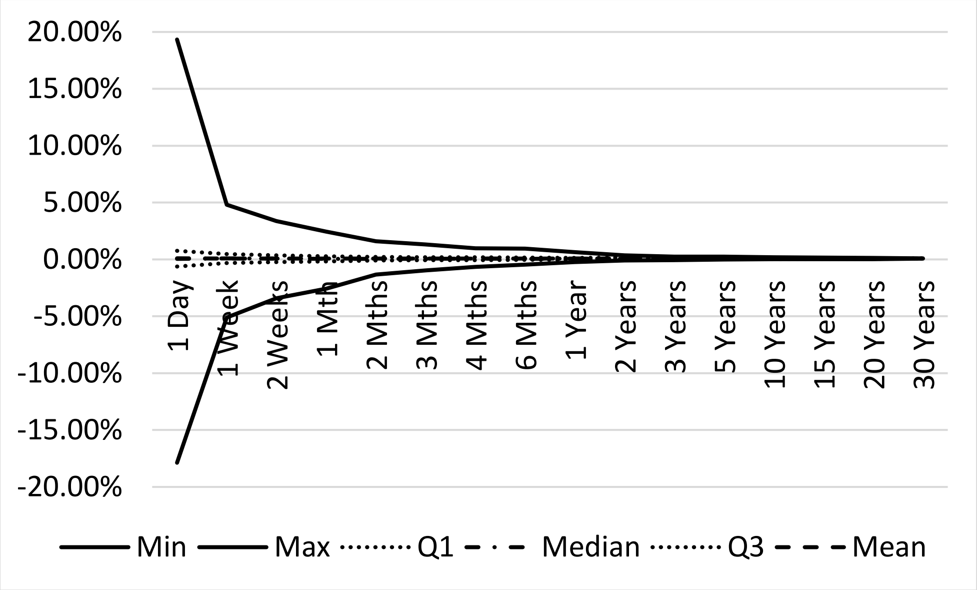

Figure 7. Moving boxplot of daily equivalent returns (nominal)

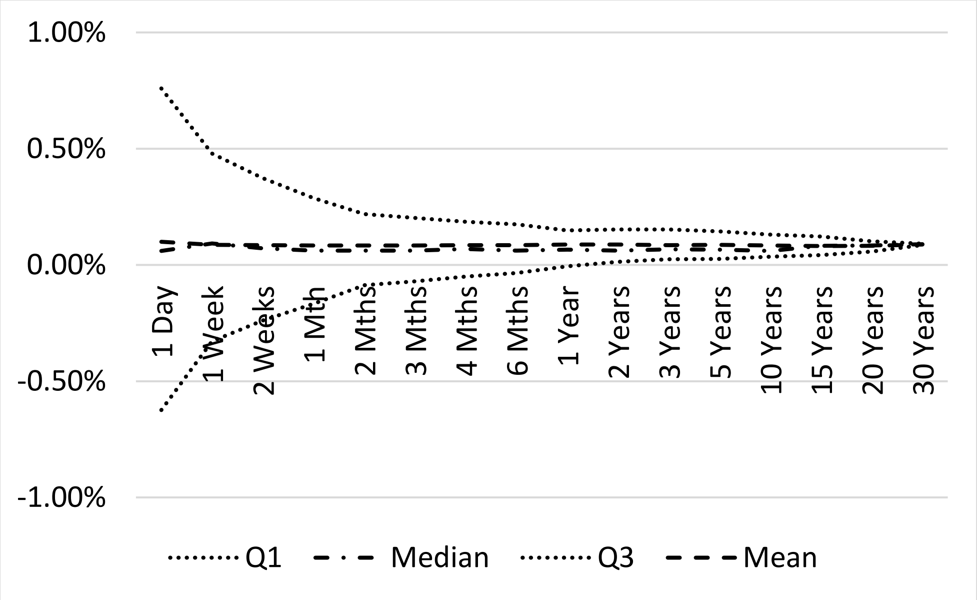

Figure 8. Moving interquartile range of daily equivalent returns (nominal)

These results are very promising for an investor with long investment horizons, but one of the main counterarguments to decreasing long-run variance should still be considered: the value at risk becoming greater than what could be compensated by the reduction in risks. In order to determine the likelihood of an investor losing money or being below a target return rate, the probabilities of such occurrences can be estimated based on the sample probabilities. In order to measure such risks, the data can be combed, and the times investments failed to beat a certain fixed income goal can be counted. This can be done for each investment length and the results can be plotted across the investment horizons. This allows constructing the sample probabilities of each investment horizon surpassing a certain fixed income target return as in Figure 9.

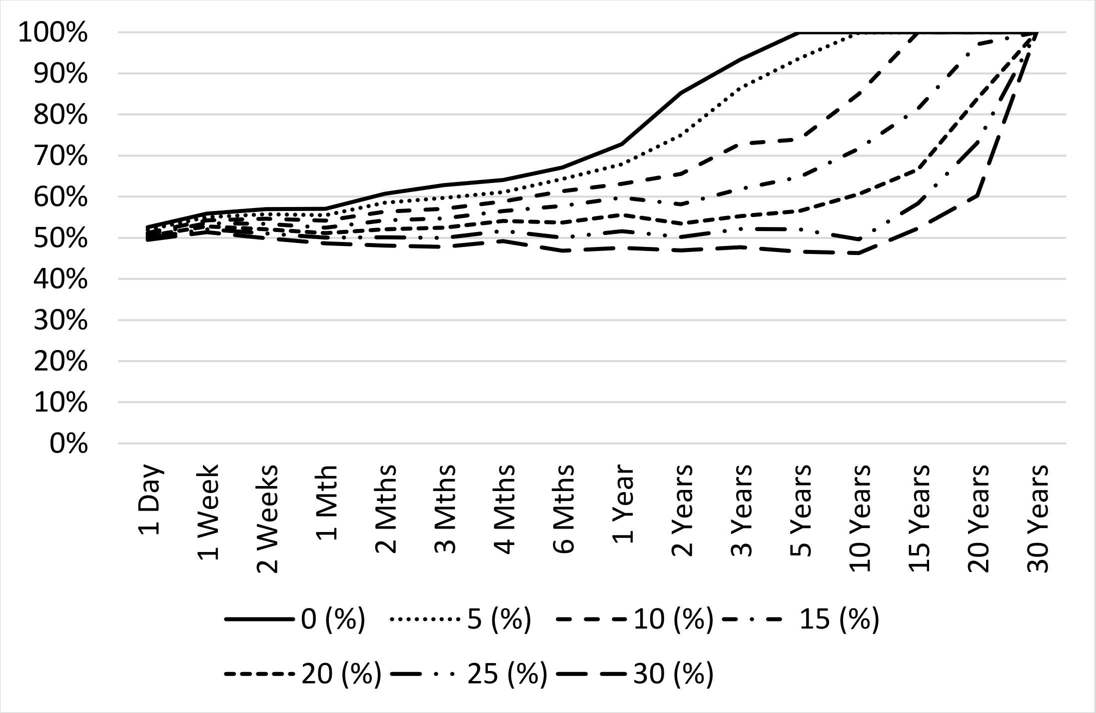

Figure 9. Probabilities of beating fixed interest rates (nominal)

This figure shows that the chances of an investor making money or beating any of the selected interest rates starts off around 50-50; however, it becomes 100% as the investment horizon lengthens. This applies to larger return rates with more of a delay as would be expected. Nevertheless, it becomes nearly impossible to fall below target rates with a long enough investment horizon. Financial downturns may become more likely in the long run; however, their relative length compared to that of the entire investment horizon becomes short, allowing the investor to reach their goals.

The results from nominal returns so far seem to support the conventional wisdom. For a more complete picture, the real returns are considered in the next sections.

3.4 Real returns findings

This part of the analysis is essentially the same except for the fact that the underlying data used consists of real returns derived from the nominal data using a daily consumer price deflator.

Looking at Figure 10 and Figure 11, a similar picture can be seen to the nominal returns case for the unconditional variance and the Sharpe ratio. Figure 12 shows that the mean returns are relatively constant here as well, hence the main driver of the explosive increase in the Sharpe ratio is once again the reduction in variance over long investment horizons.

Figure 13 and Figure 14 also paint a similar picture to their nominal counterparts. The band of possible return values narrows while the mean and median returns remain relatively constant.

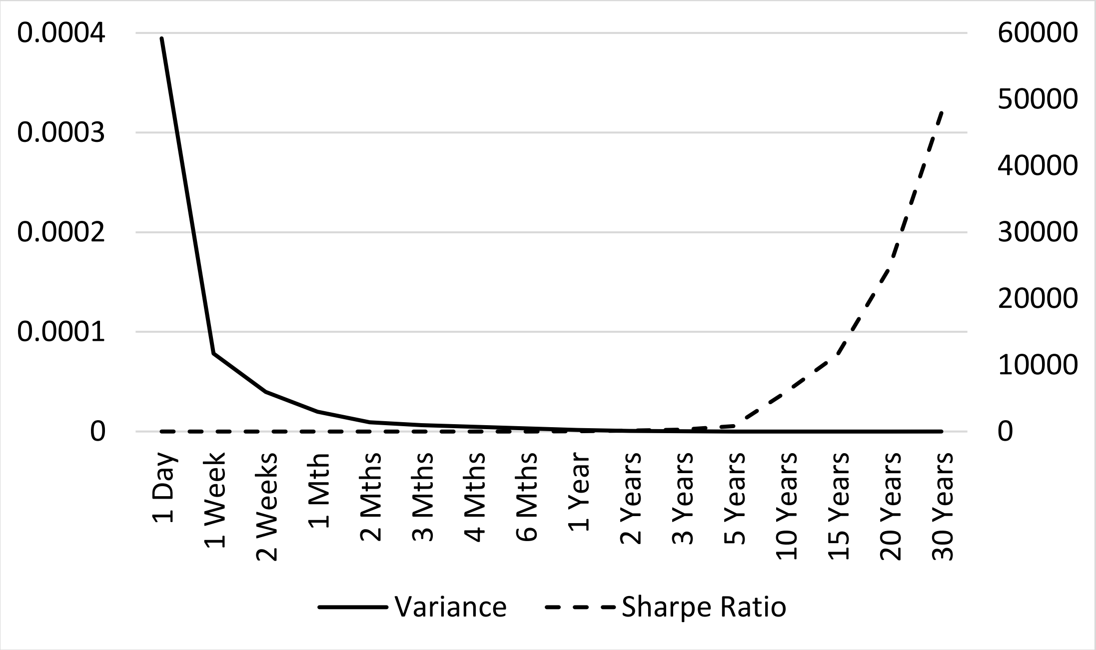

Figure 10. Variance and Sharpe Ratio (real)

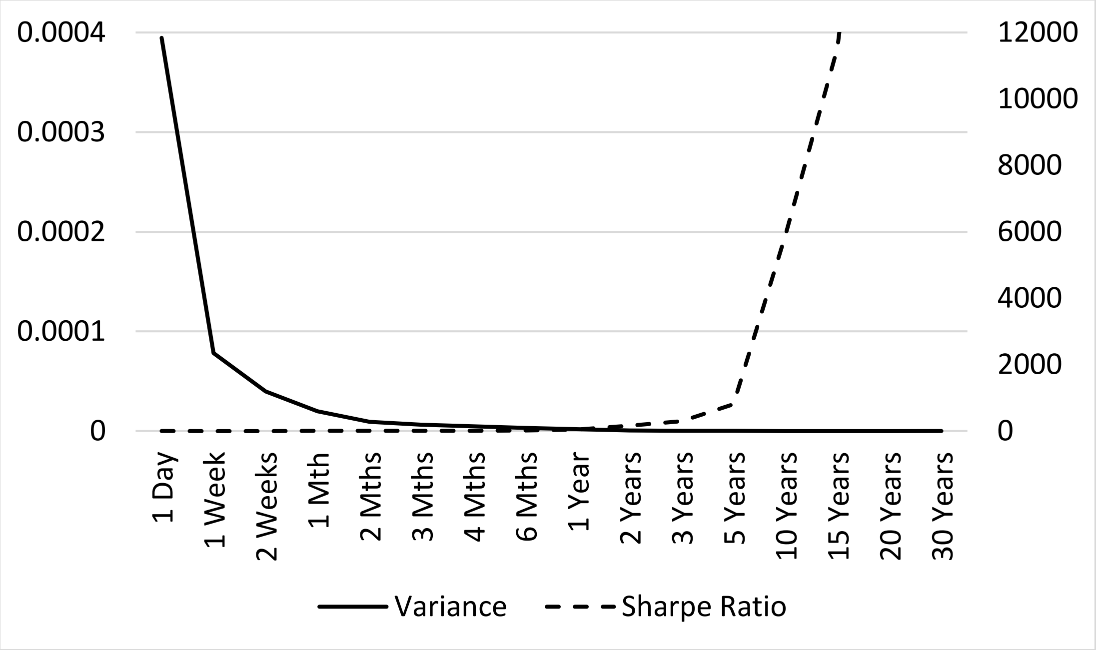

Figure 11. Variance and Sharpe Ratio (real, magnified)

The main difference between the nominal and real return results lies in the sample probabilities calculated. Figure 15 shows that up until the 5-year investment horizon, the chances of beating any of the selected target rates is basically 50-50. This is very different than the behavior displayed by nominal returns where each passing year adds to the probability of achieving a target return rate. The chances of beating one and two percent fixed returns improves past the five-year mark to reach 87% and 75% levels, respectively. The chances of beating a 3% fixed interest rate only starts improving after the 20-year mark and reaches approximately 55% at 30 years.

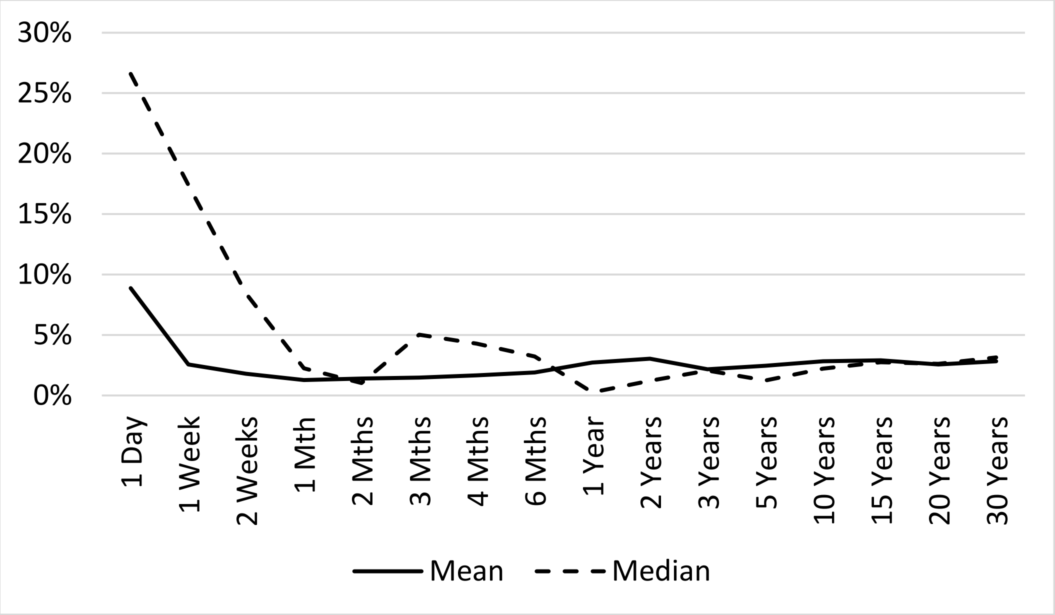

Figure 12. Annualized mean and median returns (real)

These results suggest that short-run investments are pretty much a coin toss where investors would miss their target half the time. After a few decades, this behavior starts to change: it becomes less likely to earn a high real return rate, but the minimum rate of return guaranteed increases. In short, it becomes less of a gamble and more of an investment in terms of real returns. At the end of the sample at a 30-year investment, investors are able to surpass a real annual return of 3.3% half the time.

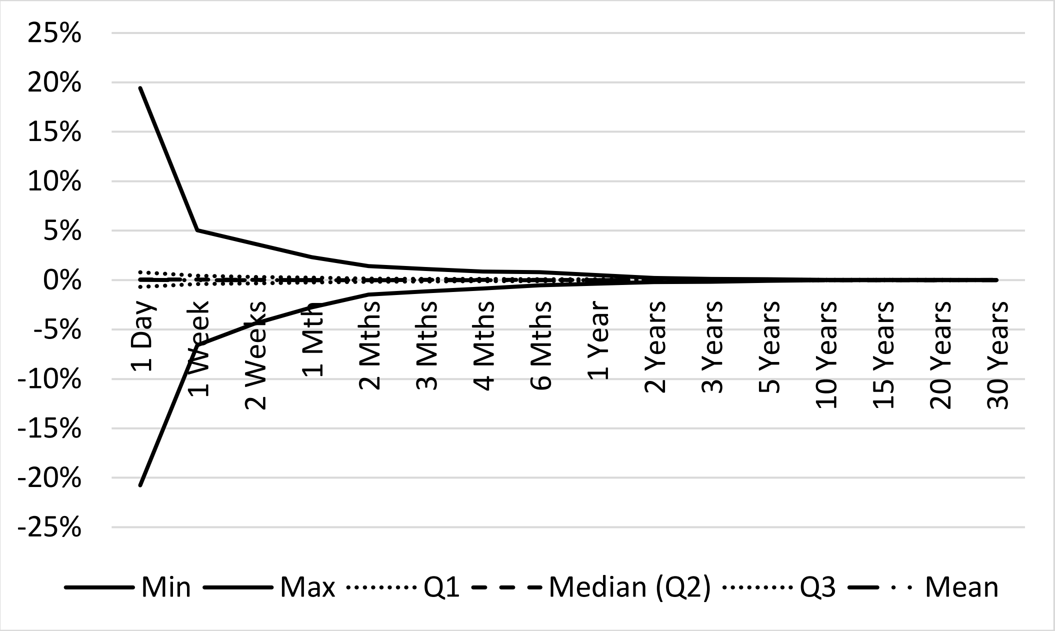

Figure 13. Moving boxplot of daily equivalent returns (real)

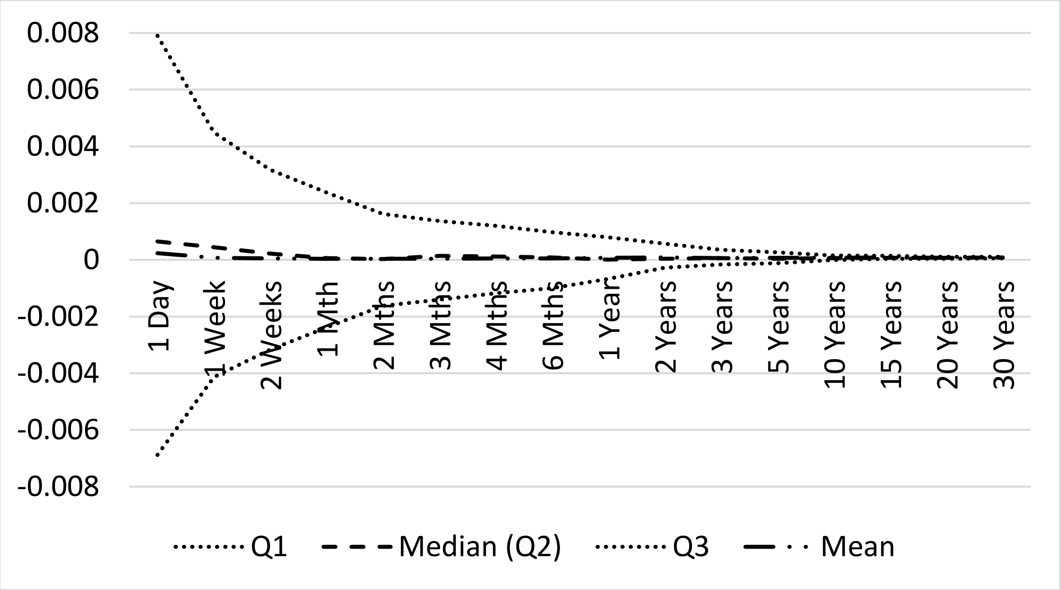

Figure 14. Moving interquartile range of daily equivalent returns (real)

In comparison, their probability of not losing money is 50% for investments horizons of one year and less; however, their chances of extreme loss and gain is also higher than a longer investment horizon.

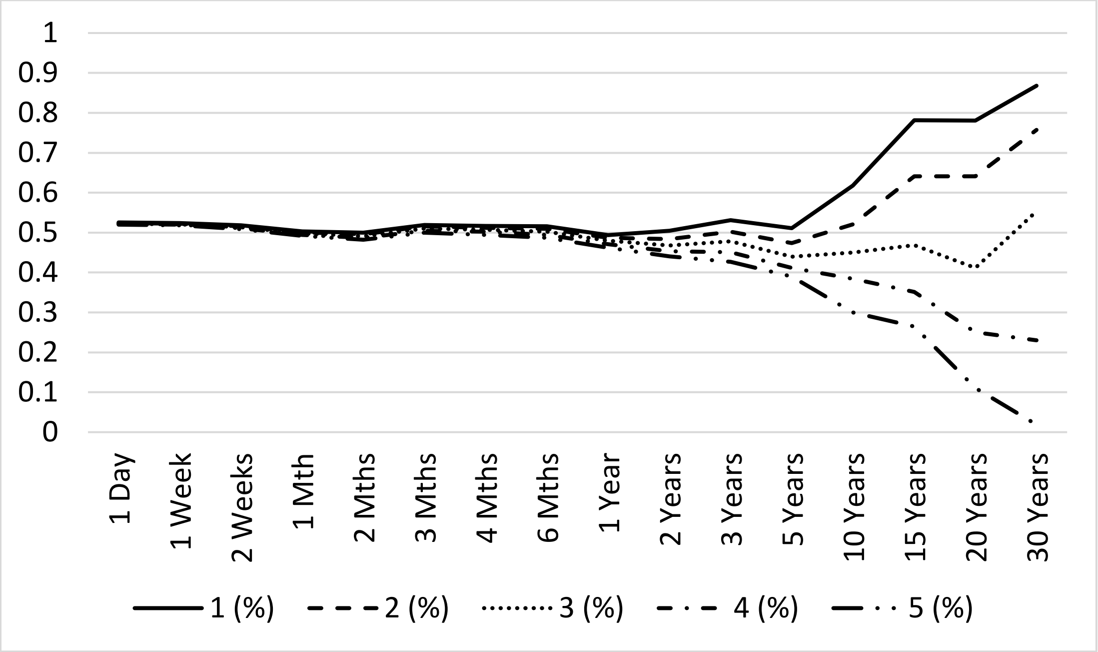

Figure 15. Sample probabilities of beating fixed interest rates (real)

3.5 Conclusions for the unconditional analysis

The results obtained here generally support longer investments being less risky than shorter ones. In addition, the mean returns are relatively constant and positive across investment horizons; therefore, investors can base their decisions almost entirely on variance regardless of their risk preference. It becomes more likely to beat nominal target return rates in the long run while the minimum real return guaranteed increases.

The analysis here is unconditional, i.e., the results are essentially summary statistics that are not conditioned upon any variables. Therefore, they are not fully guaranteed to remain the same in the upcoming decades. They should be utilized in future decision-making to the extent investors believe the underlying conditions driving this behavior will continue towards the future.

