Turun yliopisto
University of Turku

PORTFOLIO SELECTION UNDER DIRECTIONAL RETURN PREDICTABILITY

Joonas Hämäläinen

Turun kauppakorkeakoulu
Turku School of Economics

PORTFOLIO SELECTION UNDER DIRECTIONAL RETURN PREDICTABILITY

Joonas Hämäläinen

Sarja/Series A-10:2015

Turun kauppakorkeakoulu
Turku School of Economics

Custos: Professor Mika Vaihekoski
Turku School of Economics

Supervisors: Professor Mika Vaihekoski
Turku School of Economics

Professor Luis Alvarez
Turku School of Economics

Pre-examiners: Docent Hannu Kahra
University of Oulu

Professor Johan Knif
Hanken School of Economics

Opponent: Professor Johan Knif
Hanken School of Economics

The originality of this thesis has been checked in accordance with the University of Turku quality assurance system using the Turnitin OriginalityCheck service.

ISBN 978-952-249-445-0 (print) 978-952-249-446-7 (PDF)
ISSN 0357-4652 (print) 1459-4870 (PDF)

Publications of Turku School of Economics, Series A

Suomen yliopistopaino Oy – Juvenes Print, Turku 2015

ABSTRACT

Traditional investment portfolio optimization aims at maximizing the expected return and minimizing the variance of the portfolio. This approach generally requires estimation of expected returns on assets, which has been a difficult task in practice. Instead of estimates for expected returns, investors typically have directional forecasts, that is, views on whether the price of an asset is going to go up or down in the future. Incorporating directional views into portfolio selection has not been explored rigorously before.

This dissertation introduces a new theoretical framework that allows incorporating directional views into mean-variance portfolio optimization. The implications of this model are explored analytically and with computer simulations. Empirical studies and trading simulations contained in the dissertation are conducted utilizing recent stock market data.

The developed model indicates that when directional forecasts are reliable, investors prefer assets with high volatility for higher expected returns. Negative correlation between asset returns is not preferred, assuming that it is linked to correlated forecasts. In addition, correlation between absolute values of returns matters explicitly for portfolio variance.

In addition to providing a solution to portfolio selection with directional forecasts, the developed framework can increase portfolio performance in terms of Sharpe ratios compared to simpler alternative models. Care should be taken when estimating the accuracies of the directional views, in order to benefit from the model in practical investment management.

Keywords: portfolio choice, investment decisions, diversification, security markets, stock returns

TIIVISTELMÄ

Perinteinen sijoitusportfolion optimointi pyrkii maksimoimaan portfolion odotetun tuoton ja minimoimaan tuoton varianssin. Tämä lähestymistapa vaatii tyypillisesti sijoituskohteiden odotettujen tuottojen estimointia, mikä on osoittautunut vaikeaksi tehtäväksi käytännössä. Sen sijaan, että sijoittajat estimoisivat tuottojen odotusarvoja, heillä on tyypillisesti näkemys sijoituskohteiden tuottojen suunnasta, eli siitä, nouseeko vai laskeeko sijoituskohteen hinta tulevaisuudessa. Suuntanäkemysten sisällyttämistä portfolion valintaan ei ole tutkittu aikaisemmin.

Tämä väitöskirja esittelee uuden teoreettisen mallin, joka mahdollistaa suuntanäkemysten käyttämisen odotusarvo-varianssi-pohjaisessa sijoitusportfolion optimoinnissa. Mallin implikaatioita tarkastellaan analyyttisesti sekä tietokoneella tehtävien simulaatioiden avulla. Väitöskirjan sisältämissä empiirisissä tutkimuksissa sekä kaupankäyntisimulaatioissa käytetään viimeaikaisia havaintoja osakemarkkinoilta.

Malli osoittaa, että luotettavien suuntaennusteiden tapauksessa korkean volatiliteetin kohteita tulisi preferoida korkeamman odotetun tuoton saavuttamiseksi. Negatiivinen korrelaatio kohteiden tuottojen välillä ei ole hyödyllinen ominaisuus olettaen, että se on kytköksissä korreloituneisiin suuntaennusteisiin. Lisäksi sijoituskohteiden itseisarvotuottojen välisellä korrelaatiolla on suora vaikutus portfolion varianssiin.

Kehitetty malli tarjoaa ratkaisun portfolion muodostamiseen suuntaennusteiden valossa, ja sen avulla voidaan saavuttaa muun muassa korkeampia Sharpen suhdelukuja verrattuna vaihtoehtoisiin malleihin. Suuntanäkemysten tarkkuuden estimointi vaatii tulosten perusteella huolellisuutta, jotta mallista voidaan hyötyä käytännön salkunhoidossa.

Asiasanat: sijoitusportfolion valinta, sijoituspäätökset, hajauttaminen, arvopaperimarkkinat, osaketuotot

ACKNOWLEDGEMENTS

I am immensely thankful to a lot of people. First, my supervisors at Turku School of Economics – Professor Mika Vaihekoski and Professor Luis Alvarez – from the beginning, you believed in me, took me seriously, treated me more like a colleague than a student, and supported me throughout this project. I would especially like to express my gratitude to Mika: it was in your lectures about asset pricing and portfolio theory that I got the spark to study finance more seriously. Your office door was always open and you had time for questions. I can say with full confidence that without you, I never would have started this project.

Second, I would like to thank the Department of Accounting and Finance at Turku School of Economics for providing me with a secure and motivating position throughout the doctoral studies. I am grateful for the rather big responsibility of teaching in several courses that were suited for my interests and helped develop my skills as a teacher. I would especially like to thank the head of the department, Professor Kari Lukka, and Professor Hannu Schadewitz for being supportive throughout my studies. I would also like to thank Dr. Mikko Kepsu for acting as a mentor for me and introducing me to the practicalities in the world of academia. Everyone else affiliated with the department – fellow doctoral students, more experienced colleagues, administrative staff – I feel privileged to be part of such a friendly work environment.

Third, I would like to thank the entire staff affiliated with the Graduate School of Finance (GSF) and especially the director Dr. Mikko Leppämäki for making it possible to study finance at such a high level in Finland. This is truly an exceptional program with world-class lecturers flying in to teach a group of students in Helsinki. Without the demanding coursework of GSF, the competence of many doctoral students in finance, including myself, would be considerably lower.

Fourth, I would like to express my gratitude to the pre-examiners of this dissertation, Professor Johan Knif, and Docent Hannu Kahra. Your comments improved the quality of the work and provided helpful feedback that was useful in finalizing the dissertation. In addition, a number of other people have acted as discussants for the essays in this dissertation, and provided valuable feedback in conferences and seminars. Their names are listed in the cover page footnotes of each essay. I am grateful for all the feedback and time you have invested.

During the doctoral studies, I had a chance to visit a world-class university abroad. For this, I would like to thank the Finnish Fulbright Center, Princeton

University and Professor John Mulvey, who invited me to spend 12 months as a visiting student under his supervision. Moreover, the research grants from Jenny and Antti Wihuri Foundation, Nordea Bank Foundation, OP-Pohjola Group Research Foundation, Finnish Foundation for Advancement of Securities Markets, Turun kauppaopetussäätiö, Emil Aaltonen Foundation, TOP-Säätiö, Marcus Wallenberg Foundation, and Vihanto Oy Fund, have been a great help both for arranging the visit abroad and for supporting my studies here in Finland. This support from the foundations is gratefully acknowledged. For covering the expenses arising from the publication of this dissertation, I thank the Turku School of Economics Association.

Last but not least, I would like to thank my family for supporting me throughout my studies and believing that this endeavor was worth undertaking. I am especially grateful to you for always letting me choose my own path in life and explore new things, no matter how unordinary they may have been.

Joonas Hämäläinen
Turku, 3.8.2015

TABLE OF CONTENTS

LIST OF ORIGINAL RESEARCH PAPERS

(1) Hämäläinen, Joonas – Portfolio selection with directional return estimates, preprint.

(2) Hämäläinen, Joonas – Correlation in the magnitude of financial returns, preprint.

(3) Hämäläinen, Joonas – Predictable returns and portfolio optimization: Directional versus whole return forecasts as inputs, preprint.

Part I

SYNTHESIS

1 INTRODUCTION

1.1 Motivation

Modern investment portfolio optimization originates from the work of Markowitz (1952, 1959), who introduced the concept of maximizing the expected return and minimizing the variance of a portfolio[1]. Tobin (1958) complemented Markowitz by adding a riskless asset to arrive at the separation theorem, i.e. that all investors should hold the risky assets in their portfolios in same proportions regardless of their preferences. These concepts were major breakthroughs in the field of finance and have since attracted an enormous amount of attention from both academic researchers and professionals working in the industry. Markowitz's seminal paper about portfolio selection published in 1952 and the subsequent book published in 1959 are among the most cited[2] works in the history of financial research.

Nevertheless, the application of Markowitz's framework in practice has not been a major success, mainly due to its implicit requirements of accurate input estimates. Estimation error in the inputs can potentially render the optimization procedure infeasible, and simple equal-weighted portfolios can provide similar performance, or even outperform mean-variance optimized portfolios (see e.g. Michaud (1989), DeMiguel, Garlappi and Uppal (2009)). Specifically, estimating future mean returns – a key input class in the Markowitz model – is known to be notoriously difficult in practice. To some extent, financial returns are predictable (for a recent review about stock return predictability, see Rapach and Zhou (2013)), however, estimates of expected returns remain noisy and present a problem for portfolio optimization in practice.