CHAPTER 4

PREDICTIVE VARIANCE ANALYSIS

As an answer to the aforementioned shortcomings of unconditional and conditional methods, predictive variance methods have been proposed. Predictive variance is the variance obtained from a predictive system which is a Vector Autoregression (VAR) model tying the dynamics of returns to those of expected return and predictors. Some models are known to omit predictors from the system if the dynamics of return and expected return can adequately capture the dynamics.

4.1 Data

The data used in the predictive variance analysis consists of the annual closing prices as plotted in Figure 16, the consumer price index deflator used in the previous chapter, and a set of predictors. The predictors include gross domestic product growth rate, deposit rate, and the foreign exchange rate (U.S. Dollar to New Turkish Lira).

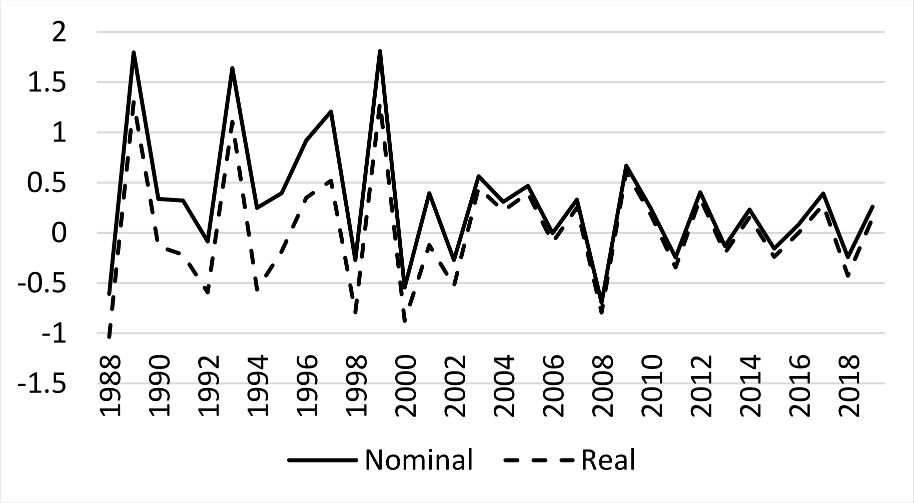

Figure 16. Log-return values (XUTUM index)

4.2 The analytical model

The standard vector auto-regression (VAR) model used in estimating the predictive variance is constructed below. This model ties the current return to the expected return and a set of predictors.

$$r_{t+1} = \mu_t + u_{t+1}$$

$$\mu_{t+1} = \alpha + \beta\mu_t + w_{t+1}$$

$$x_{t+1} = \theta + Ax_t + v_{t+1}$$

$$\text{where } \beta \in (0,1) \text{ and } \rho_{uw} < 0$$

where r_{t+1} is the next period's return, μ_{t+1} is the expected return, x_{t+1} is the k-dimension predictor vector containing the variables $x_{1,t+1}, x_{2,t+1}, \cdots, x_{K,t+1}$. α and β are the coefficients of the expected return equation. θ and A are the coefficient vectors for the predictors. u_{t+1}, w_{t+1}, and v_{t+1} are the error terms of return, expected return, and predictors, respectively.

The connection between the predictors and the other variables are not immediately apparent since their equations do not share any x variables. They are related to each other through the variance-covariance matrix of the VAR system which will be called the Σ matrix.

$$\Sigma = \begin{bmatrix} \sigma_u^2 & \sigma_{u,w} & \sigma_{u,v_1} & \cdots & \sigma_{u,v_K} \\ \sigma_{w,u} & \sigma_w^2 & \sigma_{w,v_1} & \cdots & \sigma_{w,v_K} \\ \sigma_{v,u} & \sigma_{v_1,w} & \sigma_{v_1}^2 & \cdots & \sigma_{v_1,v_K} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \sigma_{v_K,u} & \sigma_{v_K,w} & \sigma_{v_K,v_1} & \cdots & \sigma_{v_K}^2 \end{bmatrix}$$

The two key assumptions here are: *i)* expected returns are stationary with $\beta \in (0,1)$, and *ii)* the correlation between the error terms of current and expected return equations is negative: $\rho_{uw} = \frac{\sigma_{u,w}}{\sigma_u \sigma_w} < 0$. They reflect the stylized facts in the literature as noted by Cochrane (2009). The first assumption guarantees that there is momentum in expected returns, i.e., the changes in return are likely to keep their direction for a certain duration depending on how large the value of β is. The second assumption suggests that there is mean reversion, i.e., the changes in return rates are pulled back toward the mean value depending on how negative ρ_{uw} is. In combination, these effects result in a system where the long-run variance is increased by momentum and decreased by mean-reversion. The behavior of predictive variance across investment horizons heavily depends on which effect dominates in the long run. The assumptions of momentum and mean reversion are part of the stylized facts

In application, these two assumptions create a very strong dynamic that accounts for the behavior of predictive variance to such an extent that the effects of the predictors become negligible and they end up being dropped from the analysis. In such cases, the predictor equations can either be dropped from the system and a two-by-two Σ matrix is used, or they can be cancelled out by modifying the original five-by-five Σ. The latter is more practical, since it allows utilizing the same computer code for both applications.

Although the variables defined through the VAR equations and the covariance matrix can sufficiently present the dynamics of the model, intermediary and auxiliary variables must also be defined for the algorithm to utilize throughout the many iterations of the Bayesian MCMC estimation which will be more comprehensively covered in the code implementation section.

The condition for an investor to prefer long term investments is the per period variance for longer horizons being smaller than the single period variance. This is summarized in the following inequality which will be referred to as “the long-run preferability condition”.

$$\frac{Var\left(r_{t,t+k}|D_t\right)}{k} < Var\left(r_{t,t+1}|D_t\right)$$

Since long-term returns are higher, this equation by itself would almost guarantee that the Sharpe ratio is always increasing. The Bayesian MCMC analysis carried out here is essentially an analysis of the behavior of this key statistic under various investor beliefs and assumptions reflected through the predictive models and joint prior distributions imposed on them.

The per period predictive variance is calculated by obtaining its posterior values through thousands of iterations and then simply integrating over them conditional upon all the other variables. This variance can be decomposed into five separate components following the methodology proposed by Pastor and Stambaugh (2012) in order to determine the main drivers of predictive variance over different investment horizons. These five components are i) i.i.d. uncertainty, ii) mean reversion, iii) uncertainty about the future expected returns, iv) uncertainty about the current expected returns, v) estimation risk. Out of all the five, only mean reversion is expected to have a negative value, i.e., reducing the overall variance. Mean reversion is mainly driven by the ρ_{uw} parameter which denotes the negative correlation between the error terms of return and expected return equations. Hence, of all the parameters in the model, ρ_{uw} and β together have the greatest influence on the final result. The five components of variance can be calculated for each investment horizon and then plotted for a comparative analysis.

4.3 The Bayesian MCMC algorithm

This part summarizes the algorithm steps performed during a single run of the analysis. A more detailed version is provided in Chapter 5.

4.3.1 Algorithm steps

Bayesian estimation methods are often perceived to have a shroud of mystery to them, to the point of being described as “black box” methods. In order to dispel any such concerns, this section explains the steps of the Bayesian MCMC methodology employed here in detail. The software implementation contains numerous additional variables and detailed steps; therefore, it is covered in its own section in Chapter 5 to avoid unnecessary digression.

The Bayesian analysis is a numerical estimation method that starts with researcher specified (or derived) prior distributions for the parameters in the system. It updates these distributions by iterating the model many times to produce the posterior distributions which in turn can be used to arrive at the estimates for the unknown parameters. Random draws for the necessary parameters are performed at each iteration step from their own distributions. Since these draws would yield different values at each run of the estimation, the numerical results would not precisely be the same; however, they should be quantitatively close enough and qualitatively in the same vein. Inconsistent results between different runs of the algorithm would signal problems in estimation and would require further investigation.

Due to its temporal nature, each iteration of the algorithm consists of two parts: forward iteration and backwards sampling. Forward iteration uses the current values of the parameters and variables to calculate their new values for the next period. It runs until the final period is reached and stops. Backwards sampling picks up from where the forward iteration left and draws an error term for the final period from the updated distributions. It iterates backwards and creates a series of expected returns until it reaches the first period. Taking the difference between these estimated values and the observed ones gives a set of residuals. These residuals are used in updating the parameters for the next step of the iteration. The next iteration carries out the same calculations with the updated values.

The very first iteration uses the initial values set by the researcher or values obtained from prior distributions depending on the nature of the parameter. In an ideal setting, the initial values would not matter after countless iterations since the distributions would approximate their true forms which are singular. In practice, however, limited iterations are run; therefore, the initial values matter to a certain extent. Even in the case of limitless iterations, problems may occur in calculating the ratio of very small parameters having distribution spreads greater than their magnitude. In such cases, prior elicitation has a crucial role and the initial values must be justified to ensure unbiased results.

Some parameters can be expressed analytically as a function of other variables; therefore, their initial value would depend on the initial values of the variables of this function. For parameters that can be modeled as unconditional or conditional statistics, the initial values can be substituted by estimating them through conventional methods. For example, an AR(1) analysis can be run to find the α and β coefficients of the VAR system of individual equations and use them as their initial values. Other parameters can be approximated by running the analysis for a couple of rounds and calibrating the results in accordance with the difference between the prior and posterior distributions. The rest of the parameters cannot be outright determined using the existing values; therefore, they can be given different values to illustrate different scenarios. This is what is referred to as imposing beliefs on the model in Bayesian analysis.

The covariance matrix in the time-varying Σ case cannot be determined in a straightforward manner. A Cholesky decomposition is performed to transform it into a linear system involving standard normal distributions. The steps followed here are the same as Carvalho et al (2018) which are based on the methods developed by Daniels and Pourahmadi (2002), Chib, Nardari, and Shephard (2002), Jacquier, Polson, and Rossi (2002), Prado, Ferreira and West (2010), and West and Harrison (1997).

The final results can be achieved through the posterior distributions. These distributions are constructed by recording the parameter values at each iteration step. The predictive variance estimate is found by integrating over its posterior distribution by conditioning it on other variables.

4.3.2 Assessing the goodness of fit

Bayesian MCMC does not have a shorthand statistic such as R-Squared to summarize how well the model fits the data. Numerical estimation methods with random draws yield different estimates after each run of the algorithm. Depending on the nature of the data and the models used, the outcomes may be sensitive to the changes in the priors of a few key parameters and the degree of this sensitivity may not be directly measured. To make matters worse, the statistic of interest is the variance of stock returns which is not directly observable in the data as a time series.

In order to assess the goodness of fit in Bayesian models, posterior predictive checks have been proposed. This method simulates samples based on random draws from the posterior and prior distributions of the Bayesian estimation and checks the degree of similarity between the behavior of simulated draws and the actual data. It is similar to an F-Test for the model where the behavior of the actual data is expected to fall inside a certain confidence interval and reject the model if it exceeds these limits significantly. Rubin (1984) and Gelman, Meng, and Stern (1996) can be reviewed for further details on the practice.

The actual data equivalent of predictive variance is unconditional variance. The test performed here is in the same spirit as Carvalho et al. (2018). The steps are provided below.

i. Simulate expected return and return series thousands of times by randomly drawing the coefficients and covariance matrices based on their prior and posterior distributions.
ii. Calculate the unconditional variances of each simulated series.
iii. Construct the confidence intervals based on the simulated unconditional variances.
iv. If the actual unconditional variance line falls outside the confidence intervals based on priors, then the model would be invalid since it fails to replicate the properties observed in the actual data.
v. The more the actual unconditional variance line is contained by the confidence intervals based on posteriors, the better the fit.
vi. Since the variable of interest is the predictive variance rather than the return values themselves, it is sufficient that the behaviors of the simulated and actual unconditional variances are alike: the actual return values need not overlap.

4.4 Nominal returns findings

The analysis is run under three beliefs regarding the system:

i. Weak Negative Error Term Correlation: $\rho = -0.5$
ii. Strong Negative Error Term Correlation: $\rho = -0.9$
iii. Time-varying Σ matrix: Σ_t, $\rho = -0.9$

Predictive variance analysis results using the nominal return data are plotted in Figure 17 and they suggest that long-run investments are more preferable for stocks in Borsa Istanbul. The long run preferability condition is not violated under any of the beliefs and stronger negative error correlation results in lower risk.

It should be kept in mind that the lowest point of each line does not constitute an optimal point. The comparison condition requires comparing the predictive variances of single-period investments against multi-period ones. The predictive variance of a single period investment is the first value on the horizontal axis for each scenario since it corresponds to a one-year investment. Therefore, if they are extended as horizontal lines, the desirability of long-term investments can be determined based on whether they remain under this line throughout the investment horizon or not. The multi-period investments would retain their preferability over single-period ones until they cross above the straight lines. The returns are increasing in the long run; therefore, the optimal point would depend on investor preferences and it may be beyond the horizon constructed here.

Extending the trajectories of all three cases beyond the available horizon, it can be seen that the weak ρ_{uw} case would be the quickest to cross its initial value. The predictive variance for the time-varying Σ case starts at the bottom, but it quickly crosses the strong ρ_{uw} line.

The predictive variance for any case would eventually take an upward turn and cross its horizontal initial value line, highlighting the increased estimation risk faced by the investors when forecasting for longer periods. It should be kept in mind that the point where this reversal happens is still not the optimal point as variance is not the only determining factor. The expected value of returns should also be accounted for since it increases with the investment duration.

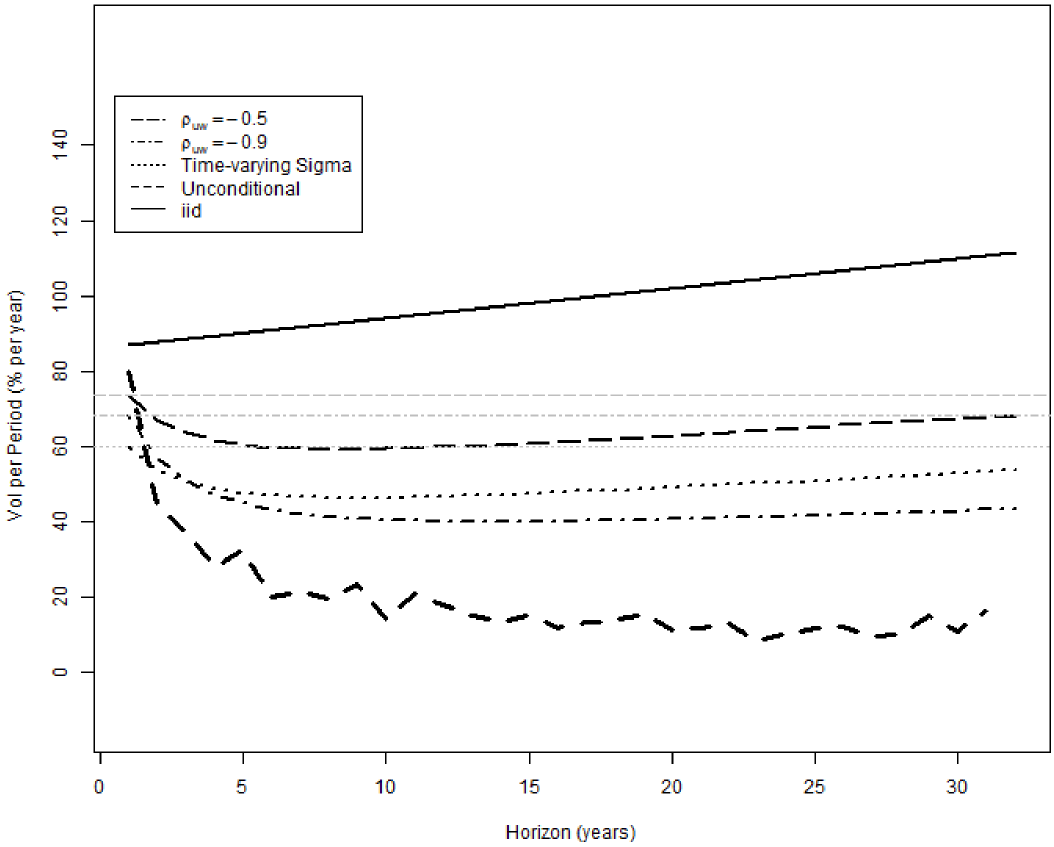

Figure 17. Predictive variance estimations, nominal

Combining these features, the graph can be divided into two regions for each scenario by drawing a vertical line at the dipping point of predictive variance. On the left side, the variance and expected returns would be both increasing, making investments preferrable without contest. On the right side, however, the variance starts to slowly increase while expected returns continue to increase; therefore, the extent of long-term investment preferability would depend on the risk preferences of individual investors.

The components of predictive variance for the weak and strong negative correlation cases are displayed in Figure 18 and Figure 19. They are plotted in the same manner as the predictive variance, displaying the behavior of each component over different investment horizons for each of the scenarios. Figure 20 shows the same for the time-varying Σ case.

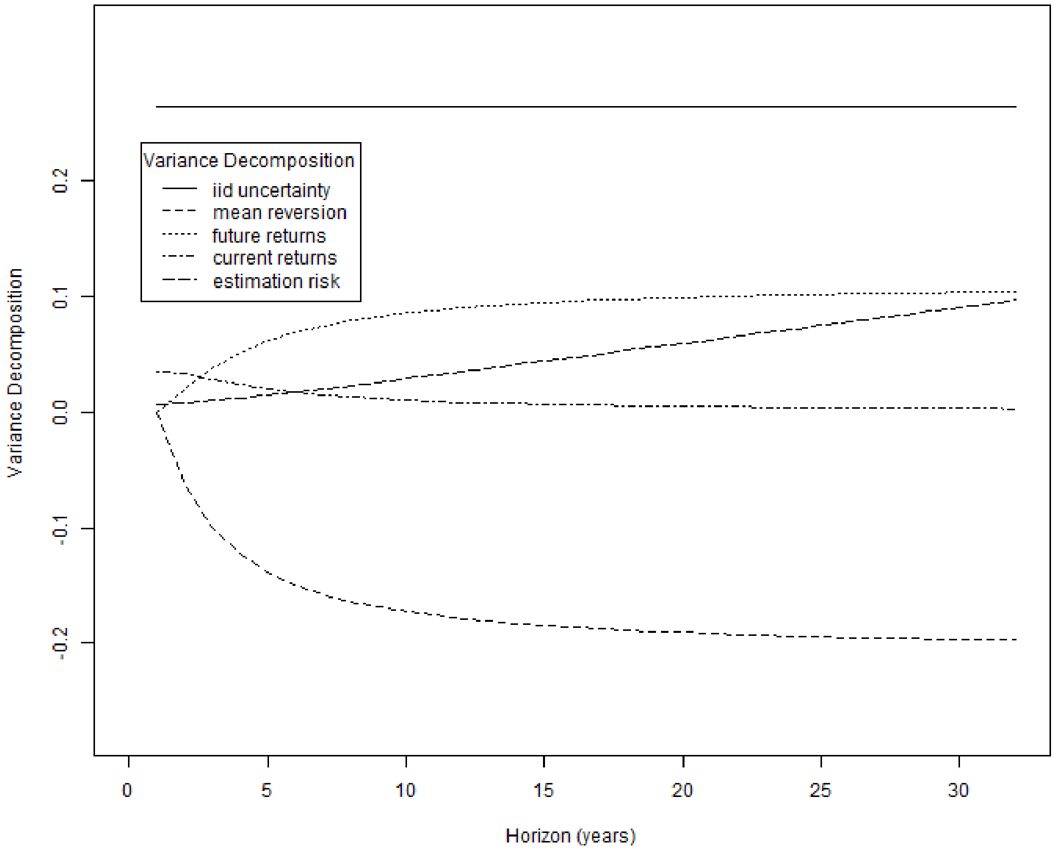

Figure 18. Components of predictive variance for $\rho_{uw} = -0.5$, nominal

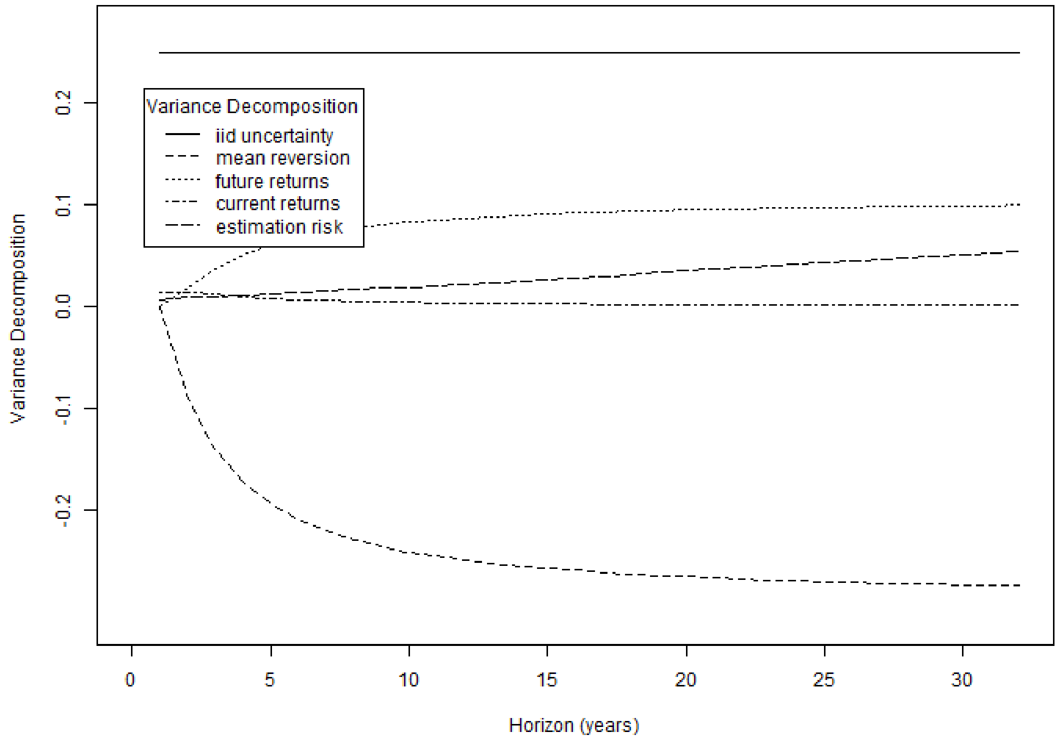

Figure 19. Components of predictive variance for $\rho_{uw} = -0.9$, nominal

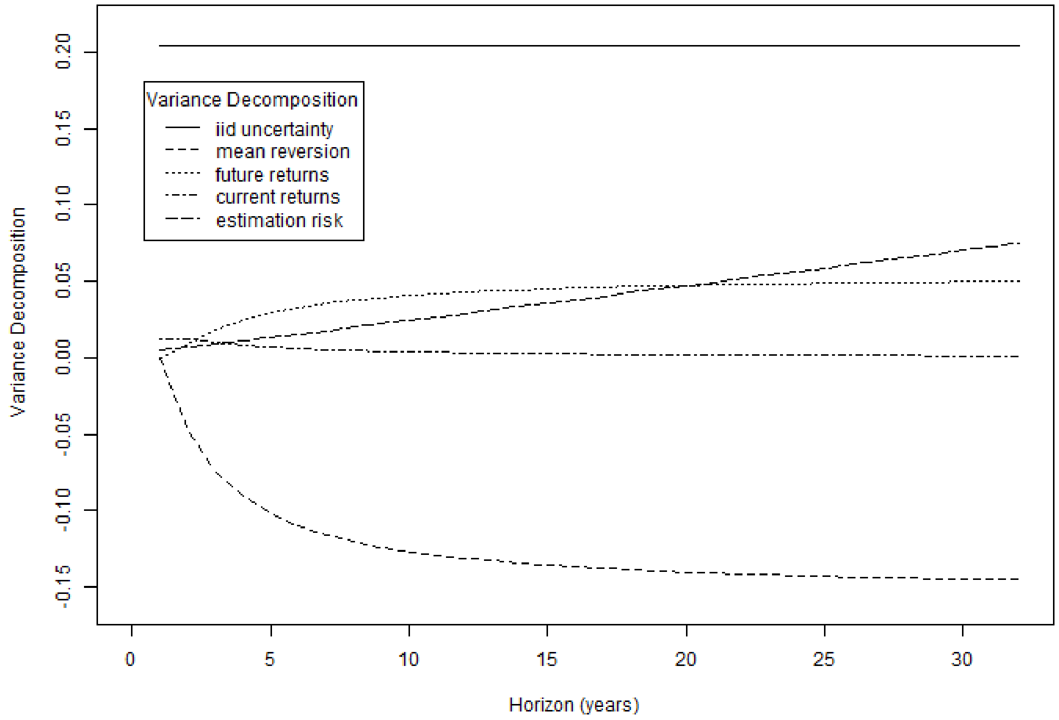

Figure 20. Components of predictive variance, time-varying Σ, nominal

The results obtained from the variance decomposition are as expected. There is a constant level of risk stemming from the i.i.d. component of the predictive variance. Mean reversion is the only negative component and its qualitative behavior is the same under all three scenarios. The only difference is that its negative magnitude is greater in the case of strong negative correlation belief which is a direct result of this assumption. Estimation risk linearly increases as time passes as it becomes more difficult for the investor to forecast the future. The risk regarding the future returns climbs quickly in the beginning but it settles at a nearly constant level, suggesting that long-term returns follow a narrower band and therefore are easier to predict. This observation is in line with the results obtained from the unconditional analysis. Lastly, the risks in predicting the current returns start from a positive value and approaches zero as the investment horizon lengthens. It is a sensible result, since the uncertainties regarding the current return estimations would be expected to have the majority of their impact in the short run.

The variance decomposition graph of the time-varying Σ case deserves further attention. It has a greater negative correlation value than the weak ρ_{uw} case; however, its mean-reversion component is smaller. Its predictive variance line would be expected to be above the weak ρ_{uw} case if it were not for the reduction in i.i.d. uncertainty. This suggests that momentum and mean-reversion are not the only two factors at play in terms of determining the predictive variance: assumptions accounting for more modeling shortcomings and resulting in more precise estimates also play an important role.

The prior and posterior distributions of key parameters are plotted in Figures 21-25. Comparing the priors against posteriors, it would be observed that the constant coefficient of the expected return equation α has a near zero value. The posterior distributions for β under all assumptions concentrate around values lower than their initial counterparts. There is little reason to believe that β is close to one. In fact, in previous trials calibrating the results, the priors higher than 0.5 for β always resulted in much lower posterior distributions.

All posterior distributions are more narrowly dispersed than the priors. The values for the error term variances σ_u^2 and σ_w^2 are small but not near-zero as feared initially. The low-signal-to-noise ratio problem plaguing most of the Bayesian MCMC research is not as great a problem here as it had been for previous studies. As a result, the system is less likely to be sensitive to changes in the initial values than comparable studies. Section 4.7 investigates this sensitivity further and provides supporting evidence.

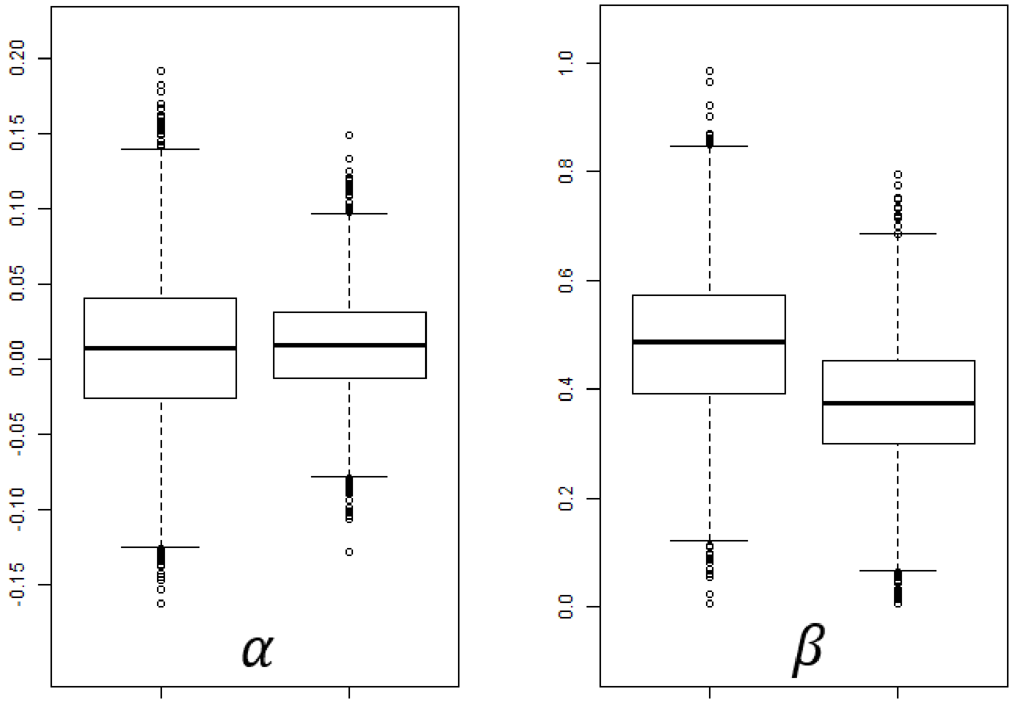

Figure 21. Priors and posteriors of μ coefficients, $\rho_{uw} = -0.5$, nominal

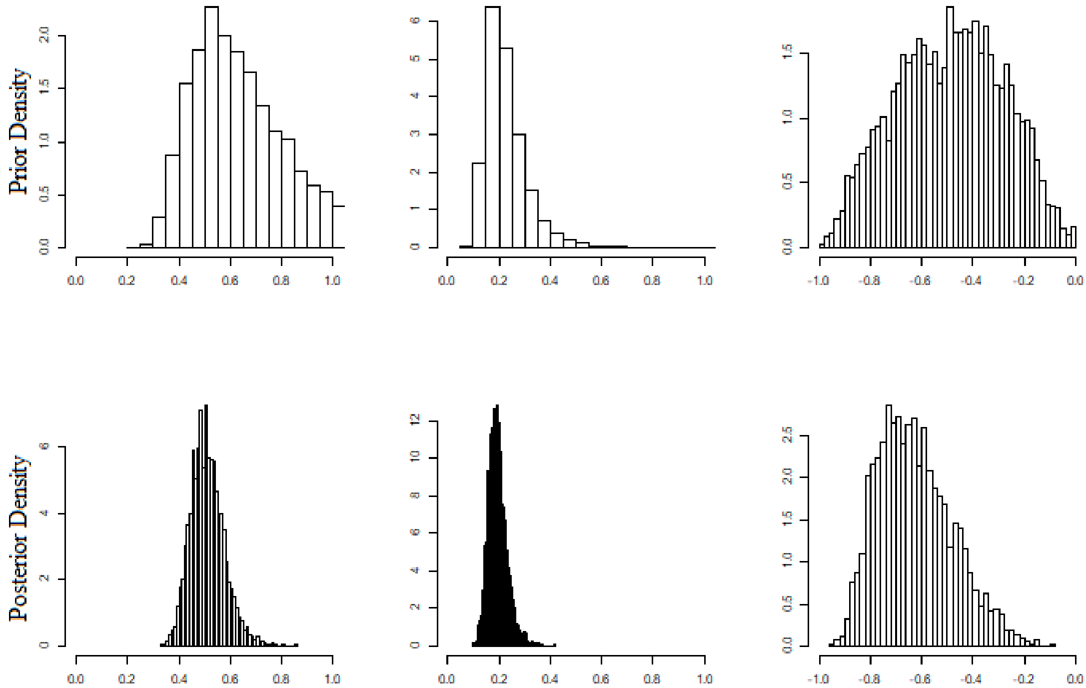

Figure 22. Priors and posteriors of σ_u^2, σ_w^2, and ρ_{uw}; $\rho_{uw} = -0.5$, nominal

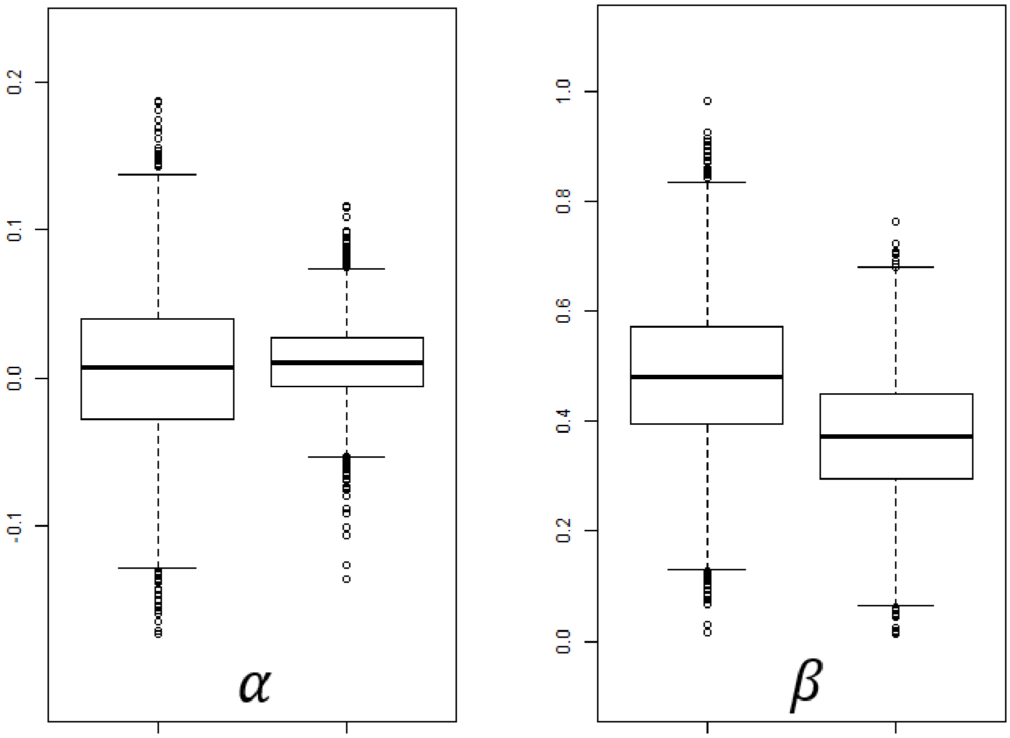

Figure 23. Priors and posteriors of μ coefficients, $\rho_{uw} = -0.9$, nominal

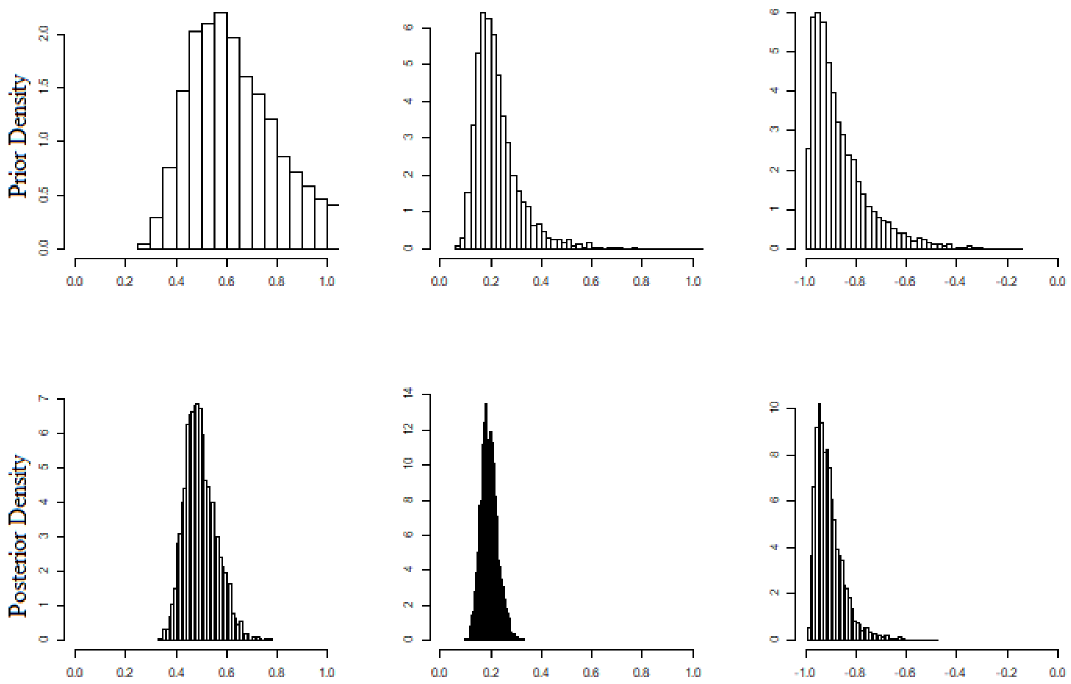

Figure 24. Priors and posteriors of σ_u^2, σ_w^2, and ρ_{uw}; $\rho_{uw} = -0.9$, nominal

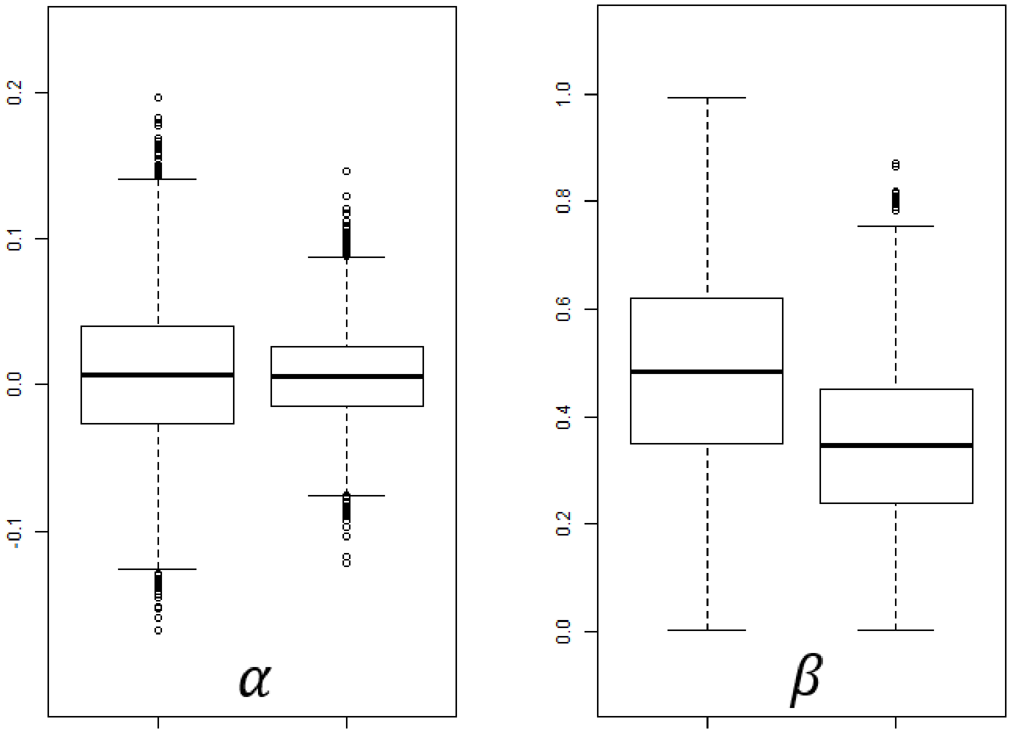

Figure 25. Priors and posteriors of μ coefficients, time-varying Σ, nominal

Goodness of fit for nominal returns

Two questions arise regarding the results obtained under the three scenarios: i) whether they fit the observed results well enough, and ii) whether one of them provides the best fit. As mentioned in Section 4.3.2; there is no shorthand statistic for this. The results regarding the goodness of fit are based on the visual examination of the simulated draws.

Simulating the return series for each case by drawing the parameter values from their respective prior and posterior distributions allows constructing the moving confidence interval bands across investment horizons. Following the reasoning in Section 4.3.2, whether the unconditional variance intervals obtained from the simulated draws contain the actual unconditional variance line is checked. Failing to meet this condition using the prior distributions would suggest the model is misspecified for the horizons the line falls beyond the intervals. After establishing model validity, the extent to which the lines are contained within the intervals constructed based on the posterior distributions can be checked. Figures 26-31 display the evolution of these confidence intervals over time for all cases.

The results obtained from both the prior and posterior draws of the weak ρ_{uw} case suggest that it fails to model the features of the dataset adequately. The unconditional variance falls outside the confidence intervals constructed using the prior distributions. In the posterior predictive checks, the confidence interval bands haphazardly contain the unconditional variance line. These two pieces of evidence suggest that the weak ρ_{uw} case falls short of explaining the data. It can be argued that the momentum in the system must be higher to adequately represent the data.

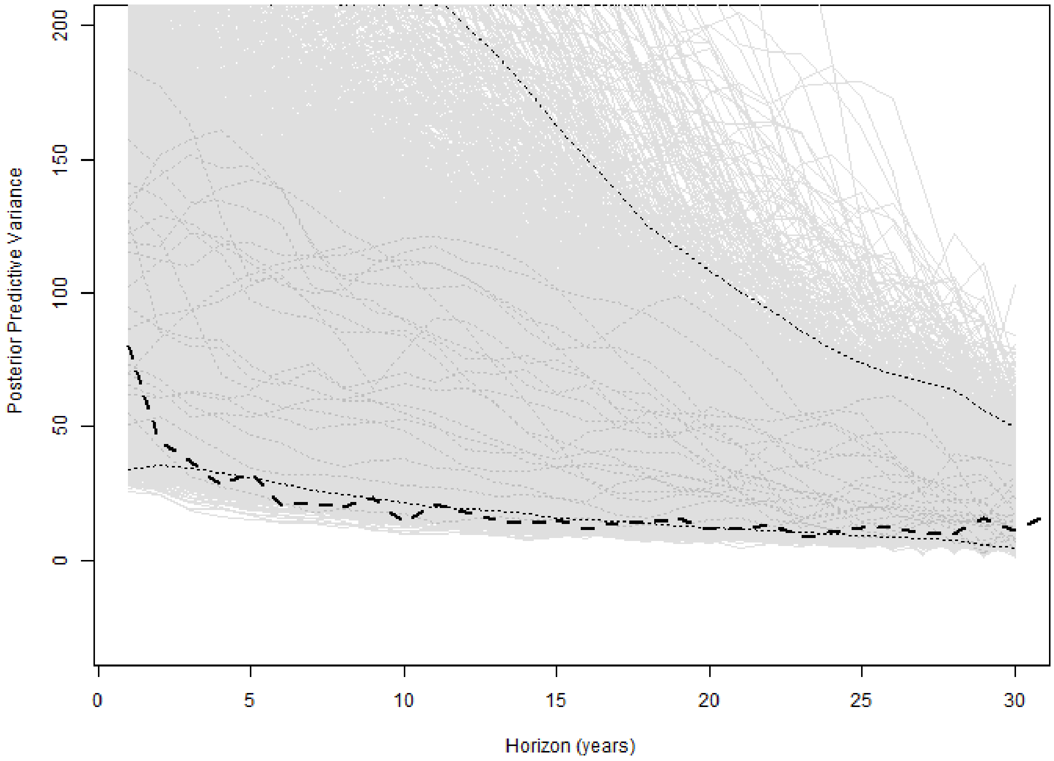

Figure 26. Prior-based confidence intervals, $\rho_{uw} = -0.5$, nominal

Turning to the strong ρ_{uw} case, it is observed that both the prior and posterior checks result in the unconditional variance line being contained within the confidence intervals, therefore the null hypothesis that the strong ρ_{uw} case adequately models the patterns observed in the data is not rejected. It should also be pointed out that out of all three scenarios, the strong ρ_{uw} case results in the narrowest confidence intervals based on the posterior distributions. In terms of prior distribution based confidence intervals, the time-varying Σ case yields the narrowest bands. The resulting shape is similar to the strong ρ_{uw} case; however, more points are left outside the confidence interval lines. This suggests that the strong ρ_{uw} case is more valid in terms of capturing the characteristics of the data.