On the other hand, previous research has shown that the mere *directions* of returns can be predicted out-of-sample, often with significant accuracy. Notable recent studies documenting this directional predictability include Bekiros (2010a), Nyberg (2011), and Chevapatrakul (2013). In essence, leaving out the magnitude component from return estimates considerably simplifies the task of the forecaster. Directional predictability can also be considered theoretically sound amidst near unforecastable mean returns: Christoffersen and Diebold (2006) show that conditional return signs can exhibit dependence even if the mean returns are independent (i.e. not forecastable).

---

[1] To be accurate, Roy (1952) introduced a similar concept, however, Markowitz is often recognized as the father of modern portfolio theory.

[2] Google Scholar gives a combined estimate of over 32 000 citations as of July 2015.

Perhaps more importantly, the behavior of investors in practice resembles that of making directional forecasts on asset returns. For example, simple buy/sell recommendations are exceedingly common in analyst reports, and in general, it appears that investors often have views only on the directions, and not the mean returns[3]. There is evidence that when mean forecasts are provided by professionals, they may contain profitable information linked to directional accuracy, however, by traditional measures such as the root-mean-squared error, their accuracy can be poor (Leitch and Tanner (1991)). Estimating the mean return of an asset can be considered a near impossible task, and therefore is often avoided. The most prominent forecasts appear to be about the directions, or signs, of asset returns.

The two premises introduced above – evidence of ouf-of-sample directional predictability, and typical investor behavior in expressing views on future returns – form the foundation for the research in this dissertation. While directional predictability cannot be directly compared to the forecastability of mean returns, it could be said that directional forecasting is easier than forecasting the mean returns, and the behavior of individuals and professionals in the industry appears to support this view. Faced with these premises, it is natural to ask how optimal portfolios should be formed when the investor has directional estimates for future returns. An answer to this question has not been explored, to the author's knowledge, in the past.

A directional estimate approach is not only natural from the investor's standpoint, but also presents a possible solution to the aforementioned portfolio optimization problem: Leaving out the direct estimation of mean returns and instead focusing on the direction of future returns is a way of circumventing the main problem present in expected return estimation that has been plaguing portfolio optimization over the past decades.

1.2 Research objectives, methodology, and structure

The dissertation aims to provide a readily applicable solution to forming optimal portfolios with directional return estimates, and to examine the implications that follow. First, a framework to incorporate directional estimates into mean-variance optimization needs to be developed. Second, the proposed solution will affect the behavior of a rational investor and portfolio performance compared to alternative solutions. These implications need to be explored, as is the case with any new theoretical work. These are the two main research objectives of the dissertation.

The research methodology employed can be divided into two parts: First,

[3] An example of directional views among individual investors is the weekly sentiment survey of the American Association of Individual Investors conducted since 1987.

two out of the three essays in the dissertation contain a theoretical section where results are derived by the way of analytical reasoning and mathematical notation. Second, the implications of analytical results are examined in simulation studies, utilizing both artificially generated data and authentic stock market data. This combination of theory, simulation and empirical analysis is a natural way to conduct the studies and to answer the research questions presented in each essay.

All computer simulations and empirical analyses are conducted using the programming language R and several additional packages that are available for the core program. Of these packages, the most heavily utilized and worth mentioning are the *nloptr* package by Jelmer Ypma, which facilitates nonlinear optimization with constraints, and the *doParallel* package by Steve Weston & Revolution Analytics, which employs parallel computation for faster problem solving. In addition, Wolfram Mathematica is employed for solving complicated problems analytically.

The data in the empirical sections of the essays come from Thomson Reuters Datastream. The main data in all studies are from the U.S. stock market during the 21st century, in order cater for the widest audience and to focus on the recent history of the financial markets. The use of data is described and reasoned in more detail in each of the essays separately.

The dissertation is comprised of three essays. Each of these essays answers a specific set of research questions and they are presented in chronological order as the subsequent studies build on the results presented earlier. In the first essay, a novel framework for portfolio optimization is introduced by decomposing asset returns into three components. Utilizing this decomposition, analytical results for conditional mean, variance and covariance between asset returns are derived. These results can readily be applied in portfolio optimization, and the essay examines what types of assets are preferred by the investor under directional return predictability. The essay also compares portfolio performance using the proposed framework against simpler alternatives.

The second essay of the dissertation examines correlation between the absolute values of asset returns (hereafter, magnitude correlation). This term emerges from the analytical results in the first essay and can have an impact on conditional portfolio variance. Since this correlation is a relatively unknown concept in financial research, the second essay aims to fill the gap in knowledge. The essay establishes an analytical link between Pearson return correlation and magnitude correlation under bivariate normality. The empirical properties of magnitude correlation in the U.S. stock market are examined in a comprehensive fashion, and an interpretation for magnitude correlation in financial context is presented. Finally, the second essay explores methods for out-of-sample estimation of this correlation from a portfolio management perspective.

The third essay aims to test the framework proposed in the first essay in out-of-sample trading simulations. Moreover, the specific issue of estimating directional accuracy is addressed. The third essay also asks whether the investor would benefit from using noisy information about mean returns instead of simply extracting the mere signs of mean return forecasts for portfolio optimization. This is evaluated by comparing the proposed framework against alternative models utilizing mean return forecasts with varying accuracy.

The contribution of this synthesis is to collect the main findings of the essays and place the information into proper context in the research field. In order to accomplish this, related previous research is examined in more detail in the second chapter. The main findings of the essays are summarized in chapter three. Finally, chapter four highlights the key results of the essays, presents future research opportunities, and offers concluding remarks.

2 PREVIOUS RESEARCH AND POSITIONING OF THE DISSERTATION

A significant contribution of the dissertation is to develop a new framework for portfolio optimization in order to utilize directional return estimates as inputs. The concepts and research questions that follow are built on this proposed framework, which is introduced in the first essay. As such, the dissertation steps partly into new grounds with no immediate base of previous research. The study builds on the ability to predict the directions of asset returns, and therefore, this area and its notable previous studies are explored in more detail. Previous research on traditional mean-variance optimization and some relevant extensions are also briefly surveyed in order to position the dissertation in light of the portfolio optimization literature.

2.1 Directional return predictability

Directional return predictability equates to an ability to forecast the direction, or sign, of price changes with some accuracy. The direction can take only one of the two discrete values: up (+1) or down (-1). As such, these type of forecasts are by nature much simpler and arguably easier to make than providing forecasts of future mean returns. Moreover, Christoffersen and Diebold (2006) show that theoretically, even if conditional mean returns are independent (i.e. not forecastable), there can still be sign dependence in the returns. Thus, even if mean returns could not be forecasted, it does not rule out directional predictability.

Christoffersen and Diebold (2006) focus on directional predictability arising from the distribution itself. However, in reality, forecasts can be derived from any source of information, not just the return distribution. Empirically, out-of-sample directional accuracy of forecasts varies depending on the investment horizon and the model used to produce forecasts. For monthly stock returns, out-of-sample accuracy[4] can range between 55% and 65% (e.g. Pesaran and Timmermann (2002), Chevapatrakul (2013)), and for monthly market index returns, correct sign forecasts up to 60% have been documented (e.g. Pesaran and Timmermann (1995), Leung, Daouk and Chen (2000), and Nyberg (2011)).

For weekly market index returns, Bekiros and Georgoutsos (2008b) report directional accuracies ranging between 52% and 59% (depending on the forecasting model). For daily stock returns, Skabar (2013) achieved a directional

[4] The accuracy referred to is the percentage of time the forecasting model produces the correct sign.

accuracy of over 52%. Finally, for daily index returns, a directional accuracy close to 55% can be achieved (see e.g. Bekiros and Georgoutsos (2007), Bekiros and Georgoutsos (2008a), Bekiros (2010a), Bekiros (2010b)). These figures are notable and imply that the forecasting models could be used to generate economically significant profits.

Stocks are not the only asset class to exhibit notable directional predictability. Currency exchange rates have been shown to be directionally forecastable as well: Elliott and Ito (1999) document monthly directional accuracy of nearly 55%. Evidence of forecastability in the daily interval is offered in Kuan and Liu (1995). More recent studies about directional predictability of currency exchange rates worth mentioning are Chung and Hong (2007) and Preminger and Franck (2007). The portfolio selection framework developed in this dissertation can naturally be applied to investments in any asset class. Stock return predictability has the most research available, and therefore results in that area are emphasized.

The dissertation does not attempt to take part in the discussion of which forecasting models work best in generating accurate forecasts. Instead, tools and concepts for portfolio optimization are presented when the investor has directional estimates on future returns available – any type of directional forecasting model can be used to arrive at this setting. Since profitable trading strategies or forecasting models have a lot of their value basing on secrecy, it is reasonable to expect that most working models are not published in academic research. Therefore, the models and accuracies presented in the aforementioned studies may well be just the tip of the iceberg. This gives even stronger motivation for the research conducted in the dissertation.

The behavior of professional analysts or investment managers can tell something about the other side of the table: Investment professionals often express directional views, such as buy/sell recommendations, instead of disclosing information about mean returns. It appears that estimating future mean returns is a very difficult task, which many tend to avoid. Moreover, there is evidence that return forecasts evaluated by standard error measurement tools such as the root-mean-squared error can be relatively poor, whereas at the same time they can be profitable, which can be linked to the directional accuracy of these forecasts (Leitch and Tanner (1991)). Taylor (1980) has developed a model to forecast the trend of price changes. His results indicate that even though the accuracy in terms of mean squared error can be low, the directional accuracy of the forecasts can be significant.