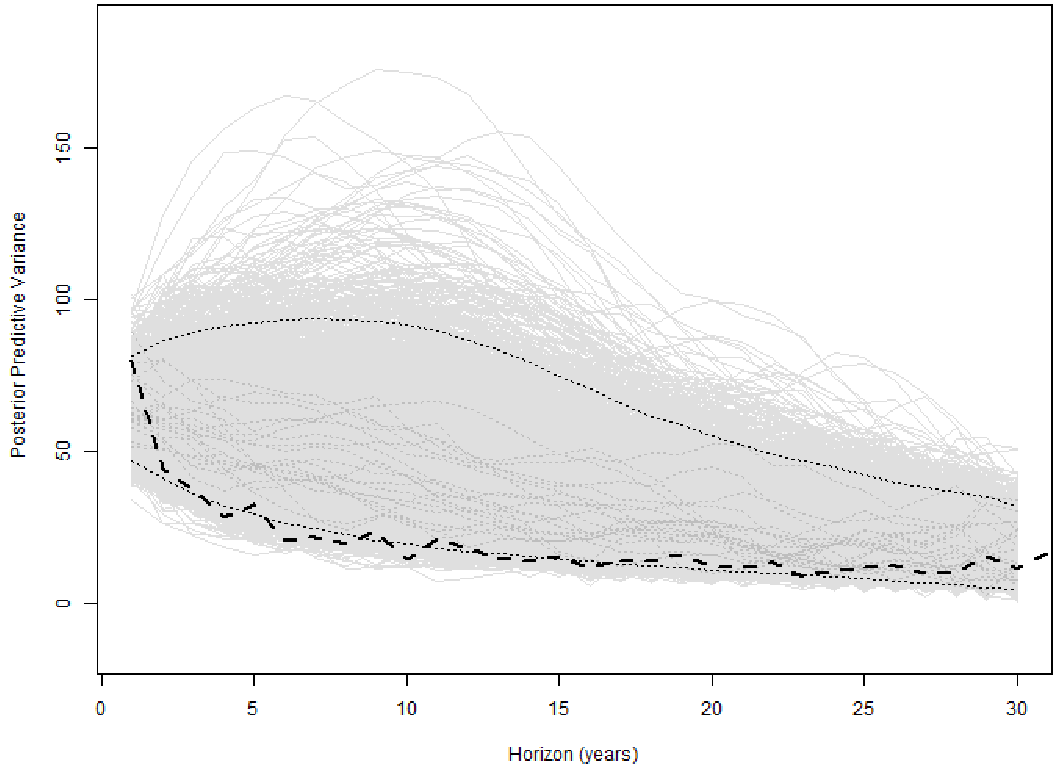

Figure 27. Posterior-based confidence intervals, $\rho_{uw} = -0.5$, nominal

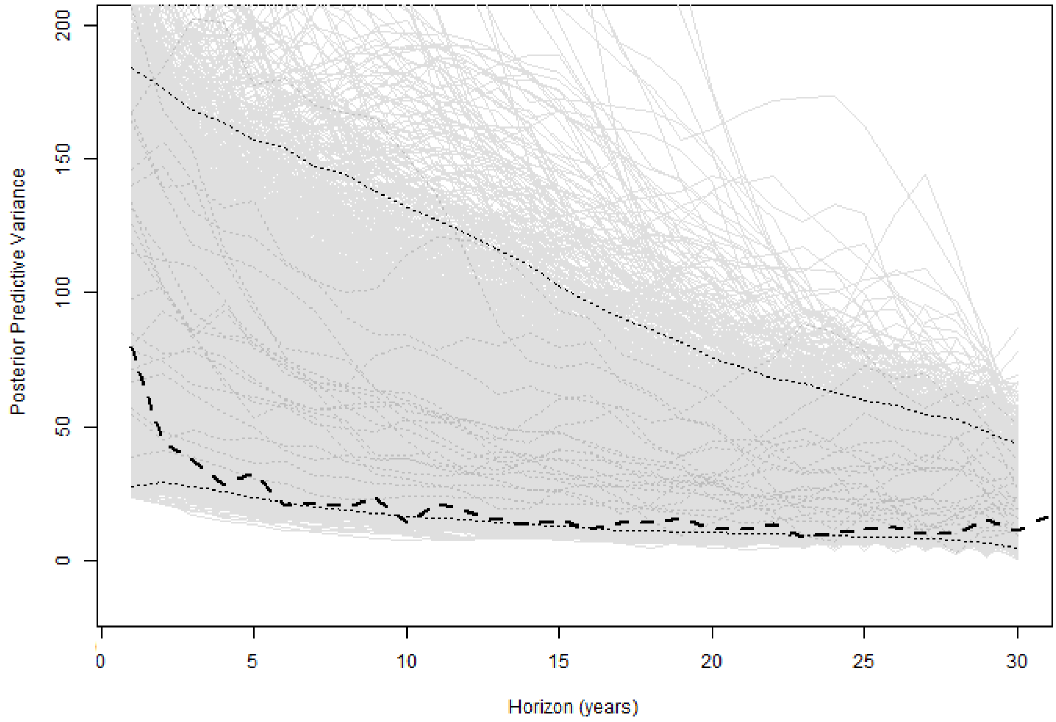

Figure 28. Prior-based confidence intervals test, $\rho_{uw} = -0.9$, nominal

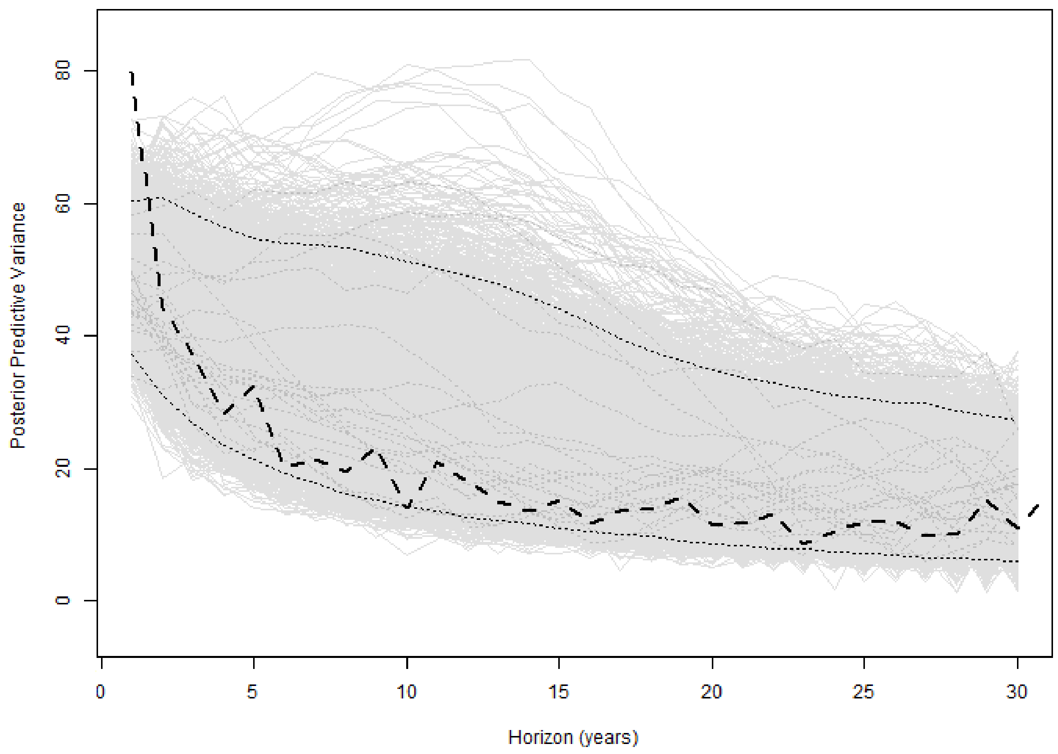

Figure 29. Posterior-based confidence intervals test, $\rho_{uw} = -0.9$, nominal

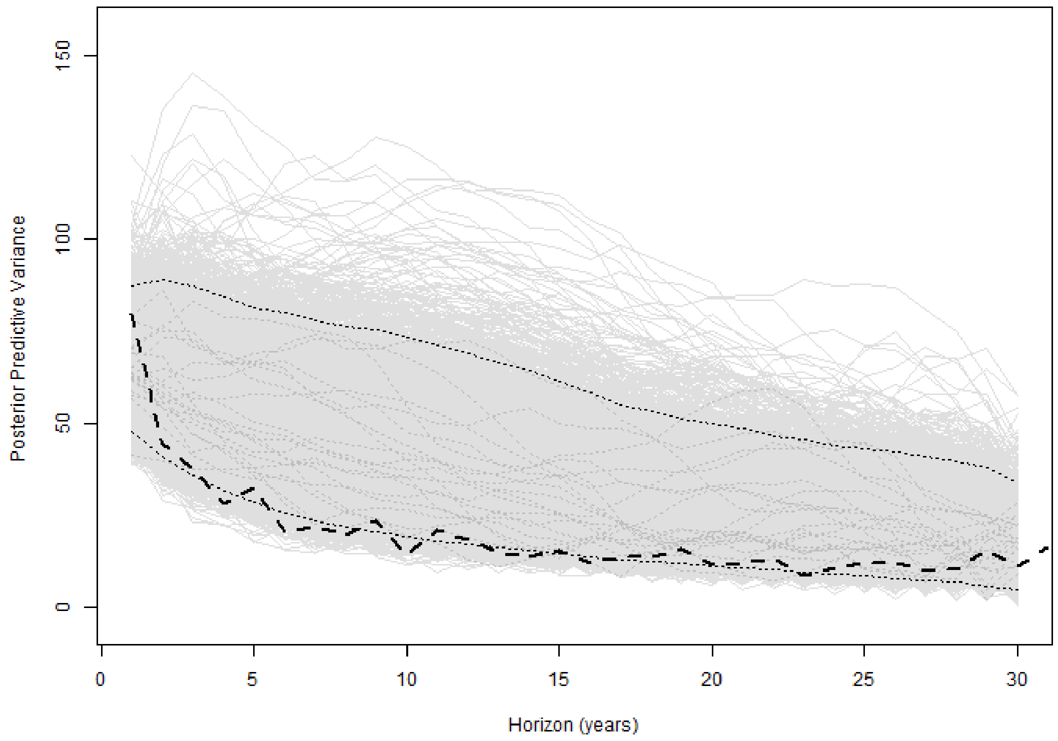

Figure 30. Prior-based confidence intervals test, time-varying Σ, nominal

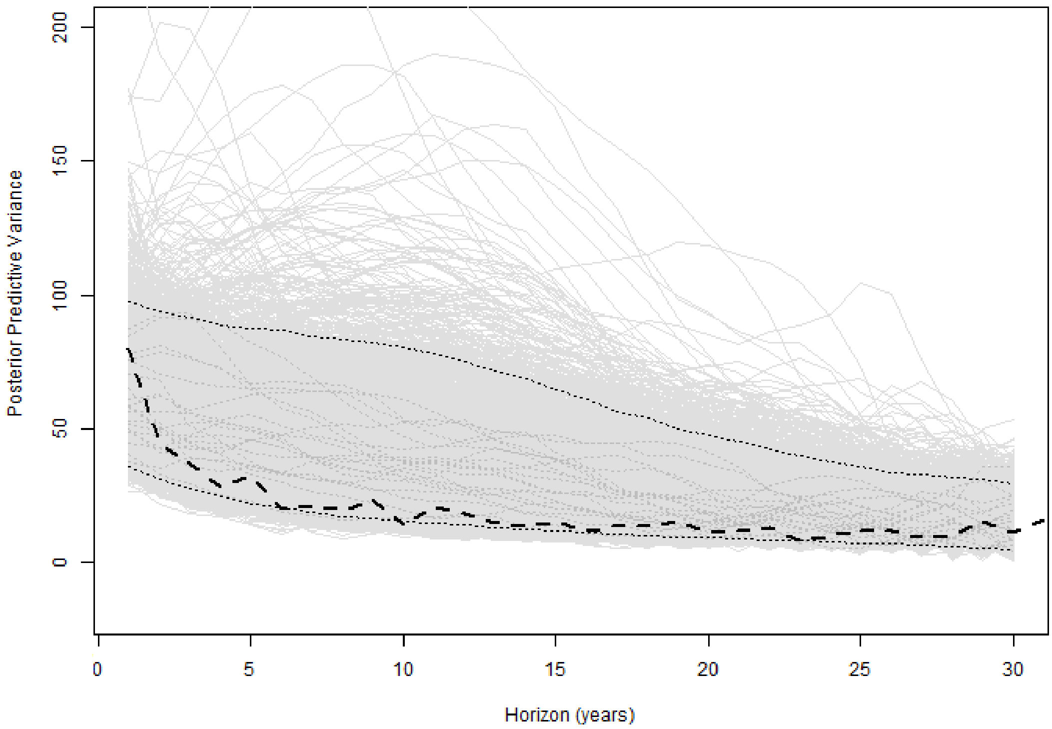

Figure 31. Posterior-based confidence intervals test, time-varying Σ, nominal

The time-varying Σ case fails to capture the unconditional variance line within its prior based confidence interval bands; therefore, it can be rejected as a valid model for the data. Its posterior results perform better than that of the weak ρ_{uw} case; however, they are not as good as in the strong ρ_{uw} case.

Considering all of the information collected, it can be concluded that the strong ρ_{uw} case provides the best and only valid fit in modeling the long-run variance behavior of nominal returns. The other two scenarios result in invalid models failing to capture the behavior observed in the data, and the amount of variation they explain is not as good as the strong ρ_{uw} case.

4.5 Real returns findings

The main theme of the findings does not change much when the real returns are considered in this section. The three scenarios, their parameters and the remaining initial values and priors are kept the same. The only difference is the dataset used and the initial value of β=0.48 as opposed to β=0.46 in the nominal returns case.

Figure 32 shows that there are small differences in the behavior of lines representing each scenario. The weak and strong ρ_{uw} cases bot start at the same point but the strong ρ_{uw} line decreases as time passes and starts climbing slower than the weak ρ_{uw} case.

The time-varying Σ line does not cross the strong ρ_{uw} line as it did in the nominal returns case and remains below the other two. The weak ρ_{uw} line is the only one coming close to violate the long run preferability condition.

The humped profile of the unconditional variance line is a result of a few abnormal observations in the dataset introduced by the CPI deflation calculations that could not be smoothed out because of the small sample size. This will introduce some difficulty in terms of determining the goodness of fit later on. It should be noted that this behavior is not specific to Turkey. Similar profiles are displayed by Canada, Australia, Germany and Japan when considering real returns as noted by Carvalho et al. (2018).

Figures 33-35 show that there are some differences in the real returns case regarding the α and β coefficients of the expected return equation. Unlike the nominal case, the α values here are not near-zero and cluster around 0.15.

The β values are also estimated to be a bit lower with around 0.4 values except for the time-varying Σ scenario that yields a very low β posterior value of near 0.3 despite having the same prior value clustering around 0.5.

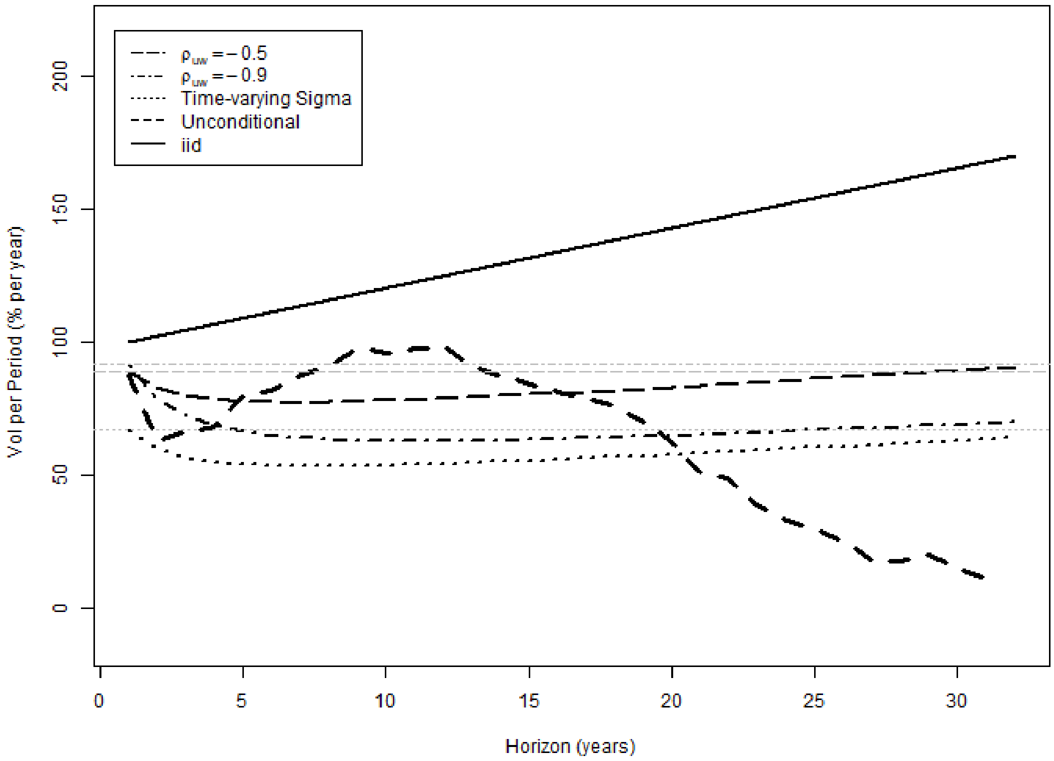

Figure 32. Predictive variance estimations, real

The prior/posterior behavior of the remaining Σ and ρ_{uw} parameters are identical to the nominal returns case, therefore they will not be reiterated or plotted here.

Regarding the components of predictive variance, the same three graphs are plotted as in the nominal returns case denoting the decomposition of predictive variance for each scenario.

Figure 36 shows that the weak ρ_{uw} case has a similar decomposition profile except that all positive contributors to variance are scaled up while the mean-reversion component line remains identical. This explains why the results obtained from the weak ρ_{uw} line suggest a higher predictive variance in the real returns case.

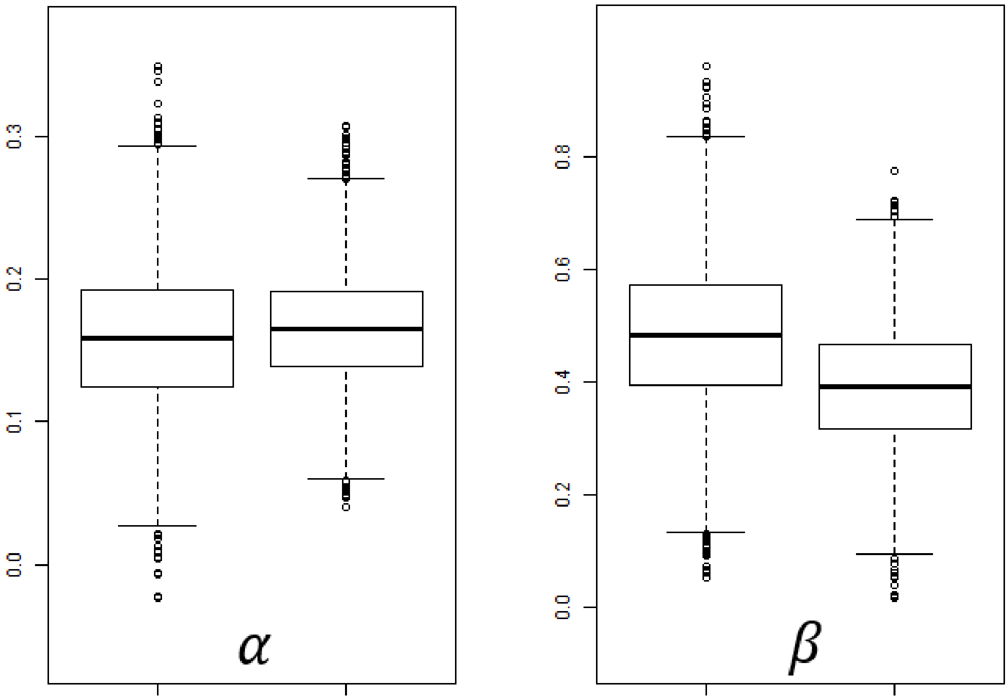

Figure 33. Prior and posterior distributions of μ coefficients, $\rho_{uw} = -0.5$, real

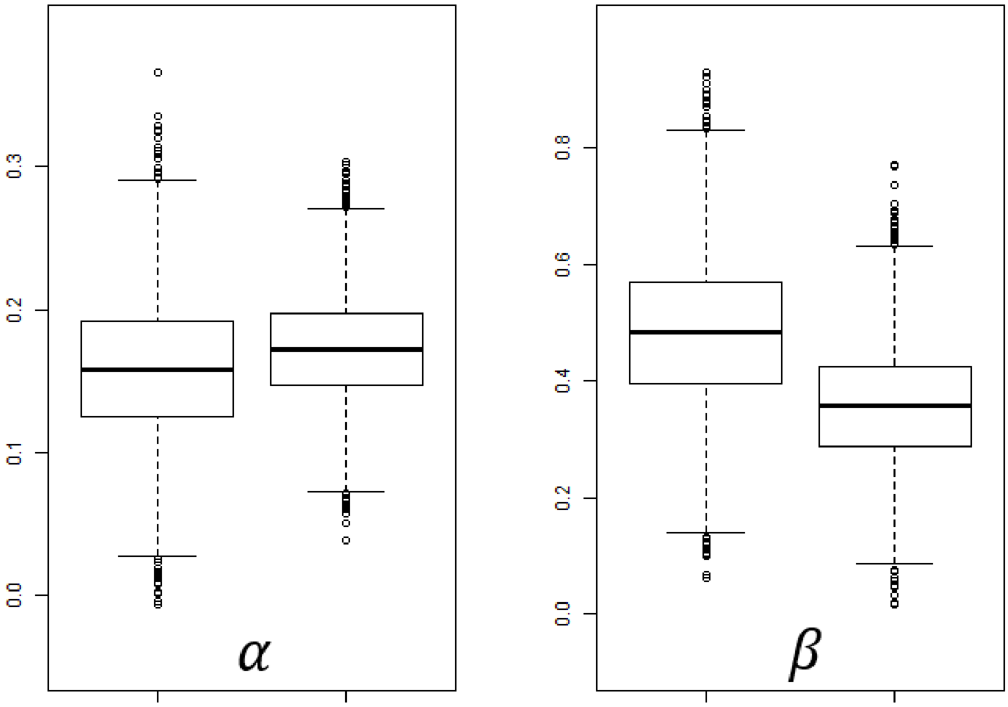

Figure 34. Prior and posterior distributions of μ coefficients, $\rho_{uw} = -0.9$, real

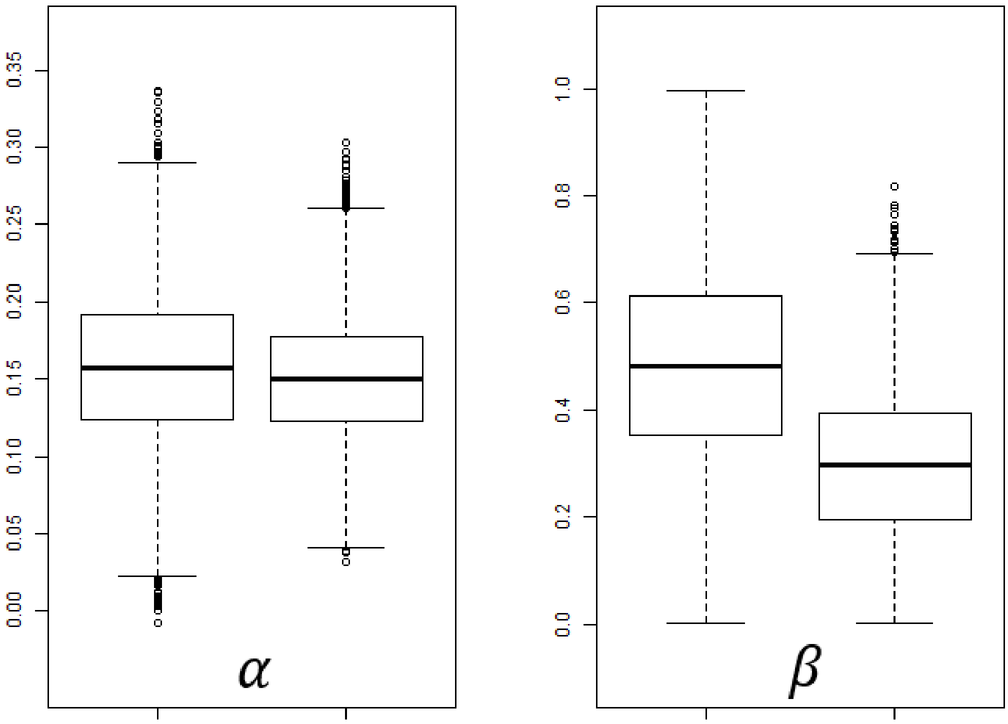

Figure 35. Prior and posterior distributions of μ coefficients, time-varying Σ, real

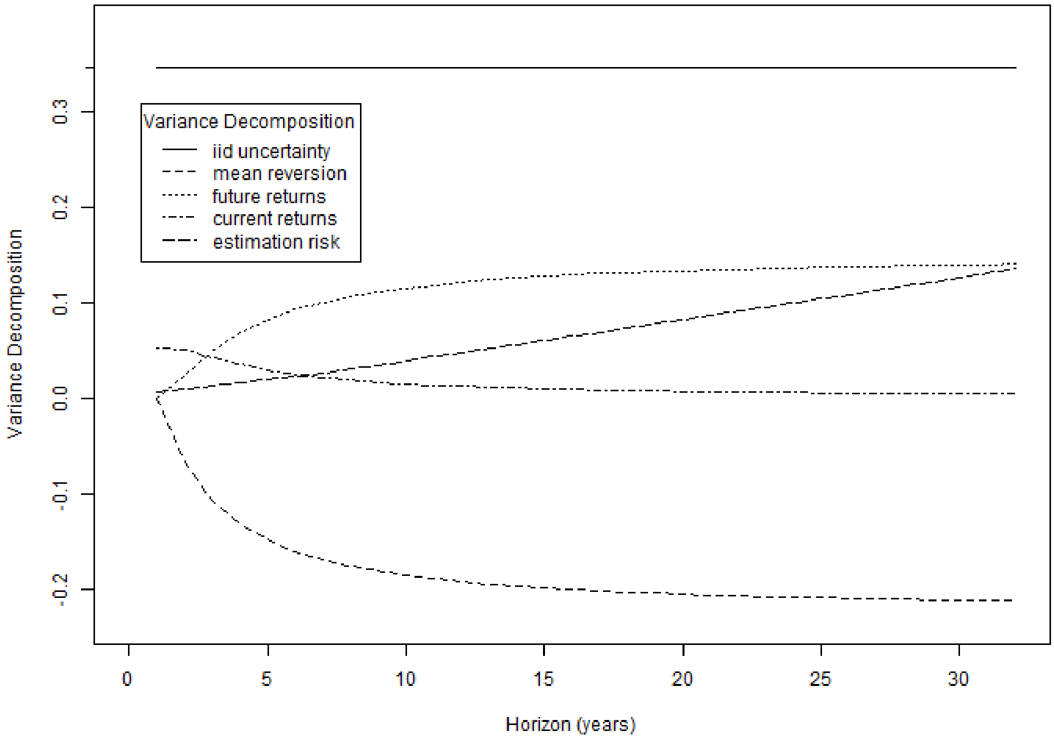

Figure 36. Components of predictive variance for $\rho_{uw} = -0.5$, real

The strong ρ_{uw} case is the opposite: Figure 37 shows that the mean-reversion component has a greater negative magnitude than the nominal case; therefore, reducing the overall value of the variance even more. It also drops the i.i.d. components from 0.4 to 0.24. All other components are near identical to the nominal returns case. Surprisingly, the overall results obtained are weaker in terms of supporting the long-run investments compared to the nominal returns analysis. The mean-reversion component may have a great contribution in terms of ratios; however, this does not automatically translate as a magnitude reduction in the final value of the predictive variance. This example makes a good case for why variance decomposition alone is not sufficient, and the actual value of variance needs to be calculated.

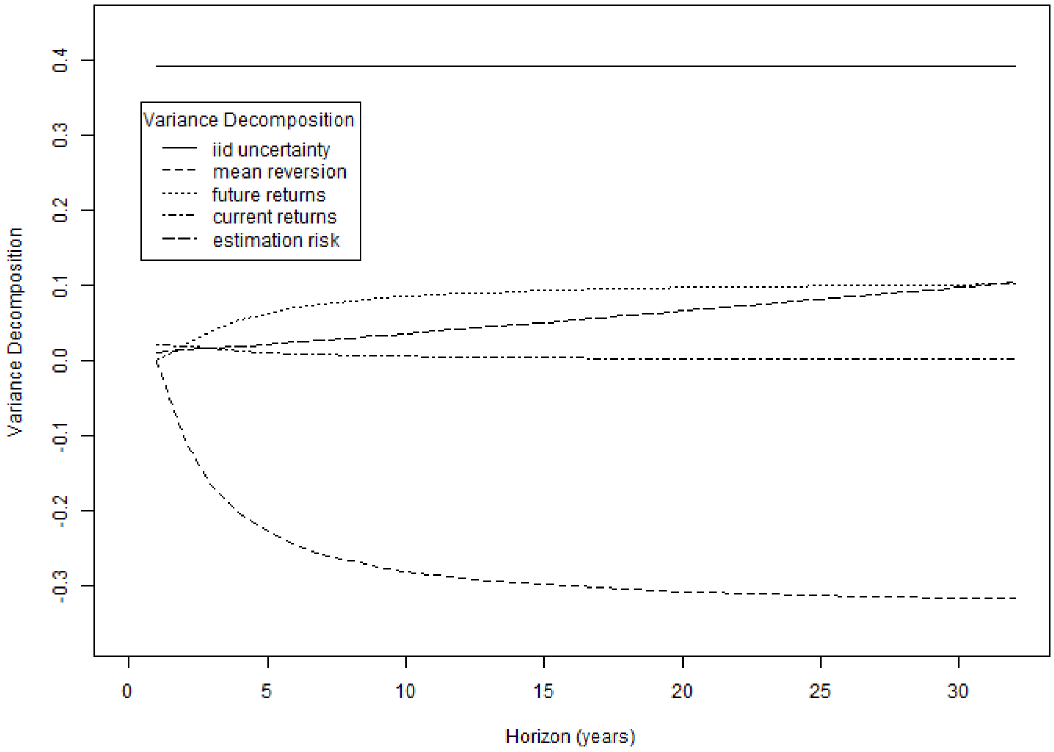

Figure 37. Components of predictive variance for $\rho_{uw} = -0.9$, real

The time-varying Σ scenario as shown in Figure 38 resembles the results of the nominal returns analysis, with the exception of having a slightly greater i.i.d. value. The predictive variance line estimated by the time-varying Σ case remains under the other two in the final graph, once again emphasizing that the decomposition of ratios does not translate to a one-to-one change in the final result.

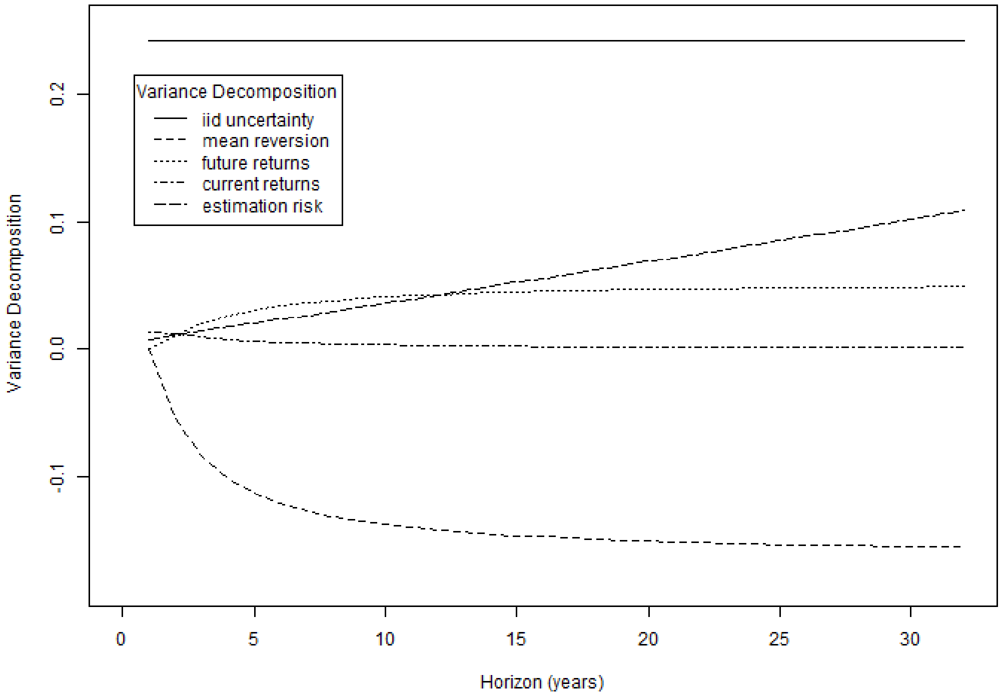

Figure 38. Components of predictive variance, time-varying Σ case, real

Goodness of fit for real returns

As mentioned before, the awkward shape of the unconditional variance obtained from the real returns introduces difficulties in assessing the goodness of fit of the three scenarios. This problem can be ameliorated by imagining a smoother curve or commenting on the posterior predictive checks by dividing the investment horizon into separate intervals. Either way, this section of the analysis cannot be performed under ideal conditions; therefore, its nature should be considered more in the spirit of speculative commentary and suggestions rather than deterministic results. Figures 39-44 plot the confidence intervals for all cases.

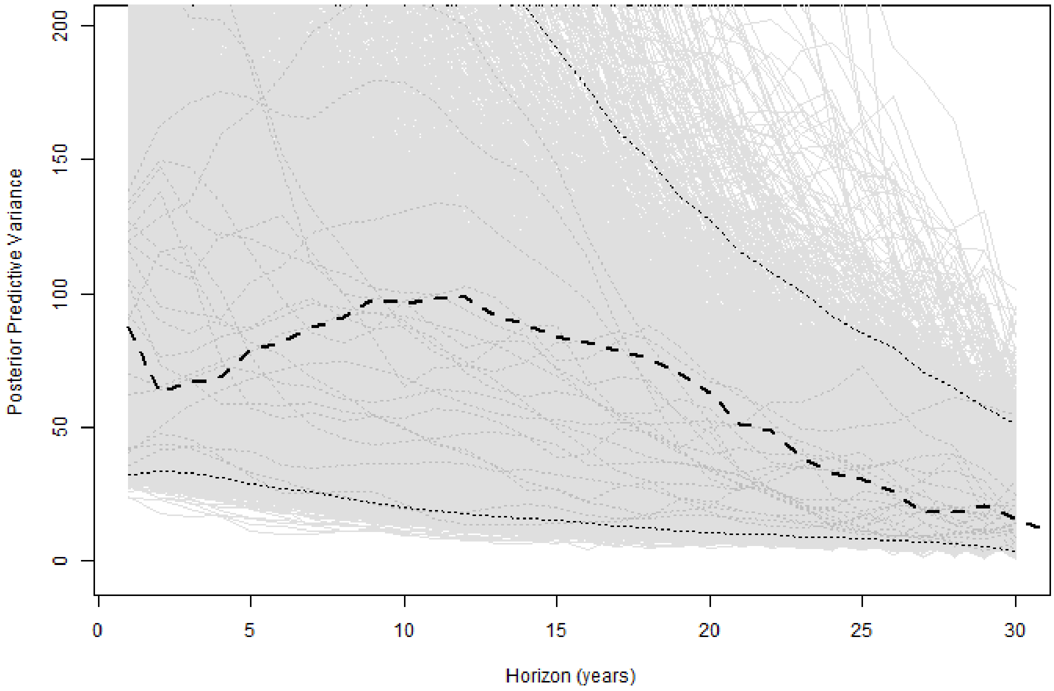

Figure 39. Prior-based confidence intervals, $\rho_{uw} = -0.5$, real

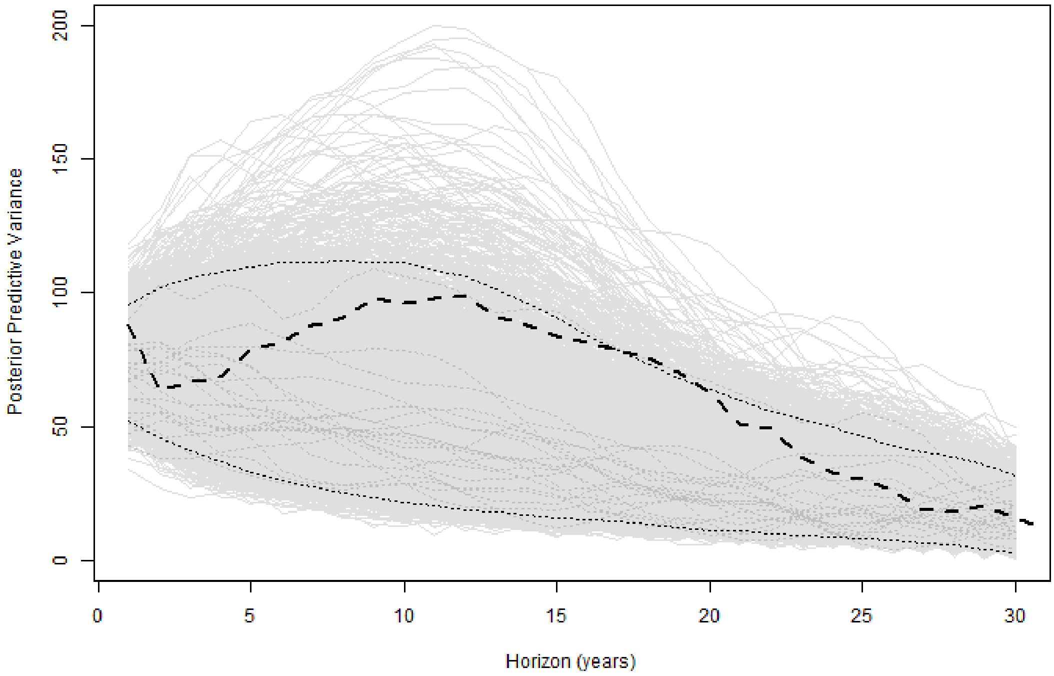

Figure 40. Posterior-based confidence intervals, $\rho_{uw} = -0.5$, real

Starting with the weak ρ_{uw} case, it is observed that both the confidence bands constructed using the prior and posterior simulated draws contain the unconditional variance, giving ample reason to not reject the null hypotheses that the weak ρ_{uw} scenario adequately explains the behavior observed in the model.

The evaluation of the strong ρ_{uw} case is a bit more problematic. The prior simulated draws result in confidence interval bands containing the unconditional variance; therefore, it can be concluded that the strong ρ_{uw} scenario presents a valid model. Its goodness of fit is very hard to assess: In the beginning and the end of the investment horizons, the bands contain the unconditional variance line very close to their center, suggesting a very good fit. The problematic hump in the middle section violates the confidence intervals and without any adjustments would force the conclusion that the strong ρ_{uw} case cannot adequately explain the patterns observed in the data with medium term investment horizons.

An adjustment can be proposed here assuming the hump is a result of non-representative behavior and draw a soother line to connect the beginning and the end of the unconditional variance line and argue that had the data maintained the behavior it displays at the beginning and end sections in its midsections, the strong ρ_{uw} case would have provided a very good fit. This result; however, would amount to little more than simple speculation without any way to confirm the proposal. One supporting piece of evidence would be obtained by referring back to the Chapter 3 and using the unconditional variance obtained through the rolling samples method. It would provide a very smooth decreasing curve that would fit inside the confidence intervals created by the strong ρ_{uw} estimations. Aside from that, there is little recourse to reconcile the results with any suspicions.

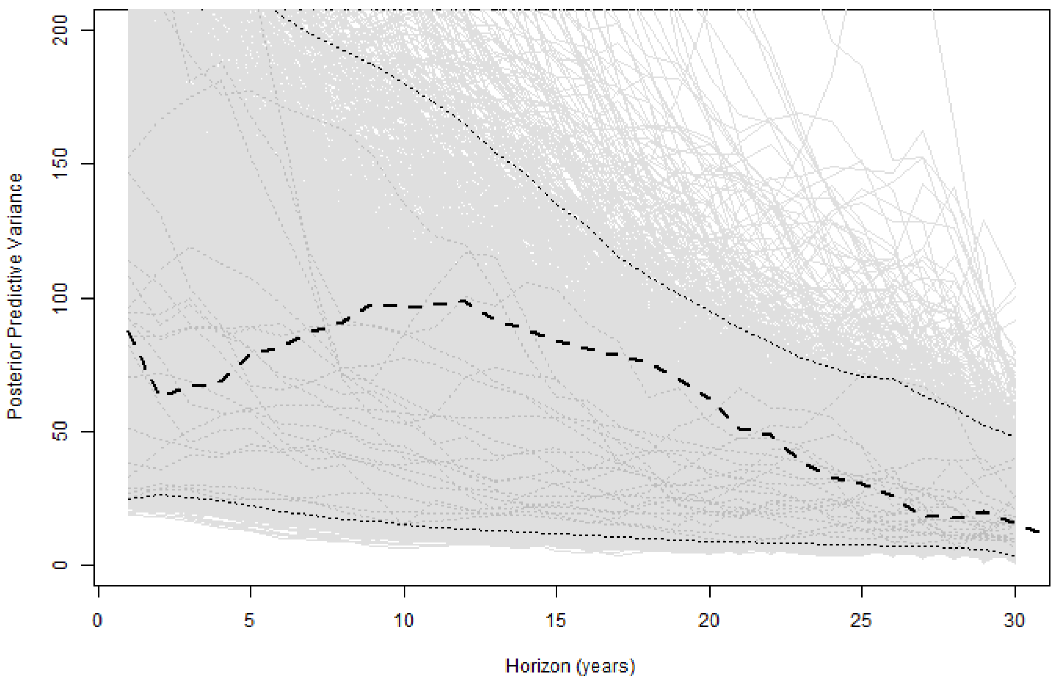

Figure 41. Prior-based confidence intervals, $\rho_{uw} = -0.9$, real

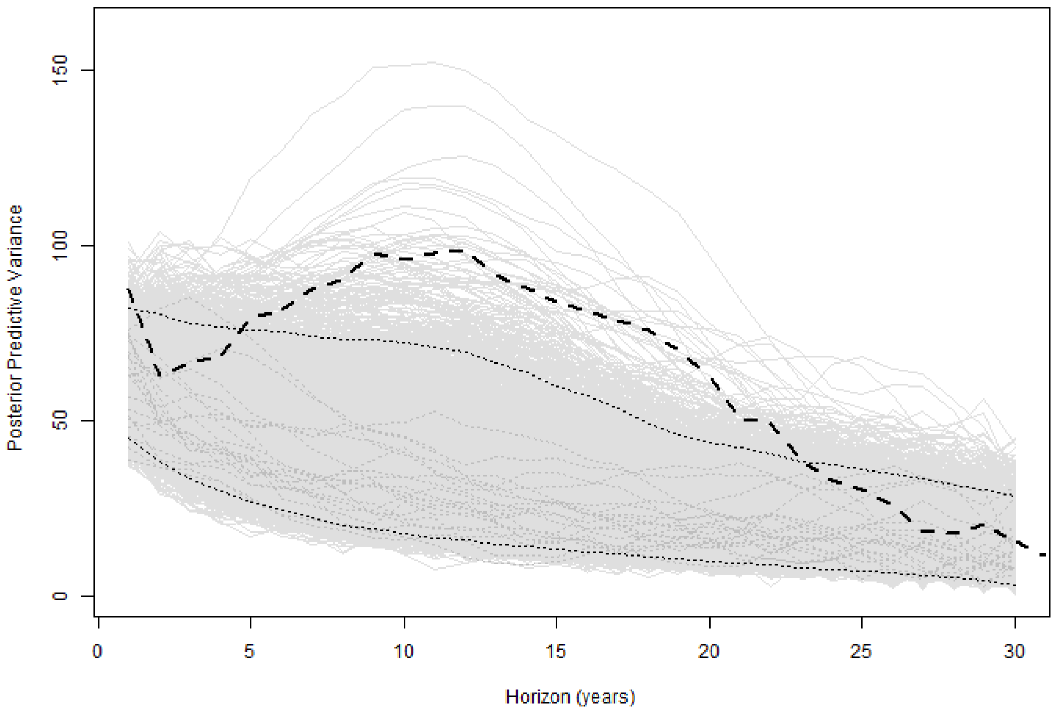

Figure 42. Posterior-based confidence intervals, $\rho_{uw} = -0.9$, real

The same problems are encountered in the time-varying Σ case, with the prior confidence intervals failing to contain the unconditional variance. The posterior predictive checks are more promising than the strong ρ_{uw} case since more of the unconditional variance line is contained by the confidence intervals. The time-varying Σ case can be argued to explain a third of the investment horizon.

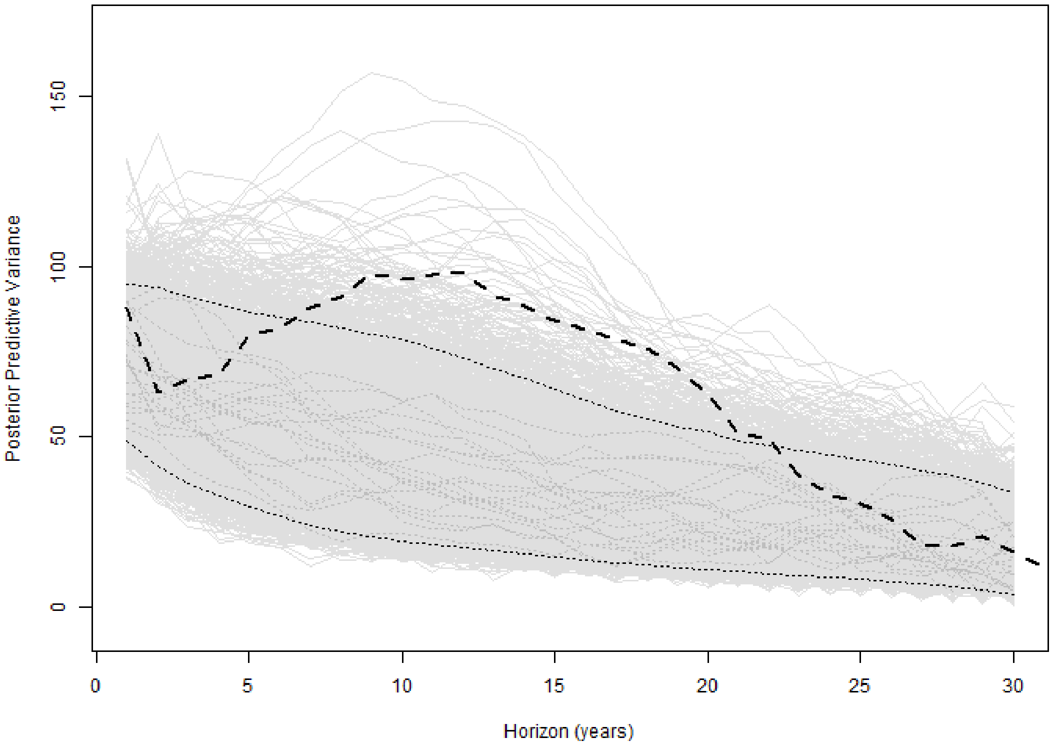

Figure 43. Prior-based confidence intervals, time-varying Σ, real

Most of this awkward shape is likely due to a few anomalies stemming from the lag occuring in how stock prices react to the inflation. Without any further data on the subject, the capacity to improve the goodness of fit results is limited. This section of the thesis should best be revisited after more years have passed and their results added to the dataset which is expected to smooth out the unconditional variance of the real returns case.

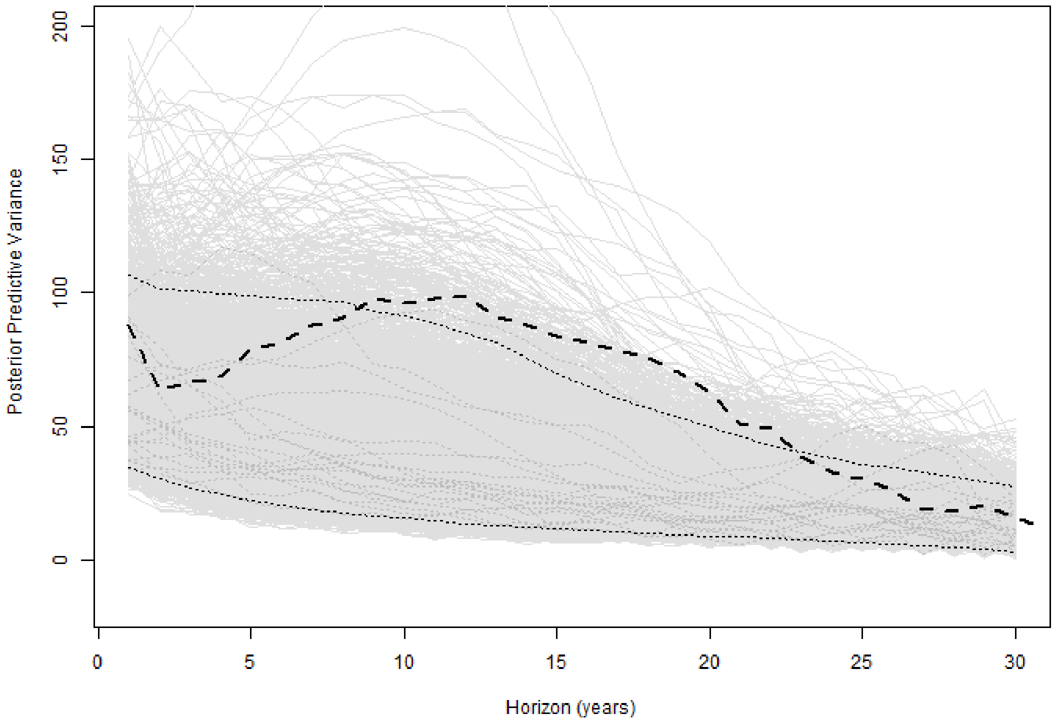

Figure 44. Posterior-based confidence intervals, time-varying Σ, real

4.6 Using predictors

The predictors include gross domestic product growth rate, foreign exchange rate between the U.S. dollar (USD) and New Turkish Lira (TRY) and the deposit rate.