Womack (1996) and Barber, Lehavy, McNichols and Trueman (2001) have examined analysts' buy/sell recommendations and the following stock price development. According to these studies, there appears to be value in the discrete or directional forecasts made by analysts. These studies on the practical side of the financial industry speak for directional accuracy of forecasts, although less directly than the aforementioned research that documents out-of-sample predictability more concretely.

Overall, the combination of the documented out-of-sample accuracy and the behavior of professionals in the field tend to indicate that directional estimates of asset returns are the most prominent type of forecasts. Therefore, there should be strong demand for a model that incorporates directional estimates into investment decision making in the form of portfolio optimization.

2.2 Portfolio optimization

Traditional investment portfolio optimization is based on a mean-variance tradeoff as introduced in Markowitz (1952). Tobin (1958) later complemented the result by adding a risk-free rate in the analysis, arriving at the separation theorem, i.e. that all investors should hold risky assets in the same proportions in their portfolio regardless of their preferences. However, portfolio optimization has come a long way since then and new extensions have been produced at a steady rate. For a recent review of some of the advancements in the field of portfolio optimization, see Kolm, Tütüncü and Fabozzi (2014).

The portfolio optimization framework which this dissertation develops takes the approach of Markowitz (1952) as the mainframe and examines portfolio formation in the mean-variance universe. The proposed framework produces conditional estimates for the expected return vector and the variance-covariance matrix. These conditional estimates are derived by decomposing asset returns into three components: the forecasted direction, the outcome of the directional forecast, and the magnitude of the return. This makes it possible to bypass direct estimation of the mean return and instead break down the estimation into simpler parts. As a result, directional forecasts can be used as inputs in mean-variance optimization. To the best knowledge of the author, this type of an approach in portfolio selection is original.

Earlier research on related topics include the studies about market timing and portfolio selection by Jensen (1972), Grant (1978), Pfeifer (1985), and, more recently, Hallerbach (2014). Market timing generally deals with a simple setting where the investor chooses between a risky investment and a riskless asset, attempting to time the market. While this approach has an element of similarity to the proposed framework in this dissertation, the latter is much more complex and complete in the sense that it describes the conditional comovement between different assets and allows portfolio optimization with any number of assets to be performed.

Portfolio optimization literature in related areas also contains studies about how predictable returns affect portfolio choice in general. For example, when a regression model with predictive variables is utilized for stock returns, Kandel and Stambaugh (1996) provide solutions for asset allocation. Barberis (2000), and Campbell, Chan and Viceira (2003) examine the implications of return predictability for investors with long investment horizons. Ait-Sahalia and Brandt (2001) incorporate predictive variable effects into portfolio weights in a direct fashion. Ang and Bekaert (2002), Guidolin and Timmermann (2007) and Tu (2010) study portfolio selection in a regime switching environment. While these approaches are loosely related to the topic of this dissertation and worth mentioning here, they are still very different from the approach taken in this dissertation.

Elton and Gruber (1987) study portfolio choice rules when assets are grouped into expected return categories. Perhaps closest to the framework proposed in the dissertation comes the work of Black and Litterman (1992), where views (on a continuous scale) about asset returns are incorporated into portfolio optimization by blending them with a prior of the expected return vector. The variance-covariance matrix is not addressed in Black and Litterman (1992), but Qian and Gorman (2001) offer an extension for this. However, the portfolio optimization framework presented in this dissertation is fundamentally different as it focuses on *directional* views, and it does not base on any of the aforementioned research. In essence, the proposed framework was developed from scratch, in response to the need of utilizing simple, directional return estimates in portfolio optimization.

2.3 Other related strands of research

A few additional areas of research are worth mentioning. First, the developed portfolio selection framework is based on an idea that asset returns can be decomposed into a sign component and a magnitude component. These types of decompositions (although not exactly in the same format) have been presented in, for example, Rydberg and Shepard (2003), Christoffersen and Diebold (2006), and Anatolyev and Gospodinov (2010). Their use of the decomposition is in the context of forecasting returns, whereas in this dissertation, a similar decomposition is utilized to derive analytical results for portfolio selection.

Second, the developed framework results in correlation between absolute values of asset returns affecting portfolio variance explicitly. The second essay of the dissertation delves deeper into this matter, which seems to be a neglected area in financial research. However, it should be noted that correlation of magnitudes is not an unknown concept in general, as studies in, for example, medical physics have examined correlation between magnitudes of changes (see e.g. Ashkenazy, Ivanov, Havlin, Peng, Goldberger and Stanley (2001) and Ashkenazy, Havlin, Ivanov, Peng, Schulte-Frohlinde and Stanley (2003)).

Third, the dissertation makes a minor methodological contribution to conducting trading simulations under directional predictability. The first essay presents a method for generating multivariate random variables that are governed by directional accuracies in the form of probabilities. This makes it possible to run trading simulations with directional forecasts without specifying any particular model for the generation of the forecasts. This approach is based on Leisch, Weingessel and Hornik (1998), and by augmenting their method, pairwise dependent directional forecasts can be generated utilizing a multivariate normal distribution.

3 MAIN FINDINGS OF THE ESSAYS

3.1 New framework for portfolio selection

Estimation of expected returns is notoriously difficult. Instead, investors typically have views only on the *directions* of price changes. In the first essay of the dissertation, a new framework for portfolio selection is developed to utilize directional forecasts as inputs in optimization. The framework is based on an idea that the return on asset i can be decomposed as follows:

$$r_i = s_i D_i |r_i|, \tag{1}$$

where $s_i \in \{-1, 1\}$ is the directional forecast, $D_i \in \{-1, 1\}$ is a random variable denoting whether the directional forecast is correct (+1) or not (-1), and $|r_i|$ is the absolute value of the unconditional return on asset i. If the unconditional asset returns are assumed (or approximated) to be symmetrically distributed about mean zero[5], in which case it is plausible to assume that the outcomes of the forecasts do not depend on return magnitudes, simple expressions for conditional expected return, variance, and covariance between an asset pair can be analytically derived, as is shown in the essay.

The equations for conditional mean, variance and covariance can be generalized for any number of assets. The end result, readily applicable for mean-variance optimization, is presented here: The conditional[6] expected return vector

$$\mathrm{E}[\boldsymbol{r}|\Omega] = \boldsymbol{s} \odot (2\boldsymbol{\zeta} - \mathbf{1})\mathbf{M}, \tag{2}$$

where $\boldsymbol{s}$ is a column vector containing the directional forecasts, $\boldsymbol{\zeta}$ denotes a vector of directional accuracies of the forecasts, and $\mathbf{M} \equiv \mathrm{diag}(\boldsymbol{\mu}_{abs})$, i.e. a diagonal matrix containing the means of the unconditional absolute returns. The operator $\odot$ denotes the element-wise (Hadamard) product. The conditional variance-covariance matrix

$$\boldsymbol{\Sigma}|\Omega = \boldsymbol{s}\boldsymbol{s}' \odot ((2\mathbf{Z}_{sim} - \mathbf{J}) \odot \boldsymbol{\Sigma}_{abs} + \mathbf{M}\boldsymbol{\Sigma}_{\mathbf{D}}\mathbf{M}), \tag{3}$$

[5] This approximation can be justified by the general observation that empirical asset return distributions are fairly symmetrical, and, especially for short horizons (daily or weekly returns), their mean is close to zero. For example, Brown and Warner (1985) report that the daily mean return for randomly selected securities and dates is 0.06%. In addition, the investor discussed in this dissertation does not have to invest in stocks; the investable assets can be, for example, currency pairs or commodities without apparent risk premiums.

[6] Conditional on the information set Ω, which contains the directional forecasts and their accuracies.

where $\mathbf{Z}_{sim}$ is a matrix containing joint pairwise forecast accuracies, $\mathbf{J}$ is a matrix of ones, $\boldsymbol{\Sigma}_{abs}$ is the variance-covariance matrix of the unconditional absolute returns, and $\boldsymbol{\Sigma}_{\mathbf{D}}$ is the covariance matrix of the forecast outcomes (which can also be expressed using the accuracy parameters). $\boldsymbol{\Sigma}|\Omega$ is positive-semidefinite (PSD) by definition, as long as the individual components forming the matrix are sensible.

This framework enables the investor to bypass direct estimation of mean returns, and instead use the directional forecasts as inputs. A new parameter class that requires estimation is comprised of the directional accuracies in $\boldsymbol{\zeta}$ and $\mathbf{Z}_{sim}$. These can be relatively easily estimated from data by backtesting the forecasting model and measuring how often it predicts the sign correctly. The means of absolute returns, $\boldsymbol{\mu}_{abs}$, are, in general, proportional to the volatility of the unconditional asset returns, and in fact, under normality, $\mu_{abs,i} = \sigma_i \sqrt{2/\pi}$, where σ_i is the unconditional return standard deviation, and π is the mathematical constant. If the return volatility follows an ARCH/GARCH process, then the mean of absolute returns can be forecasted to a degree, which would be beneficial for the investor in this setting. In general, forecasting the expected absolute value or the volatility of returns can be considered easier tasks than forecasting the mean.

The implications of this framework are multifaceted. First, assets with high (idiosyncractic) volatility are preferred for higher expected returns[7]. Second, high levels of positive or negative correlation are not preferred for low portfolio variance, assuming that high level of correlation is linked to the investor being correct or wrong on a pair of assets' signs simultaneously. Third, correlation between the absolute values of returns has an explicit effect on portfolio variance. These implications are examined in more detail in the essay.