Figure 45 and Figure 46 show that there is very little change in results when predictors are introduced into the system. This is in line with previous studies. It can be concluded that the dynamics created by the opposing momentum and mean-reversion effects are able to explain the behavior of variance over different investment horizons.

It can be argued that more suitable predictors could be found. By their very nature better predictors would better estimate the return series and conditioning the results on them would decrease the predictive variance. The results obtained in the previous sections would still be valid in that case, since long term preferability condition would be more easily met with smaller predictive variance values.

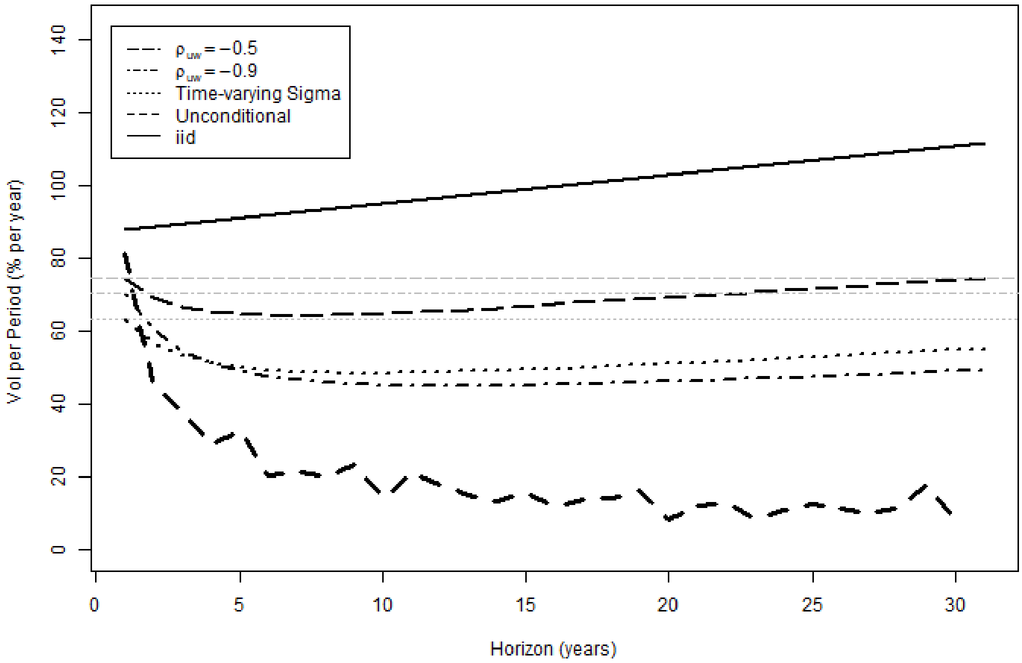

Figure 45. Predictive variance estimations using predictors, nominal

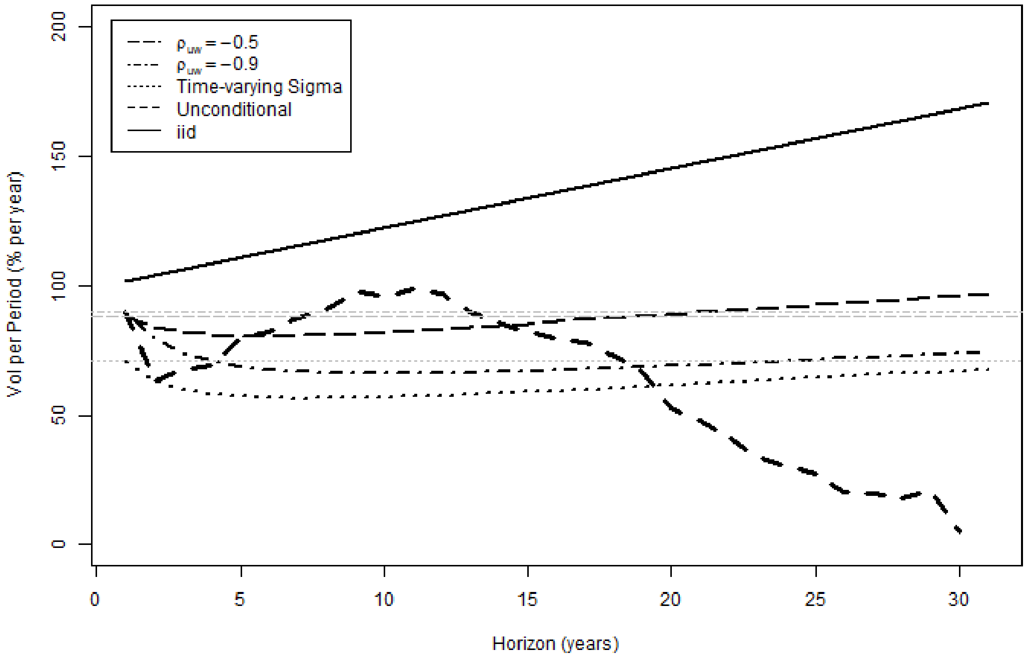

Figure 46. Predictive variance estimations using predictors, real

4.7 Under different β values

The value of β is one of the crucial drivers of the model. Therefore, it would be helpful to investigate the sensitivity of the results to changing β values. The following graphs rerun the analysis under varying priors for β. Everything else is kept the same.

Figure 47 shows that the $\rho_{uw} = -0.5$ case shows the greatest sensitivity to β. The results become more prone to be reversed with increasing β values. The long run preferability condition is violated before the end of the horizon once β crosses the 0.7 threshold.

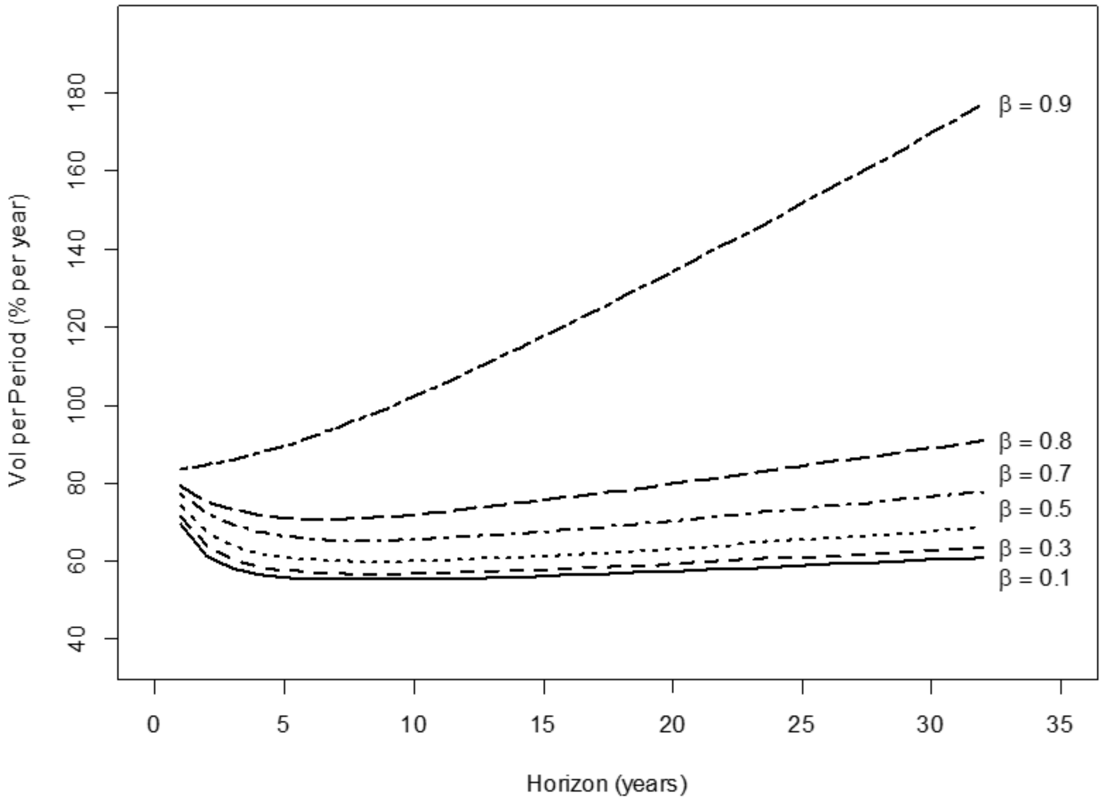

Figure 47. Estimations using different β values, $\rho_{uw} = -0.5$, nominal

Turning to the $\rho_{uw} = -0.9$ case in Figure 48, it is observed that the results are not very sensitive to changing β values and maintain the shape of the curve until β gets very large. This also signals the ability of the $\rho_{uw} = -0.9$ scenario to estimate the model with improved accuracy. An unstable estimation would be more sensitive to changing β values as in the case with $\rho_{uw} = -0.5$.

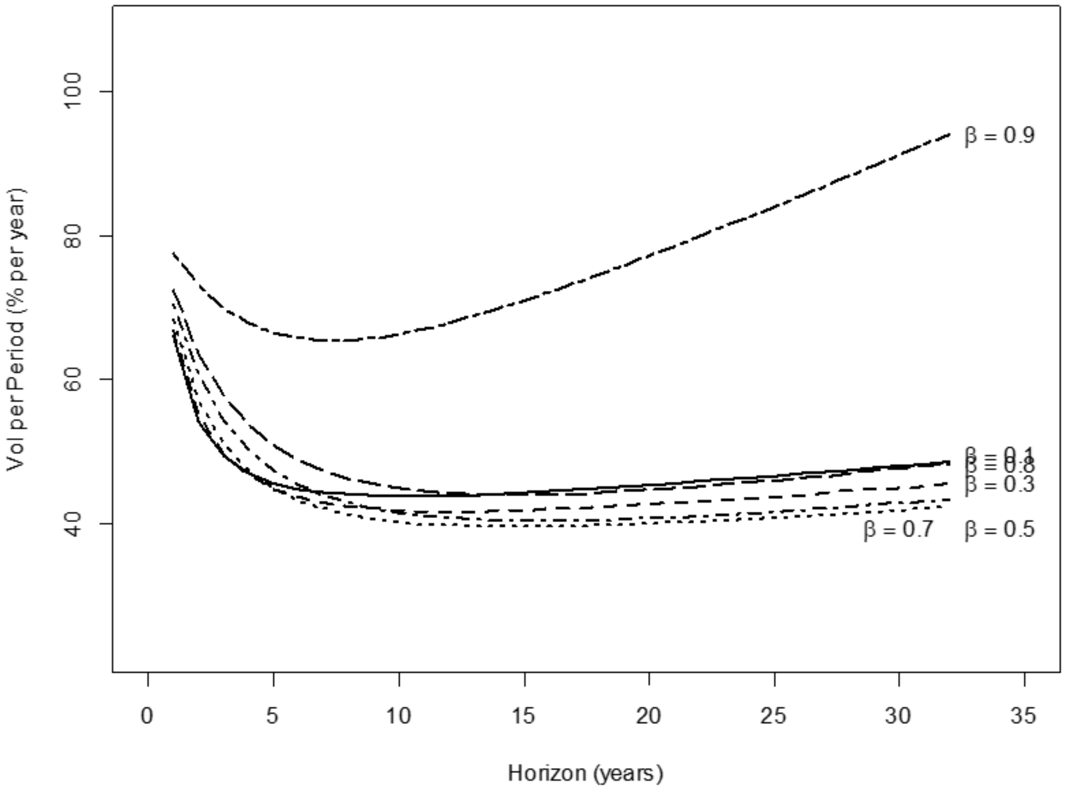

Figure 48. Estimations using different β values, $\rho_{uw} = -0.9$, nominal

The time-varying Σ case is the most stable out of all three scenarios with the shape and level of the curve showing very little sensitivity to changing β values, as shown in Figure 49. Once again, the results are only reversed for very large β values.

Further results may also be derived from the way these lines cluster under all three scenarios. The lowest risk levels are achieved for $\beta \epsilon \{0.3, 0.5, 0.7\}$. Decreasing the β to 0.1 or increasing to 0.9 also increases the risks and separates the line from the group. This clustering of the lines shows that despite the changing values of β, the results are overlapping; therefore, it can be argued that the true value of β is likely to be contained within the interval surrounding these values.

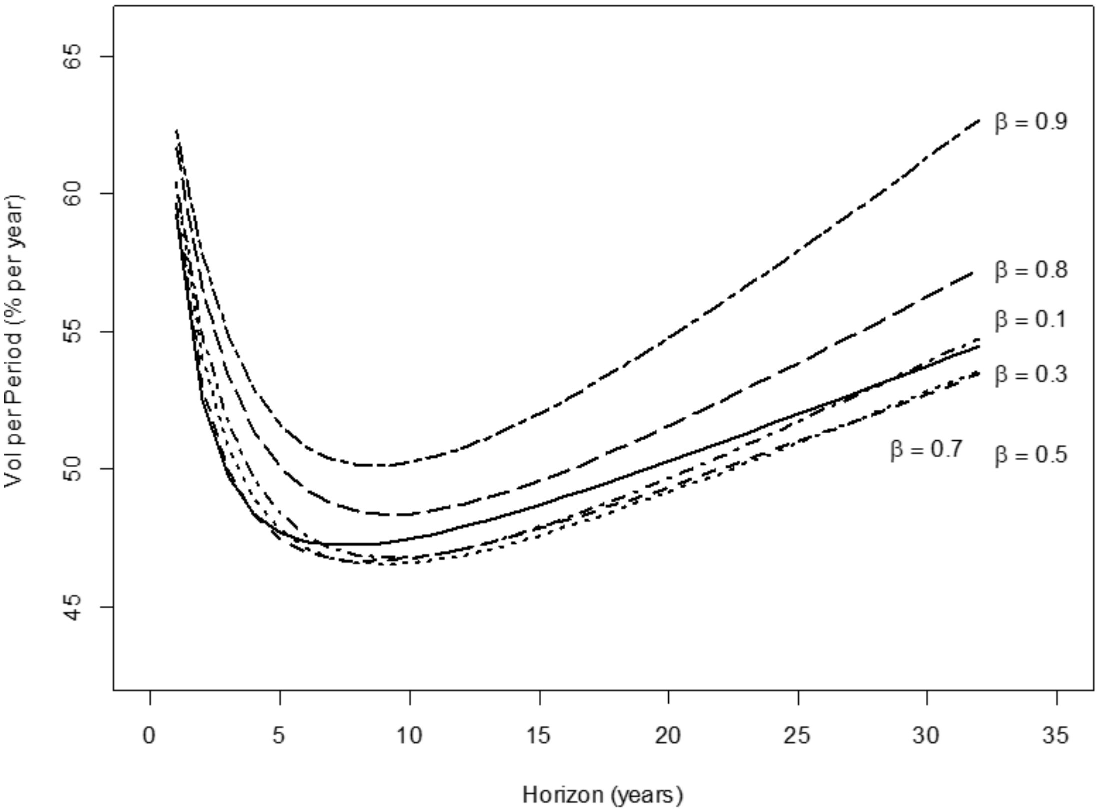

Figure 49. Estimations using different β values, time-varying Σ case, nominal

4.8 Allowing negative β values (negative stationarity)

The results obtained so far have all been under the assumption that β is between zero and one. Any values obtained outside this range would be truncated throughout the iteration. As a thought experiment, this section allows for negative β values to see whether any changes to the results would happen and how the model would handle a deviation from its standard applications.

Analyzing the sample mean returns data, it is observed that the Augmented Dickey Fuller test rejects the null hypothesis of no unit root with a 5% significance; but cannot reject it at 1% level. Continuing with the 1% level, the unit root can be removed by taking first differences. Re-running the AR(1) analysis on first differences yields a β value of -0.46. It would be interesting to investigate how the estimation behaves under these conditions, since the negative ρ_{uw} values usually are the driving force behind the mean reversion in the system. Under a negative β assumption, the expected returns would zig-zag; however, the magnitude of the shocks would still keep their momentum. It can be argued that it is sufficient for the β to have a large enough absolute value to counter the mean-reversion effect caused by the negative ρ_{uw} values.

Figure 50 shows that the mean-reversion component of predictive variance becomes smaller; however, the long-run preferability condition is not violated. In fact, the climb after the dip observed in predictive variance is much slower. All three lines move closer to each other with the time-varying Σ case maintaining a nearly identical line to the positive stationarity case. The weak and strong ρ_{uw} cases are both lowered and more horizontal.

The same results apply to the real returns case as shown in Figure 51. Combining the findings, it is observed once again that the time-varying Σ case is the most robust one out of the three scenarios. Both the strong and weak ρ_{uw} lines approximate the time-varying case. This suggests that assuming a negative β reduces the affect that negative ρ_{uw} has on the system.

The most interesting result is that the general shape and level of each line is very similar to the positive β case. Combined with the rest of the results it can be conclude that: i) the time-varying case offers the most robust estimation, and ii) negative β values may have the ability to assume the role of negative ρ_{uw} values and combine the two dynamics into a single variable. Further investigations are necessary to confirm this; however, they fall outside the scope of this thesis.

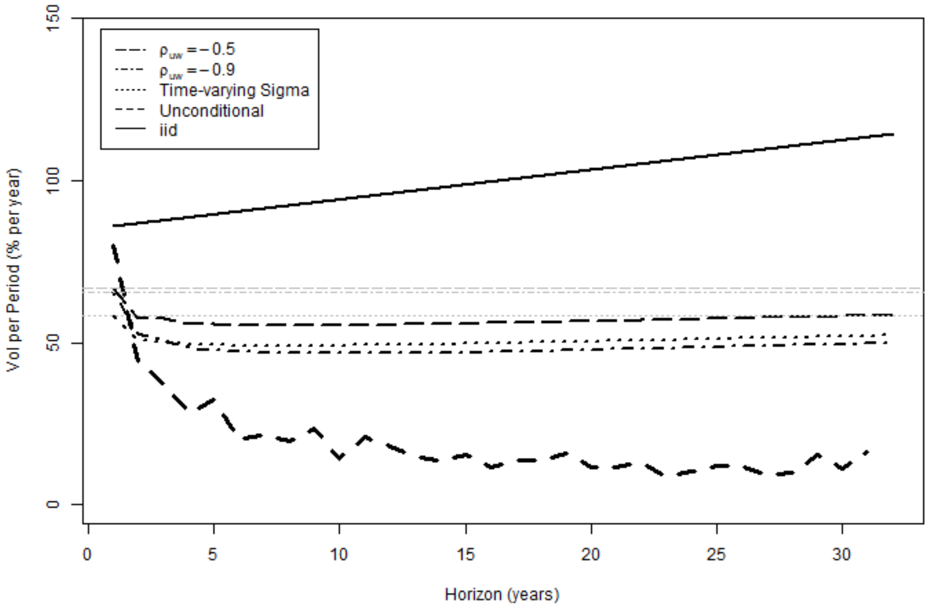

Figure 50. Estimations with negative β values, nominal

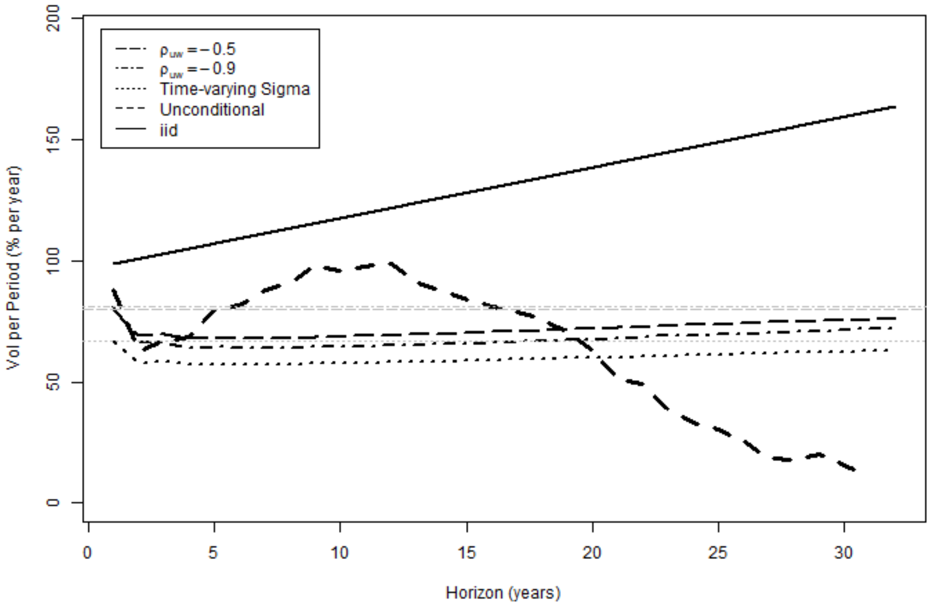

Figure 51. Estimations with negative β values, real

An important finding is that both the prior and posterior confidence intervals contain the humped unconditional variance line of real returns. This suggests that the negative case actually provides a better goodness of fit than the positive β case. It is an unexpected result and once again it is left to future studies to determine the merits of these results and their wider applicability.

The goodness of fit for the nominal returns case does not show the same drastic difference and for the most part akin to its positive β counterpart as shown in Figure 52 and Figure 53. It may be proposed that using negative β values with a time-varying Σ in humped unconditional variance cases can be a possible solution to better fit the data.