The foundation for optimal portfolio construction with directional return estimates as inputs is laid out with the analytical derivation in the first essay. To examine more complex situations, the essay conducts a simulation study to explore optimal portfolio compositions under different optimization criteria. Moreover, a performance analysis and a trading simulation show that when directional forecasts are available, the investor could benefit notably in terms of portfolio Sharpe ratios and geometric means from using the proposed framework as opposed to resorting to simpler alternatives.

An additional contribution of the first essay is to show, by augmenting the method of Leisch et al. (1998), how directional forecasts can be simulated with a multivariate normal distribution. This makes it possible to run trading simulations under directional predictability without specifying any particular forecasting model, thus making these trading simulations more general and applicable to a wide variety of scenarios.

[7] These assets are not necessarily held long in the portfolio; the directional view determines whether the asset is held long or short. Volatility is valuable for expected return because it presents an opportunity for larger profits.

3.2 Correlation in the magnitude of returns

The second essay focuses on an overlooked concept in financial research: correlation between *absolute* values of asset returns. In the portfolio optimization framework proposed in the first essay, this correlation emerged as an explicit factor affecting the conditional variance of the portfolio. Even though correlation in the absolute values, or magnitudes (hereafter: magnitude correlation), is not an unknown concept in general, it appears to be largely neglected in financial research.

In addition to affecting the conditional portfolio variance when directional return forecasts are present, magnitude correlation affects the shape of the portfolio return distribution in the case of no predictability as well. Perhaps surprisingly, not much is known about the theoretical or empirical properties of magnitude correlation in the financial context, and the second essay aims to fill this gap in knowledge. It is worth noting that in the univariate case, autocorrelation in absolute returns can arise from a process with autoregressive conditional heteroskedasticity (ARCH). In the same fashion, magnitude correlation could arise from multivariate ARCH or GARCH processes (for more information about these processes, see e.g. Tsay (2010)).

The essay explains in detail how magnitude correlation can affect portfolio return distribution. Moreover, in the bivariate normal case with zero means, it is shown that the link between Pearson correlation and magnitude correlation is fixed:

$$\mathrm{Corr}\Big[|r_i|,|r_j|\Big] = \frac{2\Big(\sqrt{1-\rho^2}+\rho\mathrm{ArcSin}(\rho)-1\Big)}{\pi-2}, \tag{4}$$

where ρ denotes the Pearson correlation coefficient, and π is the mathematical constant. In the bivariate normal case, magnitude correlation can never be negative. However, in practice, returns are often not jointly normal and this link is broken. This is why examining the levels empirically is of particular interest and can present opportunities from a portfolio manager's perspective.

The essay conducts a comprehensive examination of magnitude correlation levels in the U.S. stock market (the data are the S&P500 stock returns for daily, weekly, and monthly intervals) during the 21st century. The results show that the observed pairwise levels of magnitude correlation vary widely between different sample periods and intervals. Especially for longer horizon (monthly) returns, magnitude correlation between a pair of stocks can range from being significantly negative (close to -0.5) all the way up to near perfect positive correlation (0.9). Such variety in the observed levels is intriguing from a portfolio management standpoint. Importantly, the observed levels do not appear to be strongly linked to the corresponding Pearson correlation coefficient and the observed values can deviate widely from the implied values derived from Equation

4. This implies that there can be benefits for choosing assets in the portfolio by looking at magnitude correlation and Pearson correlation separately.

The levels of observed magnitude correlation differ widely between the earliest subsample (2003–2006, i.e. the time before the recent financial crisis) and in the later subsamples (the time during and after the crisis, 2007–2014). Overall, a general tendency of increased levels of magnitude correlation is observed in the data. Moreover, the average pairwise, or marketwide, magnitude correlation appears to be time-varying. Importantly, magnitude correlation for individual assets does not necessarily develop in sync with Pearson correlation, as is illustrated in the essay.

High positive level of magnitude correlation between two assets implies that when the return on the first asset is large in magnitude, the second asset's price is also likely to react strongly (in either direction). Moreover, when the return on the first asset is small in magnitude, the reaction of the second asset is also likely to be small. This kind of a situation for the whole market represents a scenario where markets are anxious, waiting for big news quietly, and when the news arrives, the reaction is strong for all or most assets. Since there is already a commonly used measure for market anxiety, namely, the VIX index measuring the implied volatility of stock options, the second essay compares the observed levels of abnormal[8] magnitude correlation to the VIX index.

The major spikes in the VIX index and in historical abnormal magnitude correlation appear to coincide, and in general, there appears to be a tendency for abnormal magnitude correlation to be high when the VIX is high, and vice versa. However, the strength of this link depends on the estimation method for magnitude correlation. While historical volatility portrays the VIX index in a lagged fashion, the abnormal magnitude correlation appears to react at roughly the same time as the VIX index. Overall, magnitude correlation can be interpreted as market anxiety and may serve as another proxy for it. The possibility of magnitude correlation arising from a multivariate ARCH/GARCH process seems to be supported by this view. It should be noted that the measured magnitude correlation is based on historical data, whereas the VIX index is a forward-looking measure.

Finally, the second essay evaluates out-of-sample forecasting methods for magnitude correlation. This aspect is important from a portfolio manager's standpoint, as more accurate estimates can lead to improved portfolio performance. The results show that long estimation windows or EWMA models with a low value for the gamma parameter appear to produce the most accurate forecasts of future levels of magnitude correlation.

---

[8] Abnormal in this case means the gap between the observed level and the theoretical level implied by the measured Pearson correlation coefficient, as in Equation 4.

3.3 Return forecasts and portfolio performance in practice

The third essay of the dissertation examines the application of the portfolio selection framework proposed in the first essay in practice. The simulation results in the first essay showed that notable performance increase can be achieved over simpler alternatives, however, the directional forecasts were generated in a relatively simple manner and parameter estimation error was not present in the directional accuracies. The third essay aims to take the analysis a step further by evaluating portfolio performance in a realistic setting with directional forecasts available, and furthermore, includes whole return forecasts (on a continuous scale) as an alternative to mere directional forecasts.

Whole return forecasts contain more information, however, if the information is noisy, its practical value may be negligible. Leitch and Tanner (1991) found that professional forecasts evaluated with traditional measures, such as the root-mean-squared error, can often be judged as poor, while at the same time the forecasts might be profitable, which can be connected to their directional accuracy. This kind of a view gets support in the study of Taylor (1980) as well. In portfolio optimization, even small errors in the mean return inputs can matter profoundly for the results, and thus using noisy return estimates can be detrimental.

In essence, the third essay attempts to answer two research questions: 1) When directional return forcasts are available, does it matter for out-of-sample portfolio performance which optimization framework is utilized? 2) If whole return forecasts are available, would the investor do better by extracting only the signs of these forecasts and using them as inputs in portfolio optimization, as opposed to using noisy return estimates directly as expected return inputs? To provide answers to these questions, the third essay conducts out-of-sample trading simulations and evaluates the performance of optimized portfolios. In a broader sense, the research questions are posed in order to find out whether the framework proposed in the first essay can be valuable for asset management in practice.

In order to produce robust results, the out-of-sample trading simulations do not utilize specific forecasting models but instead, a more general approach is used. The return forecasts in the study are generated by calibrating a large number of models in-sample, however, parameter estimation is done with only past data, and hence the use of these models is comparable to an out-of-sample evaluation. By taking this approach, it is certain that the models have true predictive power, but the investor does not know the values of the parameters needed for optimization. Instead, these will be estimated by observing the models' performance in past data, whereas the performance of optimal portfolios is evaluated on out-of-sample data.

In addition, since estimation methods for directional forecast accuracy appear to not have been examined comprehensively in previous research, the third essay conducts an evaluation of simple methods for out-of-sample estimation. The results indicate that constant or shrinkage estimation models produce good performance. The performance of the proposed framework appears to depend strongly on the method utilized for estimating the directional accuracies, and hence, its importance should not be underestimated.

The results of the out-of-sample trading simulation indicate that when only directional return estimates are available, or when magnitude components for the whole return estimates are pure noise, the investor can achieve higher Sharpe ratios and geometric means by using only the directional estimates as inputs. If the whole return estimates contain accurate information about the magnitudes, the investor can still be equally well or better off in the case of maximum Sharpe ratio criterion by using only the extracted signs of the forecasts as inputs in optimization. For maximum geometric mean portfolios, accurate information about the magnitudes contained in the whole return forecasts can produce better performance compared to using the mere signs as inputs.

It should be noticed that the level of accuracy present in the whole return forecasts in the simulation study may not be available in practice. Therefore, the advantage produced by the proposed framework when the magnitude components in whole return estimates contain purely (or mostly) noise could be achieved in practice as well. However, the estimation of the directional accuracies needs to be done carefully in order to achieve good performance. Overall, the results indicate that the usage of the proposed framework could be advocated for portfolio management in practice.

4 CONCLUSION

The main contribution of this dissertation is to develop a novel framework for portfolio selection in order to utilize directional return estimates as inputs in optimization. This bypasses the difficult, direct estimation of mean returns, which has been a long-standing problem for portfolio optimization in practice. In addition, the dissertation examines the implications of the developed framework and studies the feasibility of the approach from a practical investment management perspective. Previously documented out-of-sample directional return predictability and common investor behavior in practice are used as starting points and motivation for the research conducted. By developing a new framework, the dissertation steps partly into new grounds, producing research objectives that are addressed in the three essays comprising the dissertation.

When returns exhibit directional predictability, the developed framework implies that volatile assets are preferred for higher expected returns, negative return correlation is disliked by a rational investor assuming that it is linked to correlated outcomes of forecasts, and the correlation between the absolute values of returns affects portfolio variance explicitly. The developed portfolio selection framework is shown to be capable of producing notably higher Sharpe ratios and higher geometric means compared to simpler alternatives that the investor might resort to without the availability of the proposed framework.

The dissertation conducts a comprehensive study on the properties of correlation between absolute returns, or magnitude correlation, which has been an overlooked concept in financial research. The study shows that this correlation appears to be time-varying and that a wide range of pairwise levels have been observable in the U.S. stock market, presenting opportunities from a portfolio management perspective. Moreover, high levels of magnitude correlation can be interpreted as market anxiety and the observed abnormal levels appear to be loosely linked to the VIX index measuring the implied volatility of stock options.

Finally, the dissertation shows that the developed portfolio selection framework gives promising results in trading simulations compared to alternative frameworks that utilize directional return estimates in portfolio optimization. The estimation of directional accuracy of the forecasts needs to be paid close attention to in order for the proposed framework to provide good performance out-of-sample.

While the dissertation answers the proposed research questions and explores the implications of the developed portfolio selection framework, some questions could not be addressed here. For example, future research could take higher moments into account in portfolio optimization due to asymmetry of the conditional return distribution under the developed framework. It would be rewarding to derive the full distributional properties of the conditional returns in this kind of a setting. For estimating magnitude correlation, more sophisticated measurement could be utilized as the model used inevitably has an effect on the observed levels. Moreover, a class of copula functions could be developed that produce desired levels of magnitude correlation in an explicit fashion, and the role of multivariate ARCH/GARCH processes in generating magnitude correlation could be examined. Finally, to conduct further testing with the developed portfolio optimization framework, specific forecasting models could be utilized instead of the more general approach taken in the third essay. This would sacrifice generality but at the same time increase the authenticity of the study. While these ideas are intriguing, they could not be undertaken in the scope of this dissertation, and therefore, remain as possible goals for future research.

REFERENCES

Ait-Sahalia, Y. – Brandt, M. W. (2001) Variable selection for portfolio choice. *The Journal of Finance*, Vol. 56, 1297–1351.

Anatolyev, S. – Gospodinov, N. (2010) Modeling financial return dynamics via decomposition. *Journal of Business & Economic Statistics*, Vol. 28, 232–245.

Ang, A. – Bekaert, G. (2002) International asset allocation with regime shifts. *Review of Financial Studies*, Vol. 15, 1137–1187.

Ashkenazy, Y. – Havlin, S. – Ivanov, P. C. – Peng, C. K. – Schulte-Frohlinde, V. – Stanley, H. E. (2003) Magnitude and sign scaling in power-law correlated time series. *Physica A*, Vol. 323, 19–41.

Ashkenazy, Y. – Ivanov, P. C. – Havlin, S. – Peng, C. K. – Goldberger, A. L. – Stanley, H. E. (2001) Magnitude and sign correlations in heartbeat fluctuations. *Physical Review Letters*, Vol. 86, 1900–1903.

Barber, B. – Lehavy, R. – McNichols, M. – Trueman, B. (2001) Can investors profit from the prophets? Security analyst recommendations and stock returns. *The Journal of Finance*, Vol. 56, 531–563.

Barberis, N. (2000) Investing for the long run when returns are predictable. *The Journal of Finance*, Vol. 55, 225–264.

Bekiros, S. (2010a) Heterogeneous trading strategies with adaptive fuzzy actor–critic reinforcement learning: A behavioral approach. *Journal of Economic Dynamics and Control*, Vol. 34, 1153–1170.

Bekiros, S. D. (2010b) Fuzzy adaptive decision-making for boundedly rational traders in speculative stock markets. *European Journal of Operational Research*, Vol. 202, 285–293.

Bekiros, S. D. – Georgoutsos, D. A. (2007) Evaluating direction-of-change forecasting: Neurofuzzy models vs. neural networks. *Mathematical and Computer Modelling*, Vol. 46, 38–46.

Bekiros, S. D. – Georgoutsos, D. A. (2008a) Direction-of-change forecasting using a volatility-based recurrent neural network. *Journal of Forecasting*, Vol. 27, 407–417.

Bekiros, S. D. – Georgoutsos, D. A. (2008b) Non-linear dynamics in financial asset returns: the predictive power of the cboe volatility index. *The European Journal of Finance*, Vol. 14, 397–408.

Black, F. – Litterman, R. (1992) Global portfolio optimization. *Financial Analysts Journal*, Vol. 48, 28–43.

Brown, S. J. – Warner, J. B. (1985) Using daily stock returns: The case of event studies. *Journal of Financial Economics*, Vol. 14, 3–31.

Campbell, J. Y. – Chan, Y. L. – Viceira, L. M. (2003) A multivariate model of strategic asset allocation. *Journal of Financial Economics*, Vol. 67, 41–80.

Chevapatrakul, T. (2013) Return sign forecasts based on conditional risk: Evidence from the UK stock market index. *Journal of Banking and Finance*, Vol. 37, 2342–2353.

Christoffersen, P. F. – Diebold, F. X. (2006) Financial asset returns, direction-of-change forecasting, and volatility dynamics. *Management Science*, Vol. 52, 1273–1287.

Chung, J. – Hong, Y. (2007) Model-free evaluation of directional predictability in foreign exchange markets. *Journal of Applied Econometrics*, Vol. 22, 855–889.

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

Elliott, G. – Ito, T. (1999) Heterogeneous expectations and tests of effciency in the yen/dollar forward exchange rate market. *Journal of Monetary Economics*, Vol. 43, 435–456.

Elton, E. J. – Gruber, M. J. (1987) Portfolio analysis with partial information: The case of grouped data. *Management Science*, Vol. 33, 1238–1246.

Grant, D. (1978) Market timing and portfolio management. *Journal of Finance*, Vol. 33, 1119–1131.

Guidolin, M. – Timmermann, A. (2007) Asset allocation under multivariate regime switching. *Journal of Economic Dynamics and Control*, Vol. 31, 3503–3544.

Hallerbach, W. G. (2014) On the expected performance of market timing strategies. *The Journal of Portfolio Management*, Vol. 40, 42–51.

Jensen, M. C. (1972) Optimal utilization of market forecasts and the evaluation of investment performance. *Mathematical Methods in Finance*.

Kandel, S. – Stambaugh, R. F. (1996) On the predictability of stock returns: An asset-allocation perspective. *Journal of Finance*, Vol. 51, 385–424.

Kolm, P. N. – Tütüncü, R. – Fabozzi, F. J. (2014) 60 years of portfolio optimization: Practical challenges and current trends. *European Journal of Operational Research*, Vol. 234, 356–371.

Kuan, C. M. – Liu, T. (1995) Forecasting exchange rates using feedforward and recurrent neural networks. *Journal of Applied Econometrics*, Vol. 10, 347–364.

Leisch, F. – Weingessel, A. – Hornik, K. (1998) On the generation of correlated artificial binary data. *Working Paper Series, SFB "Adaptive Information Systems and Modelling in Economics and Management Sci-*

ence", Vienna University of Economics.

Leitch, G. – Tanner, J. E. (1991) Economic forecast evaluation: Profits versus the conventional error measures. *The American Economic Review*, Vol. 81, 580–590.

Leung, M. T. – Daouk, H. – Chen, A. S. (2000) Forecasting stock indices: a comparison of classification and level estimation models. *International Journal of Forecasting*, Vol. 16, 173–190.

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

Markowitz, H. M. (1959) *Portfolio selection: efficient diversification of investments*. New York: John Wiley & Sons, Inc.

Michaud, R. O. (1989) The markowitz optimization enigma: Is 'optimized' optimal? *Financial Analysts Journal*, Vol. 45, 31–42.

Nyberg, H. (2011) Forecasting the direction of the us stock market with dynamic binary probit models. *International Journal of Forecasting*, Vol. 27, 561–578.

Pesaran, M. H. – Timmermann, A. (1995) Predictability of stock returns: Robustness and economic significance. *The Journal of Finance*, Vol. 50, 1201–1228.

Pesaran, M. H. – Timmermann, A. (2002) Market timing and return prediction under model instability. *Journal of Empirical Finance*, Vol. 9, 495–510.

Pfeifer, P. E. (1985) Market timing and risk reduction. *Journal of Financial and Quantitative Analysis*, Vol. 20, 451–459.

Preminger, A. – Franck, R. (2007) Forecasting exchange rates: A robust regression approach. *International Journal of Forecasting*, Vol. 23, 71–84.

Qian, E. – Gorman, S. (2001) Conditional distribution in portfolio theory. *Financial Analysts Journal*, Vol. 57, 44–51.

Rapach, D. E. – Zhou, G. (2013) Forecasting stock returns. *Handbook of Economic Forecasting*, Vol. 2.

Roy, A. D. (1952) Safety first and the holding of assets. *Econometrica*, Vol. 20, 431–449.

Rydberg, T. H. – Shepard, N. (2003) Dynamics of trade-by-trade price movements: Decomposition and models. *Journal of Financial Econometrics*, Vol. 1, 2–25.

Skabar, A. (2013) Direction-of-change financial time series forecasting using a similarity-based classification model. *Journal of Forecasting*, Vol. 32, 409–422.

Taylor, S. J. (1980) Conjectured models for trends in financial prices, tests and forecasts. *Journal of the Royal Statistical Society. Series A (General)*, Vol. 143, 338–362.

Tobin, J. (1958) Liquidity preference as behavior towards risk. *The Review of Economic Studies*, Vol. 25, 65–86.