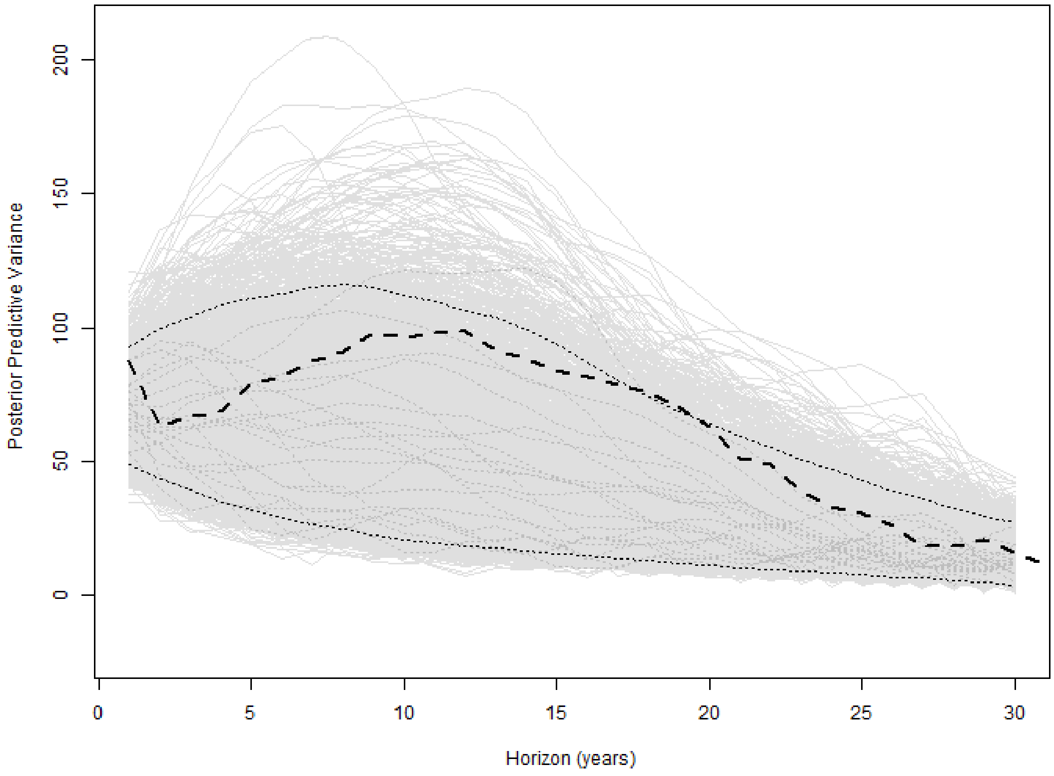

Figure 52. Prior-based confidence intervals, negative β values allowed, time-varying Σ case, real

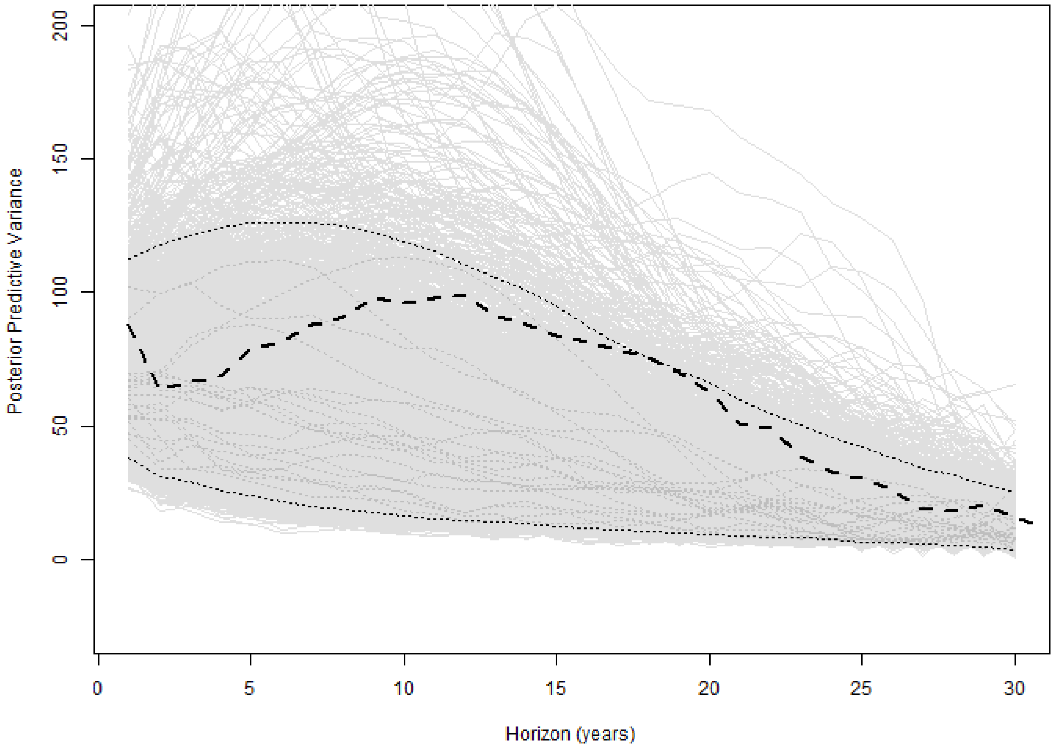

Figure 53. Posterior-based confidence intervals, negative β values allowed, time-varying Σ case, real

4.9 Consequences of the small sample size

Our sample size of 32 observations is considerably small compared to some of the previous studies using Bayesian MCMC analysis. The analysis essentially runs on two axes and the length of these axes matter in terms of estimating power. The first axis is the algorithm iterations, and its length is user-determined as 5000. The second axis is the forward iteration and backwards sampling steps that shuttle along the returns sample, the length of which is determined simply by the sample size. The number of algorithm steps can be increased; however, more annual return observations cannot be created. Therefore, there is a hard cap on how well the estimation will be able to reveal the underlying dynamics of the predictive variance. In other words, the results from the small sample should be expected to overestimate the predictive variances. As more data becomes available with time, the predictive variance estimates can be expected to be lower from the values calculated here.

In order to determine whether and how the sample size affects the results, the same analysis can be performed for a larger dataset and see what changes when the sample size is restricted to 32 consecutive observations starting at any point in the dataset. A good candidate is the U.S. dataset constructed by Siegel (2008) as it is used in previous predictive variance analyses such as Pastor and Stambaugh (2012) and Carvalho et al. (2018). This dataset has 205 observations. The most recent analysis carried out by Carvalho et al. (2018) can be replicated as a benchmark and then restrict it to see how the results vary.

The replication process reperforms their analysis, plots the main comparative variances graph and arrives at the same conclusions as Carvalho et al. (2018) as displayed in Figure 54. In the next steps, the last 32-year period is extracted as a separate sample and the same analysis performed in this restricted setting. Figure 55 shows that the results are reversed, and it becomes less preferable to invest in the long run. It appears that a smaller sample size makes it less likely to achieve results supporting long term investments.

It is hard to tell whether this guarantees that had there been more observations, the results would be even more supportive of long-term investments. It is not clear whether the results obtained from the restricted U.S. sample can be generalized the case for Turkey. Regarding the sample, it can be argued that the underlying dynamics must be strong enough to produce such results even under the restriction of a very small sample considering how the results in the U.S. data are reversed in small samples.

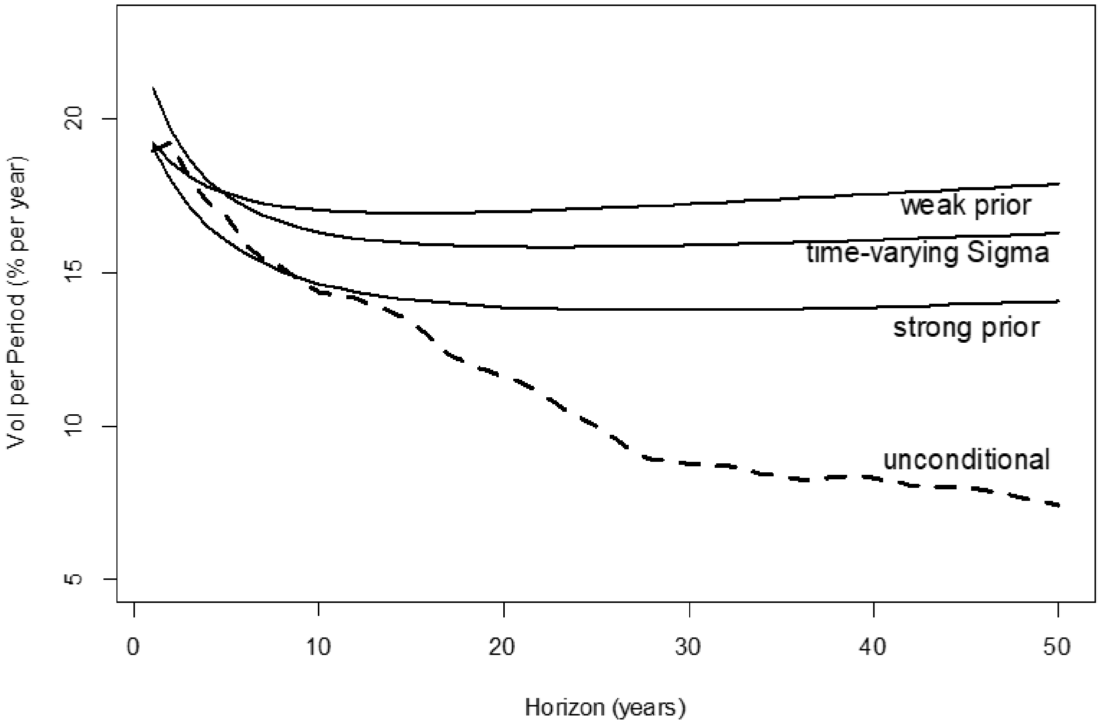

Figure 54. Predictive variance values for the entire U. S. data

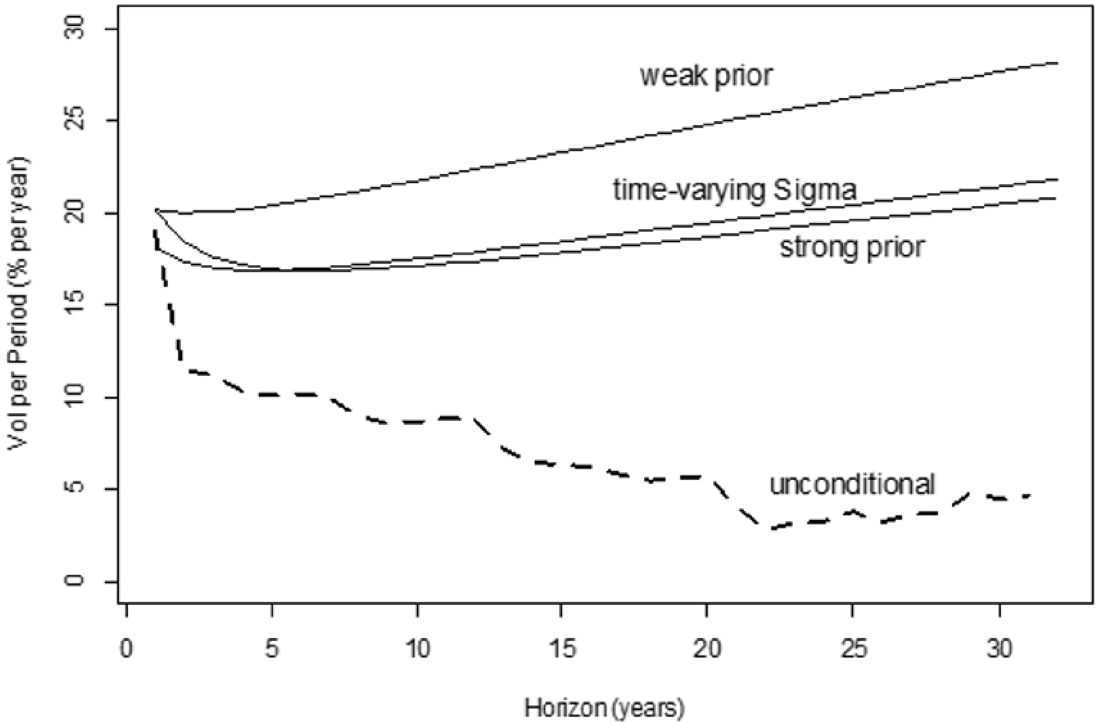

Figure 55. Predictive variance values for the restricted U. S. data

4.10 Conclusions for the predictive analysis

The predictive variance analysis performed on both nominal and real returns suggest that it is more preferable for an investor to invest in the long run under all three scenarios with varying degrees of support. The momentum coefficient β does not seem to be close to one, so the value of ρ_{uw} plays a key role in determining the overall value of the predictive variance. The time-varying Σ case presents a different pattern than the other two in that it lowers the predictive variance by not countering the momentum with a negative ρ_{uw} value but reducing the i.i.d. component through its modeling assumptions. In the nominal returns case, the strong ρ_{uw} scenario provides the best fit to the data. The goodness of fit is hard to assess in the real returns case due to the abnormal shape of the unconditional variance line introduced by the interaction between the stock prices and the CPI deflator along with the lagged response. Further adjustments and assumptions are required to determine the best fit. Without those, weak ρ_{uw} scenario becomes the winner by default.

The constant coefficient α of the expected return equation is found to be near-zero in the nominal case whereas it has a positive value of 0.15 in the real returns case. All elements of the Σ matrix and ρ_{uw} values have narrower posterior distributions than their prior distributions. The ratio of error term variances σ_u^2 and σ_w^2 is small but not near-zero. The low-signal-to-noise ratio problem may not be as detrimental as initially feared.

Including predictors in the analysis does not have too much of an effect on the final results. It can be reaffirmed that the push-pull dynamics created by the momentum and mean-reversion affects adequately capture the behavior observed in the data.

Performing the analysis under varying β values between 0 and 1 provides further insight on how robust the findings are to the changes in this key parameter of the model. Apart from the weak ρ_{uw} scenario, none of the results seem to be highly sensitive to the initial β value and the analysis seem to pull it in the prior and posterior distributions to levels near 0.5, reaffirming the suspicion that the β value is not high. As the value gets away from 0.5 in either direction, the resulting predictive variance becomes slightly higher. The results seem to change only when it approaches 0.9.

Allowing a negative β value surprisingly yields nearly identical results. The weak and strong ρ_{uw} cases approach each other while approximating the time-varying Σ line. This can be interpreted as the significance of the ρ_{uw} parameter decreasing, and its effect being absorbed by the negativity of the β parameter. This may open up new avenues for further research; however, it is left to future studies to maintain the scope.

Having a small sample size of 32 puts a hard limit on the estimation precision of the model. The number of iterations in the algorithm can be increased; however, the number of available annual return observations determining the number of forward iteration and backwards sampling steps performed along the dataset cannot be changed. In order to demonstrate whether this has a detrimental effect in the estimation capability and in what direction the results would have moved had more data been available, an experiment is conducted. The results obtained by Carvalho et al. (2018) using the U.S. data is replicated. Then the same analysis is reperformed by restricting the sample size to the last 32 years. The results are reversed, and all three predictive variance measures display very steep profiles. This can be interpreted as partial evidence that even more supportive results can be achieved for long-run investments when more data becomes available in a few years.

CHAPTER 5

THE BAYESIAN MCMC ALGORITHM DETAILED

Bayesian estimation methods are sometimes criticized for being black boxes since they involve algorithms consisting of heavily nested loops. They may appear daunting at first, but the principles behind them are fairly simple. This section aims to dispel such notions and explain in detail how the results are obtained for anyone to be able to replicate and improve upon the work presented here.

The algorithm is run in R programming language using the RStudio software. The code utilizes prepackaged function libraries in addition to the custom Bayesian MCMC functions made publicly available by Carvalho et al. (2018), building upon the work of Pastor and Stambaugh (2012). Custom portions of code have been written by the author of this thesis to extend and modify the algorithm for the specific requirements of the Turkish dataset as well as presenting additional simulations and graphs. The sources of each of these code parts will be cited along the way while explaining the inner workings of the algorithm in the following sections.

The code consists of multiple nested parts based on the hierarchy and order of the sub-analyses performed. These divisions are structured in small chunks rather than a single block of code to provide the readers with an easier understanding of the high-level logic of the analysis steps while allowing a deeper dive into the specifics in the sub-files.

The first step is to run the top-level code contained in “ThesisMainCode.R” file. It essentially lists the analyses performed and requests the key inputs from the users. Multiple other files in the same folder are referred to for each major part of the analysis. Table 1 contains the main code of the algorithm.

Table 1. The Main Code of the Algorithm

ThesisMainCode.R

```
source("CleanUp.R") # Resets everything to start from scratch
source("PrepWork.R") # Packages, libraries etc.

# Inputs:
niter = 5000 # Number of total Bayesian MCMC iterations
m0    = -0.025 # Initial mean for Mu
C0    = 0.025 # Initial variance for Mu
AlphaVariance = 0.05^2 # Initial variance for Alpha
BetaMinCutOff = 0 # '0' to force positive values, '-.9999' to allow negative stationarity
AnalysisType = "Nominal" # Whether to use the nominal or real returns
TVS = 0 # Also perform a time-varying Sigma Analysis
Predictors = 0 # Include predictors?

source("CreateSaveDir.R") # Creates a new save folder reflecting the features of the analysis

source("BetaPriorElicitation.R") # Determine the initial value of Beta from the dataset

# BetaParameter = 0.5 # Uncomment to manually enter the initial value for Beta
# BetaParameterTV = 0.5 # Uncomment to manually enter the initial value for Beta in time-varying
case
BetaVariance = min(.133,(1-BetaParameter),BetaParameter)^2
# ^^ Set the Beta variance so that it is concentrated in a suitable range

source("LoadDataset.R") # Load and prepare the dataset

source("IIDcase.R") # IID model

source("UnconditionalVariance.R") # Unconditional variance

Rho=-.5; SubAnalysisName = "Weak Negative Rho" # Weak Prior setting: rho = -0.5
source("BayesianMCMC.R")

Rho=-.9; SubAnalysisName = "Strong Negative Rho" # Strong Prior setting: rho = -0.9
source("BayesianMCMC.R")

if(TVS) { SubAnalysisName = "Time-varying Sigma" # rho = -0.9
source("BayesianMCMC_TVS.R"); TVRho = (exp(LongRunPastor)-1)*100 }

source("DrawComparisonGraph.R")

source("SaveImage.R") # Saves the necessary data for later examination
```

The first sourced file is “CleanUp.R” which essentially resets the software and clears the data environment to perform the analysis from scratch every time. The second one, “PrepWork.R”, handles the general preparations necessary for the R programming language such as loading the custom function files and libraries. Once the preparations are complete, user defined inputs are requested to begin the analysis. The default values and explanations are provided in Table 2.

Table 2. Default Values of User-Defined Parameters

Variable Name	Default Value	Explanation
niter	5000	Number of algorithm iterations
m0	-.025	Initial mean for μ
C0	0.025	Initial variance for μ
AlphaVariance	0.0025	The variance for the α parameter
BetaMinCutOff	0	The minimum value for the β parameter to ensure positive stationarity
AnalysisType	“Nominal”	Whether to use the nominal or real data
TVS	1	Whether to perform an additional analysis for time-varying Σ matrix case
Predictors	0	Whether to include predictors in the analysis

Once the inputs are defined, "CreateSaveDir.R" file is run to create a folder that will contain the analysis results of the current attempt. Its name would reflect the major properties of the current analysis settings.

The user can choose to run the "BetaPriorElicitation.R" function to determine the initial value of β from the dataset instead of manually picking a value. This step is shown to be redundant for the analysis here during the robustness checks since the results are similar for the likely values of β; however, it is kept for users who may wish to use a different dataset and have an educated guess for their initial β value.

The initial value for the variance of β is set such that three standard deviations worth of its mass would be contained within its initial value and one. This is to reflect the assumption of momentum in expected returns.

"LoadDataset.R" is tasked with loading the dataset and modifying it depending on the conditions applied such as allowing a time-varying Σ matrix or using predictors.

The rest of the files take care of the actual analysis. In a single run, five separate comparative variance measures are computed:

i. Unconditional variance
ii. i.i.d. case
iii. Predictive variance with weak $\rho_{uw} = -0.5$
iv. Predictive variance with strong $\rho_{uw} = -0.9$
v. Predictive variance with a time-varying Σ matrix with strong $\rho_{uw} = -0.9$

Each of these sub-analyses are carried out by running separate files. The associated variables are recorded in separate datasets for later comparison.

The "UnconditionalVariance.R" file calculates the unconditional variance based on the parameters entered and the data. The "IIDcase.R" file calculates the variance under the identical and independently distributed errors assumption. The last three sub-analyses handle the predictive variance analysis under different beliefs. The first two use a constant Σ matrix and only change the value of ρ_{uw}, which is the correlation between the error terms of current return and expected return equations. The last one assumes a time-varying Σ matrix and $\rho_{uw} = -0.9$. The final file "DrawComparisonGraphs.R" plots the comparison figure juxtaposing all the results on the same graph along the investment horizon axis.

Finally, "SaveImage.R" saves an image of the dataset. This is necessary, since Bayesian methods by their nature yield slightly different simulation results each time they are run. Ideally, the results would be similar. In the case that there is an anomaly, the simulated part of the data can be examined manually to identify and correct the problems that may be occurring with the algorithm.

5.1 Code implementation

This section details the inner workings of the algorithm to reveal how each of the simulated values are calculated. The implementation of the computer code is not as clear cut as the analytical model: simple calculations may require many lines of code to carry out in a software environment. It is hoped that this section would enable future researchers to replicate and advance the analysis presented here.

The code can be thought to consist of two classes: the main analysis query and the function libraries. The previous overview section summarized the general approach and order of calculations. This section will dwell on the more granular details of the algorithm. It is recommended that the latest versions of the RStudio application and the R programming language are installed as some calculations require recently introduced functionalities to both. For Mac OS users, a version of

Quartz such as XQuartz needs to be installed for the graphs to display correctly whilst running the analysis. All graphs are saved on a run-specific folder nonetheless and they can be browsed later. Runtime graphs windows can be closed without any loss.

“CleanUp.R” wipes the RStudio console and resets the global environment, i.e., clearing all variables and emptying the dataset. This ensures no leftover values are carried over from the last run of the analysis. It is assumed that the simulation will be run over multiple times by the researchers, at least while determining the initial values. Thus, starting each run fresh is necessary.

“PrepWork.R” file updates the graphic display handlers depending on the operating system (using “quartz” vs “windows” functions to open new display windows). Installation of external packages and sourcing of custom function files are carried out at this step.

The packages installed are “LongRunVol_1.0.tar.gz” from Carvalho et al. (2018) to replicate the U.S. analysis for the last 32 years whilst investigating the consequences of having a small sample size, “readxl” to read the data from Excel files, “MCMCpack” for the Bayesian MCMC analysis functions, “tseries” for the time series functions to optionally estimate the initial value of the β parameter, “pracma” for time series detrending purposes.

The custom functions file sourced consists of author-written additions to selected functions from Carvalho et al. (2018). The functions utilized from the Carvalho et al. (2018) and their purposes are provided in Table 3. Some functions are modified to account for the features of the Turkish stock returns data.

The author-written functions handle the simulation of return series for goodness of fit tests. These return series are simulated for both the constant and time-varying Σ cases.

Table 3. Custom Functions Obtained from Carvalho et al. (2018)

computeMW_PS3d	Computes the next values for the auxiliary $\boldsymbol{M}$ and $\boldsymbol{W}$ matrices.
computemC_PS3d	Computes the next values of m and C based on the M and W matrices
sampleMU_PS3d	Computes the next value of μ
sampleCOEF_MU_3d	Computes the next values of α and β
sampleCOEF_X1_3d sampleCOEF_X2_3d sampleCOEF_X3_3d	Compute the next values of $\alpha_{x_1}, \beta_{x_1}, \alpha_{x_2}, \beta_{x_2}, \alpha_{x_3}, \beta_{x_3}$
SigmaFix_5d	Computes the new Σ based on the residuals obtained from backward sampling
ComputePastorMEAN ComputePastorVAR	Computes the left and right components of predictive variance.
ComputePastorComp	Computes the decomposed components of predictive variance
computeMW_PS3d_TV	Same as “computeMW_PS3d”, modified to support a time-varying Σ matrix
sampleCOEF_MU_3d_TV	Same as “sampleCOEF_MU_3d”, modified to support a time-varying Σ matrix
DrawPriors.R	Unlike the others this is a script file containing a collection of functions used in drawing priors distributions for the variables.

A start time based on the current computer time is noted for record keeping purposes. Custom written function files are sourced, and prepackaged libraries are called to finalize the setup. The function files create and register the necessary functions that are used multiple times throughout the analysis to avoid repeating them unnecessarily within the code.

The data is imported from the “Data.xlsx” file. This file contains the observations on the annualized real and nominal log-returns along with the predictors. If the predictors are included in the analysis, the last row is dropped since the value for GDP growth for the last year would yet to be determined. Once the data is loaded, the x matrix containing the predictors is created by concatenating the columns.

User defined inputs are entered as described in the previous section. The corresponding symbolic entries are provided in Table 4.