Tsay, R. S. (2010) *Analysis of Financial Time Series, 3rd ed.* New York: John Wiley & Sons, Inc.

Tu, J. (2010) Is regime switching in stock returns important in portfolio decisions? *Management Science*, Vol. 56, 1198–1215.

Womack, K. L. (1996) Do brokerage analysts' recommendations have investment value? *The Journal of Finance*, Vol. 51, 137–167.

Part II

ESSAYS

ESSAY 1

Hämäläinen, Joonas

Portfolio selection with directional return estimates

Preprint

Portfolio Selection with Directional Return Estimates

Joonas Hämäläinen*

Abstract

Expected returns of assets are notoriously difficult to estimate. Instead, investors typically have views on the mere *directions* of returns. This paper develops a novel framework for portfolio optimization to use directional forecasts as inputs. We show that in this setting, the investor prefers volatile assets for higher expected returns, dislikes negative return correlation, and that the covariance between *absolute* values of returns matters explicitly for portfolio variance. The developed framework can produce substantially higher Sharpe ratios and geometric means compared to simpler alternatives.

1 INTRODUCTION

Traditional investment portfolio optimization has its roots in the seminal work of Markowitz (1952, 1959), who presented the concept of maximizing the expected return and minimizing the variance of a portfolio. Tobin (1958) showed that when a risk-free asset is available, the risky asset portfolio composition is the same for all investors. These concepts were major breakthroughs in the field of finance. However, Markowitz's approach has not triumphed in practice, mainly due to errors in parameter estimation (see, e.g. Michaud (1989), DeMiguel, Garlappi and Uppal (2009)). Especially the estimation of expected returns has proven to be difficult. Still, returns are known to exhibit some degree

* University of Turku, Department of Accounting and Finance, Rehtorinpellonkatu 3, 20500 Turku, Finland. Part of the research has been conducted during the author's visit to Princeton University in 2013/2014. The author is grateful for the comments of Anders Löflund, Yannick Malevergne, Mika Vaihekoski, Luis Alvarez, John Mulvey, Matti Keloharju, Matthijs Lof, Matti Heikkonen, Hannu Kahra, Johan Knif, and the seminar participants at Princeton University, Aalto University, Helsinki Center of Economic Research, University of Turku, GSF Summer Workshop (Helsinki), and the Eastern Finance Association's 2014 Annual Meeting.

of predictability (for a recent review on stock return predictability, see Rapach and Zhou (2013)).

Research on *directional* predictability of returns has been relatively active over the past decade. There is a considerable amount of evidence that the directions of stock returns can be predicted out-of-sample, often with surprisingly high accuracy – notable recent studies include Bekiros (2010), Nyberg (2011), Chevapatrakul (2013), and Skabar (2013). Moreover, Christoffersen and Diebold (2006) have shown that theoretically, even if conditional mean returns are independent (i.e. not forecastable), there can be directional dependence. Thus, it is possible that return signs can be predicted even if mean returns are not forecastable.

In practice, analysts or investment managers often do not report continuous estimates for expected returns. Instead, it appears that most investors have discrete views of an asset's price either going up or down in the near future[1]. There is evidence that when forecasts on a continuous scale *are* made by professionals, their performance evaluated by traditional measures (such as the root-mean-square error) can be poor whereas their profitability, connected to directional accuracy, can be notable (Leitch and Tanner (1991)). Forecasting the mean return accurately is a near impossible task, which many simply avoid. Therefore, the most prominent forecasts of asset returns appear to be the directions, or signs, of returns.

This paper is built on the two premises discussed above: i) evidence that return directions are predictable; ii) common type of views expressed by investors estimating the future performance of financial assets[2]. The natural question to ask is: how should optimal portfolios be formed in this kind of setting? The contribution of this paper is to develop and examine a new framework for portfolio selection – an answer to the question. The inputs for portfolio selection in this framework are based on the investor's directional views on future returns and the directional accuracies with which the views are assumed to be correct (which can be estimated, for example, by backtesting a forecasting model). This approach bypasses the difficult, direct estimation of mean returns and facilitates the use of directional return estimates in portfolio optimization.

The implications are intriguing and unconventional. For example, as opposed to the traditional mean-variance framework, large negative correlation between asset returns is not preferred because it presents the possibility of the investor being either right or wrong on both assets' directions simultaneously, thus actually increasing the variance of the entire portfolio. Moreover, it is clearly shown that investors seeking to maximize expected return should pick assets exhibiting

1 Buy/sell recommendations being an example of the typical directional views produced by analysts. These types of forecasts have been shown to contain predictive power (see, e.g. Womack (1996), Barber, Lehavy, McNichols and Trueman (2001)).

2 The investable assets dealt with in this paper need not be stocks, but can be members of any asset class such as commodities or currency pairs.

high (idiosyncratic) volatility. Third, an explicit role is played by the covariance of *absolute* values of returns, which appears to be a largely neglected area in financial literature. Finally, a performance analysis and a trading simulation show that the investor could benefit notably from using the proposed framework as opposed to simpler alternatives to form maximum Sharpe ratio or maximum geometric mean portfolios.

To the best knowledge of the author, this framework is original. Some earlier research in related areas include the studies about market timing and portfolio selection by Jensen (1972), Grant (1978), Pfeifer (1985), and more recently, Hallerbach (2014). Elton and Gruber (1987) examine portfolio decision rules when assets are grouped into categories based on expected returns. Moreover, when stock returns are forecastable using a regression model with predictive variables, Kandel and Stambaugh (1996) offer solutions for asset allocation. Barberis (2000) and Campbell, Chan and Viceira (2003) examine return predictability implications for the long-term investor. Ait-Sahalia and Brandt (2001) incorporate the effect of predictive variables on portfolio weights in a direct fashion. Finally, Ang and Bekaert (2002), Guidolin and Timmermann (2007) and Tu (2010) examine portfolio choice in a regime switching environment. Perhaps closest to the framework introduced in this paper comes the work of Black and Litterman (1992) and Qian and Gorman (2001), where views (on a continuous scale) about asset returns are incorporated into portfolio optimization. However, the approach presented here is fundamentally different, focusing on *directional* estimates for several assets, and does not base on the aforementioned studies.

An additional contribution of this paper is to show (in Appendix B), by augmenting the method of Leisch, Weingessel and Hornik (1998), how pairwise dependent directional forecasts can be generated utilizing a multivariate normal distribution. This allows trading simulations with directional predictability to be run without specifying any particular forecasting model.

The paper proceeds as follows. In Section 2, a new framework for portfolio selection, which takes directional estimates of returns as inputs, is developed. The end-product of the analytical examination is a conditional form for the expected return vector and the variance-covariance matrix. In the third section, it is examined what types of assets are preferred in this framework under different optimization criteria. In the fourth section, the theoretical performance of the framework is compared to simpler alternatives available to the investor. In the fifth section, the framework is examined in a trading simulation with authentic stock market data. Finally, the sixth section offers concluding remarks.

2 NEW THEORETICAL FRAMEWORK

2.1 Mean and variance under directional forecasts

Let r_i denote the one-period unconditional net return on asset i. The investor is assumed to have a view $s_i \in \{-1, 1\}$ on the direction, or sign, of the return. Moreover, let $\zeta_i \in [0.5, 1]$ denote the probability with which this view is correct. For simplicity, the probability ζ_i can be assumed to be independent of the actual return. Logically, the minimum value for ζ_i is 0.5, implying that the investor possesses no predictive power and the direction of the return is a coin flip. For practical applications, an estimate of this probability can be obtained, for example, by backtesting a predictive model and measuring how often it produces the correct sign. If the measured ζ_i were less than 0.5, a rational investor would simply reverse her strategy, thus achieving a higher probability of being correct.

In this setting, the return on asset i can be decomposed as follows[3]:

$$r_i = s_i D_i |r_i|, \tag{1}$$

where D_i is a random variable denoting whether the directional forecast is correct or not:

$$D_i = \begin{cases} 1 & \text{with probability } \zeta_i \\ -1 & \text{with probability } (1 - \zeta_i). \end{cases} \tag{2}$$

This holds for any distributional assumption for the unconditional returns. Utilizing this decomposition, implications for portfolio optimization with directional forecasts are derived in what follows.

In order to make the approach analytically more tractable, the unconditional returns r_i are assumed to be symmetrically distributed around mean zero. This fairly well approximates financial returns, especially when the investment horizon is sufficiently short[4]. Notice that the zero-mean assumption does not rule out directional predictability as in Christoffersen and Diebold (2006), because the power of the forecasts can come from any source other than the unconditional distribution itself. It is also important to notice that the *conditional* distribution does not have a mean zero if $\zeta_i > 0.5$.

It is now possible to derive a simple expression for the conditional expected return and variance, as the value of the absolute return is independent of the predicted sign s_i and the outcome D_i is assumed to be independent of the magnitude

---

[3] Similar decompositions of returns into a directional component and a size component have been utilized (in different context) in Rydberg and Shepard (2003), Christoffersen and Diebold (2006), and Anatolyev and Gospodinov (2010).

[4] In general, empirical stock return distributions are fairly symmetrical. Moreover, the investable assets need to be stocks offering a risk premium, but can be, for example, currency pairs or commodities. The zero-mean assumption is a sufficient condition that simplifies things without sacrificing too much in return.

of the return (i.e. the investor cannot forecast the directions of large returns better than the directions of small returns, or vice versa). In what follows, the conditioning information set, Ω, is understood to contain the forecasted signs s_i and the probabilities governing the accuracy of these forecasts. Utilizing the decomposition in Equation 1, the investor's conditional expected return

$$\mathrm{E}[r_i|\Omega] = s_i\zeta_i\mathrm{E}[|r_i|] - s_i(1-\zeta_i)\mathrm{E}[|r_i|] = s_i(2\zeta_i - 1)\mu_{abs,i}, \tag{3}$$

where $\mu_{abs,i} \equiv \mathrm{E}[|r_i|]$. Note that if $\zeta_i = 0.5$, meaning that the sign of the return is a coin flip, the investor is expected to make a zero profit based on the directional forecast because there is no predictive power. The value of $\mu_{abs,i}$ is proportional to the unconditional volatility of the return, and hence its estimation can be compared to estimating the volatility[5]. In the special case of a normal distribution, this actually reduces exactly to estimating just the volatility of the unconditional distribution, as will be shown later.

The conditional variance

$$\begin{aligned}\mathrm{Var}[r_i|\Omega] &= \mathrm{E}\left[(s_iD_i|r_i|)^2|\Omega\right] - (s_i(2\zeta_i-1)\mu_{abs,i})^2 \\ &= \mathrm{E}\left[r_i^2\right] - (2\zeta_i-1)^2\mu_{abs,i}^2 \\ &= \sigma_i^2 - (2\zeta_i-1)^2\mu_{abs,i}^2,\end{aligned} \tag{4}$$

where σ_i^2 is the unconditional variance of r_i.

If the unconditional distribution for r_i is normal[6], it is possible to write $\mu_{abs,i}$ in terms of the unconditional volatility: The probability distribution function of the absolute value of a normal random variable with mean μ and variance σ^2 is known to be

$$f(x) = \frac{1}{\sigma\sqrt{2\pi}}\exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right) + \frac{1}{\sigma\sqrt{2\pi}}\exp\left(-\frac{(x+\mu)^2}{2\sigma^2}\right), \quad (x \geq 0). \tag{5}$$

This is known as the *folded normal distribution* (Leone, Nelson and Nottingham (1961)). Subsequently, the expected value of $|r_i|$ for a zero-mean distribution,

$$\mathrm{E}[|r_i|] = 2\int_0^\infty \frac{|r_i|}{\sigma_i\sqrt{2\pi}} e^{-\frac{|r_i|^2}{2\sigma_i^2}}\, d|r_i| = \sigma_i\sqrt{2/\pi}, \tag{6}$$

where π is the mathematical constant. In this case, the conditional expected return[7]

$$\mathrm{E}[r_i|\Omega] = s_i(2\zeta_i - 1)\sigma_i\sqrt{2/\pi} \tag{7}$$

---

[5] The expected value of the absolute return can be conditional on the information set Ω if the directional forecasts contain information about the joint distribution of unconditional returns. This can be the case if directional forecasts are derived from true returns, which is addressed in Section 2.5. However, conditioning the absolute value on the information set here would be an unnecessary drag and likely to cause more confusion than clarity. Therefore, expectations of the absolute or squared unconditional returns are thought of as being independent of Ω.

[6] The returns for several assets need not be *multivariate* normal.

[7] In a simple market timing context (the investor has only one risky asset available for investment) and under normality, Hallerbach (2014) has recently arrived at similar solutions.

and the conditional variance

$$\mathrm{Var}[r_i|\Omega] = \sigma_i^2 - (2\zeta_i - 1)^2\sigma_i^2(2/\pi). \tag{8}$$

Immediately from these results it is evident that an investor maximizing the conditional expected return will prefer assets that exhibit high (idiosyncratic) volatility and whose return directions can be predicted with a high accuracy. It needs to be emphasized that these assets are not necessarily held long in the portfolio – the forecasted direction determines whether the asset is held long or short. Volatility is valuable in this context because it presents an opportunity for greater profits. Additionally, it can be seen that high unconditional volatility increases the conditional variance, whereas a higher probability of being correct will decrease the variance, as is intuitive to expect.

In these results, conditional expected return is broken down into parts which include either the mean absolute return or the unconditional volatility (depending on whether normality is imposed). The estimation of these variables may be much easier than forecasting mean returns, especially if the return volatility follows an ARCH/GARCH process, which are known to be descriptive of financial return behavior. This is of course beneficial from the investor's point of view, and can be considered an advantage of the framework. An important feature of this framework is to break down the mean return into simpler parts that can be estimated more easily.

It is worth mentioning that the conditional distribution is a piecewise distribution. The conditional return directions can be thought of as being drawn separately from the return magnitudes, and for forecastable asset returns this results in a slightly skewed conditional distribution, weighing the positive values more. In case the probability $\zeta_i = 0.5$, the resulting conditional distribution is exactly symmetrical and if the unconditional returns are normally distributed, then the conditional distribution is exactly normal. While it is possible to take into account higher moments in portfolio optimization[8], this option is left for future research. Restricting to examine only the first two moments in this paper is justified for simplification purposes, and to compare the framework against alternatives using traditional mean-variance optimization criteria.

2.2 Covariance of returns under directional forecasts

To characterize the conditional joint distribution of asset returns from a mean-variance perspective, it is necessary to determine the conditional covariance between returns. As in the univariate case, the outcomes of the forecasts D_i, D_j are assumed to be independent of the unconditional return magnitudes. The conditional covariance involving assets i and j, $\mathrm{Cov}[r_i, r_j|\Omega] =$

[8] Preference for higher moments has been explored in e.g. Scott and Horvath (1980).

$s_i s_j \mathrm{E}[D_i D_j|\Omega]\mathrm{E}[|r_i||r_j|] - \mathrm{E}[r_i|\Omega]\mathrm{E}[r_j|\Omega]$. For the first part of this equation, consider the following definition:

$$D_i D_j = \begin{cases} 1 & \text{with probability } \zeta_{ij} \\ -1 & \text{with probability } (1-\zeta_{ij}). \end{cases} \tag{9}$$

The first alternative realizes if both views are simultaneously correct or wrong, so that $\zeta_{ij} \equiv \mathrm{P}(D_i = D_j) \in [0,1]$, which can be estimated from data, for example, by backtesting a forecasting model. The second alternative takes place if one of the views is correct and the other one is not. Under this definition, $\mathrm{E}[D_i D_j|\Omega]\mathrm{E}[|r_i||r_j|] = \zeta_{ij}\mathrm{E}[|r_i||r_j|] - (1-\zeta_{ij})\mathrm{E}[|r_i||r_j|]$. Moreover, since $\mathrm{E}[|r_i||r_j|] = \mathrm{Cov}[|r_i|,|r_j|] + \mathrm{E}[|r_i|]\mathrm{E}[|r_j|]$, we get the following, simplified form:

$$\mathrm{E}[r_i r_j|\Omega] = s_i s_j (2\zeta_{ij} - 1)\Big(\mathrm{Cov}[|r_i|,|r_j|] + \mu_{abs,i}\mu_{abs,j}\Big), \tag{10}$$

and subsequently,

$$\mathrm{Cov}[r_i, r_j|\Omega] = s_i s_j \Big((2\zeta_{ij} - 1)\mathrm{Cov}[|r_i|,|r_j|] + [(2\zeta_{ij} - 1) - (2\zeta_i - 1)(2\zeta_j - 1)]\mu_{i,abs}\mu_{j,abs}\Big). \tag{11}$$

When $i = j$, ζ_{ij} is equal to one, and the above equation converges to the variance formula presented in Equation 4. Notice that $(2\zeta_{ij} - 1) - (2\zeta_i - 1)(2\zeta_j - 1)$ is actually equal to $\mathrm{Cov}[D_i, D_j]$, so that we can write

$$\mathrm{Cov}[r_i, r_j|\Omega] = s_i s_j \Big[(2\zeta_{ij} - 1)\mathrm{Cov}[|r_i|,|r_j|] + \mathrm{Cov}[D_i, D_j]\mu_{abs,i}\mu_{abs,j}\Big]. \tag{12}$$

If the unconditional marginal distributions are normal with mean zero, then the above equation takes the following form:

$$\mathrm{Cov}[r_i, r_j|\Omega] = s_i s_j \left[(2\zeta_{ij} - 1)\mathrm{Cov}[|r_i|,|r_j|] + \mathrm{Cov}[D_i, D_j]\frac{2\sigma_i\sigma_j}{\pi}\right]. \tag{13}$$

Furthermore, in the normal case,

$$\mathrm{Cov}[|r_i|,|r_j|] = \mathrm{Corr}[|r_i|,|r_j|]\sigma_i\sigma_j\left(1 - \frac{2}{\pi}\right), \tag{14}$$

as $\mathrm{Var}[|r_i|] = \sigma_i^2(1 - 2/\pi)$, a property of the folded normal distribution with zero mean.

From Equation 12, several intriguing features emerge under the assumption that the investor prefers low, or even negative, values for covariance. This is the case when the portfolio weights for all assets are in the same direction as the forecasts, i.e. the investor does not bet against the forecasted directions. In this case, the forecasted directions s_i and s_j in Equations 12 and 13 need not be included, and the two terms inside the brackets remain. The first of these terms is non-negative if $\zeta_{ij} \in [0.5, 1]$. The covariance between absolute returns, $\mathrm{Cov}[|r_i|,|r_j|]$, is of special interest here. If the covariance, or correlation, between absolute returns is negative, the value of the entire first term is lower, which is preferable to an investor following the directional forecasts. It is important to note that a high level of correlation between the *absolute* values of returns does not necessarily imply high correlation between the unconditional returns, or vice versa. Assets with zero (Pearson) return correlation can very well exhibit high correlation between the absolute values of the returns when the form of the joint distribution is not restricted.