Once the inputs are defined, "CreateSaveDir.R" file is run to create a folder that will contain the analysis results. The folder name would reflect the major properties of the current settings: whether a time-varying Σ matrix is used, predictors are included, or negative β values are allowed. The folder name also includes the starting time. Creating a separate folder for each run makes it convenient to compare the results of changes to initial settings. It also allows for an easier viewing of the graphs outside RStudio.

The first comparative variance calculated comes from the i.i.d. model. The i.i.d. model takes the return series, calculates the mean and variance of them as initial inputs to create the i.i.d. series. The iteration at each step draws new values for μ from a normal distribution with parameters updating based on the residuals obtained in the previous iteration. The residuals are obtained by taking the difference between the actual and simulated returns. The simulated returns are essentially obtained by performing random draws from a normal distribution with μ and Σ as its mean and variance, respectively. The final i.i.d. variance measure is obtained by taking the variance of all the simulated returns for each horizon. The resulting vector is in log terms, so the comparative percentage i.i.d. variance is found by taking its exponent, subtracting one and multiplying with 100.

Table 4. User-Defined Entries

$i \epsilon \{1, I\}$ where $I = 5000$ is the number of iterations.

r is the return series.

$x = [x_1, x_2, x_3]$ is the predictor matrix where x_1, x_2, x_3 are the predictor vectors.

$\bar{x} = [\bar{x}_1, \bar{x}_2, \bar{x}_3]$ is the vector of predictor means.

$t \epsilon \{1, T\}$ where $T = length(r)$ is the number of observation periods.

$$\rho_{uw} = \begin{cases} -0.5 & \text{if weak correlation assumed} \\ -0.9 & \text{if strong correlation assumed} \end{cases}$$

$m_0 = -0.25$ is the initial expected value of μ.

$C_0 = 0.25$ is the initial variance of μ.

$\sigma_\alpha = 0.0025$ is the initial variance of α.

$\beta_{min} = 0$ is the minimum cutoff value for β to ensure momentum.

The second comparative variance calculated is the unconditional variance. This is essentially a rolling windows calculation similar to the one in Chapter 3. It is actually possible to replace the unconditional variance calculated here with the one computed there. The main difference here is taking the logarithm and using only end-of-the-year data instead of daily returns.

After the i.i.d. and unconditional cases are completed and their results recorded, the predictive variance analyses for all relevant cases are performed. They all share the same structure with differences in their initial parameters. The time-varying Σ case requires additional calculations to be performed to capture the evolution of the covariance matrix across periods and iterations.

The predictive variance analysis starts by defining the coefficients and covariance matrices of μ and the predictors. Predictors can be excluded from the model by setting their coefficients zero and covariances near-zero in the corresponding matrices.

The first step is to determine the value of β. It is a crucial step since changes in β have a bearing on the final results. Running thousands of iterations approximates the true value of the β to a certain extent; however, it is not fully guaranteed since the actual dataset is barely above 30 observations. The algorithm reaches its end too soon at every iteration step to give a stable result without a good starting point, i.e., the prior for β.

Considering that the true value of β can never be known, the next best thing is to exhaust all available methods and pieces of information there are to obtain at least partial results and check whether they would agree on a similar range of values. The initial value is only part of the process: a variance for β must also be entered, so it is in practice a range of possible values to define a distribution. The algorithm will attempt to approximate the true value of β by updating the distribution at each iteration step to fit the data; however, the degree of its precision will depend on how good the prior is due to the small sample size. In a nutshell, the prior distribution supplied for β must be good enough to compensate for the smallness of the sample size to sufficiently allow the iterations to properly approximate the true value.

The value of β can be calibrated by running the analysis a few times and updating the initial value based on the difference between the prior and posterior distributions. Calibrating in this manner would likely result in overfitting the data since a second sample of returns to test against does not exist; therefore, it should not be solely relied upon. Calibration should rather be treated as partial information. For additional information, the dataset can be examined for the general behavior of expected returns. The problem here is that expected returns are not directly observable. The only clue available is the behavior of actual returns. As a proxy, a sample of expected returns can be constructed by taking the cumulative average of returns. An AR(1) analysis can then be run to determine the value of β. In this specific case, a preliminary step to detrend the sample μ series is also performed. The results from these steps are in qualitative agreement in that the β is not too high. Calibration suggests that the value of β is around 0.2 while AR(1) analysis suggests a value of 0.48. As a conservative estimate, the initial value of β is set as 0.5 and a large initial variance is entered to still allow fitting higher values if necessary. The final results under all three beliefs suggest a low β around 0.3. In order to illustrate how the outcome varies depending on different β priors, further analyses are carried out and their results are plotted for comparison. In a nutshell, higher β priors result in higher predictive variance; however, they also have worsened goodness of fit, once again signaling a low β value.

The user can choose to run the "BetaPriorElicitation.R" function to determine the initial value of β from the dataset instead of picking a value through trial and error. μ is unobservable, therefore it is hard to estimate an initial value for β. This elicitation code constructs a sample of average μ values from the data and runs separate autoregressive analyses of order one using i) μ values themselves, ii) their detrended versions, and iii) their first differenced versions. This step is not crucial by any means since it is shown during the robustness checks that the results do not change for many likely values of β; however, this option is kept here to provide the users with a starting point for their analysis when using a different dataset.

The user can enter β values manually in order to override its values for the time-varying and constant Σ cases. Once the value of β is determined, its variance is automatically calculated such that three standard deviations worth of its mass would be contained within its initial value and one, reflecting the assumption of momentum in expected returns.

$$\sigma_\beta = min\big(0.133\,,\ \ \beta_0\ \ ,\ \ (1-\beta_0)\big)$$

The initial value of α_x values for both μ and the predictors can be obtained analytically based on the β_x values. The β_x values for the predictors are determined through simulation and calibration.

$$\alpha_0 = (1-\beta_0)\bar{r}$$

$$\beta_{x_1,0} = -0.219$$

$$\alpha_{x_1,0} = \big(1-\beta_{x_1,0}\big)\bar{x}_1$$

$$\beta_{x_2,0} = 0.630$$

$$\alpha_{x_2,0} = \big(1-\beta_{x_2,0}\big)\bar{x}_2$$

$$\beta_{x_3,0} = 0.851$$

$$\alpha_{x_3,0} = (1 - \beta_{x_3,0})\bar{x}_3$$

The initial values of the diagonal elements in the covariance matrix is determined by calculating the sample variances for the return and predictor series, denoted by s^2. The second diagonal element denotes the variance of expected return which is a user determined value obtained by multiplying the variance of return series, s_r^2, with (0.1). The remaining two elements are symmetrical and denote the covariance between the error terms of current and expected return. They are found by multiplying the standard deviations of the current and expected return with the ρ_{uw} determined by the beliefs chosen beforehand.

$$\Sigma_0 = \begin{bmatrix} s_r^2 & \sqrt{0.1}s_r^2\rho_{uw} & 0 & 0 & 0 \\ \sqrt{0.1}s_r^2\rho_{uw} & (0.1)s_r^2 & 0 & 0 & 0 \\ 0 & 0 & s_{x_1,0}^2 & 0 & 0 \\ 0 & 0 & 0 & s_{x_2,0}^2 & 0 \\ 0 & 0 & 0 & 0 & s_{x_3,0}^2 \end{bmatrix}$$

The auxiliary φ matrix is found by Cholesky decomposing the Σ, solving the resulting linear system and carrying out some matrix multiplications. It is mainly used in calculating the next value of the Σ matrix based on the residuals.

m and C variables denote the expected value and covariance of μ at a given iteration. M and W are auxiliary variables used for calculating the values of m and C for the next step of the iteration based on the current Σ matrix. The next values of m and C are found by running the “computeMW_PS3d” and “computemC_PS3d” functions obtained from Carvalho et al. (2018). These functions also calculate the next values of M and W.

The prior distribution values are computed by running the “DrawPriors.R” script file provided by Carvalho et al. (2018). This file contains multiple functions and some code that calculates the necessary auxiliary variables along with the prior values of the predictive variance across the investment horizon. The author-written “SimulateReturnSeries” function constructs thousands of return series based on the parameters defining the prior distributions. The associated variance values are calculated, plotted, and recorded for each investment length. These values are mainly used during the goodness of fit tests.

The system is now ready to run the main body of the algorithm. The iteration steps are calculated thousands of times. Each iteration consists of a forward filter and a backwards sampling step. The values and distributions obtained from one iteration are used in determining the initial values of the next.

Each iteration starts with the forward filter sub step. The forward filter is run across the length of return observations. It takes m and C values of this period as inputs and calculates the auxiliary variables M and W based on the Σ matrix which in turn are used in conjunction with that period’s return data to calculate the next period values of m and C. As a result, the series for m and C are obtained for each observation period. The functions used here are “computeMW_PS3d” and “computemC_PS3d”.

Once forward iteration is complete, and the final period values are reached, the backwards sampling step starts by making a random draw for the error term of the final μ value. Based on this final value of μ and the previously obtained M and W series, the μ values for remaining periods are sampled by iterating backwards across periods and evaluating the equations using the “sampleMU_PS3d” function. The next values for the coefficients of μ are obtained through the "sampleCOEF_MU_3d" function. The new β value is chosen such that it is greater than β_{min} to ensure momentum. Predictor coefficients are calculated in the same manner if they are included in the system.

By taking the differences across the observed and simulated values of the equations in the VAR system, the residual terms ε_i can be obtained. These residual terms are used in determining the new Σ matrix for the next step of the iteration.

$$\varepsilon_{i,t} = \begin{cases} r_{t+1} - \mu_{i,t} \\ \mu_{t+1} - \alpha_i - \beta_i \mu_t \\ x_{1,t+1} - \alpha_{x_1,t} - \beta_{x_1,t} x_{1,t} \\ x_{2,t+1} - \alpha_{x_2,t} - \beta_{x_2,t} x_{2,t} \\ x_{3,t+1} - \alpha_{x_{31},t} - \beta_{x_3,t} x_{3,t} \end{cases}$$

The values for each relevant variable are recorded for the current iteration for future reference. These need not all be included in the calculations in the following iterations. Due to the simulated nature of the Bayesian MCMC algorithm, anomalies may occur. Having well-kept records in such cases allows users to identify potential issues more conveniently.

Predictive variance parts for the current iteration step corresponding to each observation period are calculated and recorded before moving on to the next iteration. This step concludes the main loop of the algorithm. These calculations are performed for each iteration.

The Bayesian MCMC portion of the analysis finishes once all iterations are completed. Before continuing with the calculations, the first 40% of the iteration results are excluded from the system, since the simulation is thought to be still approximating during those iterations. The predictive variance parts corresponding to the remaining iterations are evaluated for each observation period across all iteration steps. The five components of predictive variance are calculated and recorded in the same way.

Before predictive variance analysis concludes, the priors and posteriors are reported as boxplots for comparison along with histograms. In addition, the predictive variance decomposition is plotted at this step. Once the predictive analysis is finished, the results from all steps are collected and plotted in the final graph for comparison. The final step simulates the returns for goodness of fit tests and plots them. This concludes a single run of the algorithm. “SaveImage.r” script would record the relevant variables in the same folder as the graphs for further examination.

Multiple runs of the algorithm may be necessary if the user is still calibrating the initial values or performing robustness check tests by changing the initial values. In the analysis performed here, multiple runs are performed for various β values to make sure the final results are not sensitive to its initial value. Bayesian MCMC would yield slightly different results at each run of the algorithm. The results should be similar if under normal circumstances. If not, the saved files can be examined for trouble shooting. The algorithm starts fresh at each run by clearing the global environment of all variables to prevent values obtained from different runs of the algorithm affect each other.

This concludes the explanations of the estimation algorithm used here. For any further questions, the readers are encouraged to contact the author.

CHAPTER 6

CONCLUSION

Both analyses performed here suggest that it is more preferrable for an investor to make long term investments in stocks. In the unconditional analysis chapter, it is found that the unconditional variance decreases while the mean returns are relatively constant as the investment horizon lengthens. This suggests that since all horizons have similar mean returns for the investor, the unconditional variance becomes the sole determinant of the investment decision. Investors using the unconditional variance as their risk estimate would prefer to make long-term investments. Furthermore, it becomes less likely for an investor to lose money in the long run. Sample probabilities suggest that defeating a set of increasing fixed interest rates becomes more probable as time passes. In the real returns case, beating any of the selected interest rates in the short run is a coin toss with fifty-fifty chances. In the long run, it becomes more likely to earn respectively 1%, 2%, and 3% returns while the chances of beating 4% and 5% drop. This suggest that in terms of real returns, investing in the long run becomes less of a gamble and more of an investment guaranteeing an increased, albeit capped, return.

The predictive variance analysis results for the most part echo those of the unconditional variance analysis. Investors prefer to make long-term investments in all of the cases tested. Stronger negative correlation between the error terms of return and expected return results in reduced risks as a result of mean-reversion dominating momentum in the long-run. Assuming that the Σ matrix is time-varying results in a reduced i.i.d. uncertainty in the model with mean-reversion levels similar to the weak negative correlation case, despite having a higher negative correlation.

In terms of goodness of fit, the strong ρ_{uw} scenario provides the best fit in nominal returns case. With real returns, it is hard to draw solid conclusions as the humped shape of the unconditional variance curve hinders confidence interval tests. With no adjustments or assumptions, the weak ρ_{uw} case is the winner; however, it can be argued that for the long-run investments time-varying Σ and strong ρ_{uw} scenarios provide better fits with little explanatory power for the medium-term investments.

Including predictors in the system does not result in meaningful changes, signaling that the dynamics created by the momentum and mean-reversion effects adequately model the patterns observed in the data.

The results are not very sensitive to changing the β values unless it is close to one; however, there is no evidence supporting a high β value. Instead, the β is estimated to be around 0.5. Out of all the scenarios, the time-varying Σ is the most robust to changing the initial β value.

Allowing a negative β produces unexpected results: the results are nearly identical. Both the weak and strong ρ_{uw} lines approach the time-varying Σ line which shows no change in its behavior. This suggests that negative β values may have the potential to assume the role provided by mean-reversion and impose a combined effect on the model by itself. As a byproduct, the goodness of fit problem in real returns case due to the humped unconditional variance curve is solved for the time-varying case: the null hypothesis that it adequately captures the behavior present in the data cannot be rejected.

The small sample size presents an obstacle since it limits the number of iterations that can be performed in the forward iteration and backwards sampling processes contained within each looping step of the algorithm. The tests performed on the U.S. data suggests that smaller sample sizes tend to overestimate the predictive variance. Taking this as evidence, it can be proposed that the results obtained here are indicative of a much higher long-term preferability level and this proposal can be tested in the future as more data becomes available to test for longer horizons.

Combining all of the findings, it can be concluded that it is more beneficial for an investor to make long-term investments in Borsa Istanbul. Depending on their individual beliefs about the model features and preferences, the exact proportions they allocate to stocks may change; however, the main lesson remains the same.

6.1 Suggestions for further research

Considering that there are only 32 observations in terms of annual return data, reperforming the same analysis in a few years would be enlightening on the effects of the sample size. In addition, the unconditional variance line for the real returns case in predictive variance analysis is likely to become smoother with more data, enabling more accurate goodness of fit tests.

Other predictors may be controlled for when running the analysis to further test whether the dynamics between the momentum and mean-reversion indeed sufficiently explain the behavior in the data. If such variables are found, their effects can be easily captured within the existing system since the Σ matrix consists of 2+k sized square matrix where k is the number of predictors. As argued before, inclusion of predictors is not necessary for the main lessons of this thesis to hold; however, statistically significant predictors would allow more precise estimations for the predictive variance. This new predictive variance would be lower, further improving the arguments made here. It is left to future studies to determine whether a set of such predictors exist to maintain the scope here.

The case for negative β presents an interesting avenue for further research since it seems to combine both the momentum and mean-reversion effects in a compact mechanism through itself. This is evidenced in the way weak and strong ρ_{uw} lines approach the time-varying Σ line. The negative β analysis carried out here is essentially a thought experiment relaxing the momentum assumption as a sensitivity check. Whether it is a plausible assumption and holds any merits can only be determined with future work on the subject.

Lastly, relaxing other assumptions of the model may be considered. For example, allowing for time-varying Σ values in recent years has introduced many improvements to Bayesian MCMC estimations. One wonders if more contributions can be made by introducing more features or removing the restrictions placed by the assumptions made in the literature.

REFERENCES

Akbaş, Y.E. (2013). The analysis of relationship between the rate of stock return and interest rate with nonlinear methods: The case of Turkey, *Business and Economics Research Journal, 4*(3), 1-21.

Barberis, N. (2000). Investing for the long run when returns are predictable. *The Journal of Finance, 55*(1), 225-264. doi.org/10.1111/0022-1082.00205

Bajo-Rubio, O., Berke, B., & McMillan, D. (2017). The behaviour of asset return and volatility spillovers in Turkey: A tale of two crises. *Research in International Business and Finance, 41*, 577-589. doi.org/10.1016/j.ribaf.2017.04.003

Bodie, Z. (1995). On the risk of stocks in the long run. *Financial Analysts Journal, 51*(3), 18-22. doi.org/10.2139/ssrn.5771

Bodie, Z., Merton, R. C., & Samuelson, W. F. (1992). Labor supply flexibility and portfolio choice in a life cycle model. *Journal of economic dynamics and control, 16*(3-4), 427-449. doi.org/10.1016/0165-1889(92)90044-F

Carvalho, C. M., Lopes, H. F., & McCulloch, R. E. (2018). On the long-run volatility of stocks. *Journal of the American Statistical Association, 113*(523), 1050-1069. doi.org/10.1080/01621459.2017.1407769

Chib, S., Nardari, F., & Shephard, N. (2002). Markov chain Monte Carlo methods for stochastic volatility models. *Journal of Econometrics, 108*(2), 281-316. doi.org/10.1016/S0304-4076(01)00137-3

Cochrane, J. H. (2009). *Asset pricing: Revised edition*. Princeton, NJ: Princeton University Press.

Daniels, M. J., & Pourahmadi, M. (2002). Bayesian analysis of covariance matrices and dynamic models for longitudinal data. *Biometrika, 89*(3), 553-566. doi.org/10.1093/biomet/89.3.553

Ferson, W. E., Sarkissian, S., & Simin, T. T. (2003). Spurious regressions in financial economics?. *The Journal of Finance, 58*(4), 1393-1413. doi.org/10.1111/1540-6261.00571

Fisher, K. L., & Statman, M. (1999). A behavioral framework for time diversification. *Financial Analysts Journal, 55*(3), 88-97. doi.org/10.2469/faj.v55.n3.2275

Gelman, A., Meng, X. L., & Stern, H. (1996). Posterior predictive assessment of model fitness via realized discrepancies. *Statistica Sinica, 6*(4), 733-760.

Jacquier, E., Polson, N. G., & Rossi, P. E. (2002). Bayesian analysis of stochastic volatility models. *Journal of Business & Economic Statistics, 20*(1), 69-87. doi.org/10.2307/1392199

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

Kandır, S. Y. (2008). Macroeconomic variables, firm characteristics and stock returns: Evidence from Turkey. *International Research Journal of Finance and Economics, 16*(1), 35-45.

Kasman, A., & Torun, E. (2007). Long memory in the Turkish stock market return and volatility. *Central Bank Review, 2*(2), 13-27.

Kasman, S., Vardar, G., & Tunç, G. (2011). The impact of interest rate and exchange rate volatility on banks' stock returns and volatility: Evidence from Turkey. *Economic Modelling, 28*(3), 1328-1334. doi.org/10.1016/j.econmod.2011.01.015

Kritzman, M. (1994). What practitioners need to know... about time diversification (corrected). *Financial Analysts Journal, 50*(1), 14-18. doi.org/10.2469/faj.v71.n1.4

LeRoy, S. F., & Porter, R. D. (1981). The present-value relation: Tests based on implied variance bounds. *Econometrica: Journal of the Econometric Society, 49*(3), 555-574. doi.org/10.2307/1911512

Li, D., Nguyen, Q. N., Pham, P. K., & Wei, S. X. (2011). Large foreign ownership and firm-level stock return volatility in emerging markets. *Journal of Financial and Quantitative Analysis, 46(*4), 1127-1155. doi.org/10.1017/S0022109011000202

Pastor, L., & Stambaugh, R. F. (2009). Predictive systems: Living with imperfect predictors. *The Journal of Finance, 64*(4), 1583-1628. doi.org/10.1111/j.1540-6261.2009.01474.x

Pastor, L., & Stambaugh, R. F. (2012). Are stocks really less volatile in the long run?. *The Journal of Finance, 67*(2), 431-478. doi.org/10.1111/j.1540-6261.2012.01722.x

Prado, R., Ferreira, M. A., & West, M. (2021). *Time series: Modeling, computation, and inference*. Boca Raton, FL: CRC Press.

Reichenstein, W., & Dorsett, D. (1995). Time diversification revisited. *The Research Foundation of The Institute of Chartered Financial Analysts, 46*.

Rjoub, H., Türsoy, T., & Günsel, N. (2009), The effects of macroeconomic factors on stock returns: Istanbul Stock Market, *Studies in Economics and Finance 26*(1).
doi.org/10.1108/10867370910946315

Rubin, D. B. (1984). Bayesianly justifiable and relevant frequency calculations for the applied statistician, *The Annals of Statistics, 12*(4), 1151–1172.
doi.org/10.1214/aos/1176346785

Samuelson, P. (1994). The long-term case for equities - and how it can be oversold, *Journal of Portfolio Management*, *21*(1), 17-18.
doi.org/10.3905/jpm.1994.409499

Shiller, R. J. (1981). Alternative tests of rational expectations models: The case of the term structure. *Journal of Econometrics*, *16*(1), 71-87.

Siegel, J. (2008). *Stocks for the long run*. New York: McGraw-Hill.

Sönmez, F., (2007) Does inflation have an impact on conditional stock market volatility?: Evidence from Turkey and Canada, *International Research Journal of Finance and Economics, 11*, 123-133.

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

Sum, V. (2012), Economic policy uncertainty and stock market performance: evidence from the European Union, Croatia, Norway, Russia, Switzerland, Turkey and Ukraine, *Journal of Money, Investment and Banking, 25*, 99-104
doi.org/10.2139/ssrn.2094175

Tokat, Y., & Stockton, K. A. (2006). A primer on tactical asset allocation: strategy evaluation. Valley Forge, Pa.: Vanguard Investment Counseling & Research, The Vanguard Group.

Toparlı, E. A., Çatık, A. N., & Balcılar, M. (2019), The impact of oil prices on the stock returns in Turkey: A TVP-VAR Approach, *Physica A: Statistical Mechanics and its Applications*, 535
doi.org/10.1016/j.physa.2019.122392

Umutlu, M., Akdeniz, L. & Altay-Salih, A. (2010), The degree of financial liberalization and aggregated stock-return volatility in emerging markets, *Journal of Banking & Finance*, *34*(3), 509-521.
doi.org/10.1016/j.jbankfin.2009.08.010

Vardar, G., & Aydogan, B. (2019), Return and volatility spillovers between Bitcoin and other asset classes in Turkey: Evidence from VAR-BEKK-GARCH approach, *EuroMed Journal of Business*, *14*(3). doi.org/10.1108/EMJB-10-2018-0066

Welch, I., & Goyal, A. (2008). A comprehensive look at the empirical performance of equity premium prediction, *Review of Financial Studies, 21*(4), 1455-1508. doi.org/10.1093/rfs/hhm014

West, M., & Harrison, J. (1997). Bayesian forecasting and dynamic models (2nd ed.). New York: Springer. doi.org/10.1007/b98971