The correlation between absolute returns appears to be an overlooked concept in financial literature. However, it plays an explicit part in forming optimal portfolios in this kind of a setting. High positive correlation between absolute returns implies that when asset i provides a small return (positive or negative), asset j is also likely to have a return close to zero, and vice versa, large returns tend to go hand in hand as well. Large negative absolute correlation implies that when asset i exhibits a relatively large return (positive or negative), the return on asset j is likely to be close to zero, and vice versa.

In essence, asset pairs with negative absolute return correlation "smooth" the overall portfolio return if the joint probability ζ_{ij} is high. There is a nonlinear relationship between the conditional return covariance and absolute correlation of the unconditional returns. Figure 1 illustrates this relationship in the normal case when $\sigma_i = \sigma_j = 0.1$ and $\zeta_i = \zeta_j = 0.6$ (Equation 13 multiplied by the product of forecasted signs, $s_i s_j$).

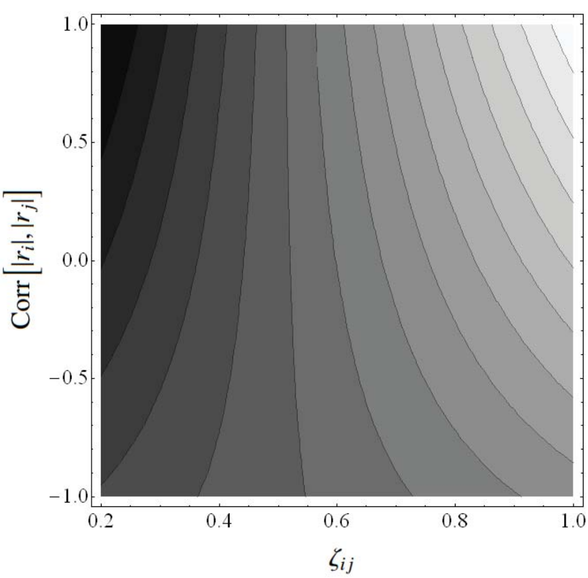

Figure 1: Conditional covariance under normally distributed unconditional returns with the investor following the directional forecasts; darker color indicates lower conditional covariance

From Figure 1, it is evident that the conditional covariance term is the largest when $\zeta_{ij} = 1$ and correlation of absolute returns is equal to 1 (the white area in the upper right hand corner). In this case, both directional forecasts are always correct or wrong at the same time, and the return magnitudes go hand in hand, creating a very variable outcome. To decrease the conditional covariance when $\zeta_{ij} \geq 0.5$, one moves down and left in the diagram — this corresponds to having lower correlation between absolute returns and also a lower value of ζ_{ij}. If $\zeta_{ij} < 0.5$, indicating that the directional forecasts being simultaneously correct or wrong is a more rare event, positive correlation between absolute returns becomes desirable and one moves up and left in the diagram for lower conditional covariance. In practice, it may be difficult to find a pair of assets that exhibit a low ζ_{ij} due to correlated asset returns and, presumably, correlated forecasts. Therefore, as a general guideline, low positive or highly negative absolute return correlation is preferred.

2.3 Properties of ζ_{ij}

The value of ζ_{ij} can be estimated by backtesting a predictive model on a pair of assets, however, its theoretical properties are also of interest. Naturally, there is an imposed structure: the joint probabilities are dependent on the individual probabilities ζ_i and ζ_j. If there is no dependence between the outcomes of two predictions, it is the case that $\zeta_{ij} = \zeta_i\zeta_j + (1-\zeta_i)(1-\zeta_j)$, which is always greater than or equal to 0.5. It seems fair to assume that the value of ζ_{ij} is higher if there is dependence between the signs of unconditional returns r_i and r_j. In this case, forecasts for these assets' return directions are likely to go hand in hand, at least if the forecasting model is sensible and takes into account relevant information.

For example, assume that two assets have a Pearson correlation coefficient of 0.8 for their unconditional returns. Asset 1 is predicted to go up in value in the next period. A sensible forecasting model will also predict Asset 2 to go up in value a majority of the time, unless some other factors in the model tilt the view. Therefore, an investor is likely to take the same directional positions on highly positively correlated assets, and to take opposite positions on highly negatively correlated assets. This, in turn, will make ζ_{ij} go up, since, if forecast for Asset 1 is correct, then there is a greater chance that the forecast for Asset 2 is also correct. Ultimately, if two assets have a correlation coefficient of ± 1, one should be simultaneously correct or wrong on both directional forecasts 100% of the time. In this case, ζ_{ij} would be equal to one. It is logical to assume that any correlation deviating from zero, be it positive or negative, increases the value of ζ_{ij}. Naturally, there are also other factors besides unconditional return correlation affecting predictions and the outcomes.

Based on the above reasoning, it is useful to define the theoretical value of ζ_{ij}

as a function of return correlation as follows:

$$\zeta_{ij} \equiv \zeta_i\zeta_j + (1-\zeta_i)(1-\zeta_j) + f\Big(|\mathrm{Corr}(\mathrm{sgn}(r_i), \mathrm{sgn}(r_j))|\Big) + \eta_{ij}, \tag{15}$$

where $f(\cdot)$ is a monotonic function that gives the absolute value of sign correlation an appropriate weight, and η_{ij} is a term representing unknown factors that can take either positive or negative values. Substituting Equation 15 into Equation 12 and assuming that the investor follows the directional forecasts when forming a portfolio (multiplying r_i by s_i is a way of expressing this) yields

$$\begin{aligned} \mathrm{Cov}[s_i r_i, s_j r_j|\Omega] = &\Big((2\zeta_i - 1)(2\zeta_j - 1) + 2f(x_{ij}) + 2\eta_{ij}\Big)\mathrm{Cov}[|r_i|, |r_j|] \\ &+ (2f(x_{ij}) + 2\eta_{ij})\mu_{abs,i}\mu_{abs,j}, \end{aligned} \tag{16}$$

where $x_{ij} \equiv |\mathrm{Corr}(\mathrm{sgn}(r_i), \mathrm{sgn}(r_j))|$. Taking the derivative with respect to x_{ij}, we get

$$\begin{aligned} \frac{\partial \mathrm{Cov}[s_i r_i, s_j r_j|\Omega]}{\partial x_{ij}} &= 2f'(x_{ij})\mathrm{Cov}[|r_i|, |r_j|] + 2f'(x_{ij})\mu_{abs,i}\mu_{abs,j} \\ &= 2f'(x_{ij})E[|r_i||r_j|] \geq 0 \end{aligned} \tag{17}$$

for any monotonic function $f(\cdot)$. Therefore, higher return correlation (positive or negative) never decreases the conditional covariance and the total variance of the portfolio, assuming that the investor does not go against the directional forecasts. Optimally, the unconditional returns should have zero correlation.

2.4 Generalization to N assets

The equations presented for conditional mean, variance and covariance are easily generalized for any number of assets. The conditional mean return vector

$$\mathrm{E}[\boldsymbol{r}|\Omega] = \boldsymbol{s} \odot (2\boldsymbol{\zeta} - \mathbf{1})\mathbf{M}, \tag{18}$$

where $\boldsymbol{\zeta}$ denotes the vector of directional accuracies, and $\mathbf{M} \equiv \mathrm{diag}(\boldsymbol{\mu}_{abs})$, i.e. a diagonal matrix containing the means of the unconditional absolute returns. The operator $\odot$ denotes the element-wise (Hadamard) product. Analogous to Equation 12, the conditional variance-covariance matrix

$$\boldsymbol{\Sigma}|\Omega = \boldsymbol{s}\boldsymbol{s}' \odot [(2\mathbf{Z}_{sim} - \mathbf{J}) \odot \boldsymbol{\Sigma}_{abs} + \mathbf{M}\boldsymbol{\Sigma}_{\mathbf{D}}\mathbf{M}], \tag{19}$$

where $\mathbf{Z}_{sim}$ is a matrix containing the joint probabilities ζ_{ij}, $\mathbf{J}$ is a matrix of ones, $\boldsymbol{\Sigma}_{abs}$ is the covariance matrix of the absolute values of unconditional returns, and $\boldsymbol{\Sigma}_{\mathbf{D}}$ is the covariance matrix of the forecast outcomes, which can also be expressed with the probabilities in $\boldsymbol{\zeta}$ and $\mathbf{Z}_{sim}$. $\boldsymbol{\Sigma}|\Omega$ is positive-semidefinite (PSD) by definition, as long as all the individual components are well defined.

2.5 Endogeneity of sign forecasts

The developed framework does not require explicit information about the correlation structure of the unconditional asset returns, although the joint probabilities can be viewed as proxying the pairwise correlations. If the sign forecasts, $\boldsymbol{s}$, are viewed as exogenous and the unconditional return signs are formed as $\text{sgn}(\boldsymbol{r}) \equiv \boldsymbol{s} \odot \mathbf{D}$, then the correlation structure of the unconditional asset returns is a result of this process (see Appendix A for a more detailed explanation). If, on the other hand, the sign predictions are viewed as endogenous, i.e. $\boldsymbol{s} \equiv \mathbf{D} \odot \text{sgn}(\boldsymbol{r})$, then the unconditional correlation structure is exogenous and the forecasts $\boldsymbol{s}$ contain additional information about the distribution of the conditional returns $\boldsymbol{r}|\Omega$. The expected values of the random variables in $\mathbf{D}$ and hence also the probabilities ζ_i and ζ_{ij}, would be affected by the assets' unconditional correlation structure. For example, assume that the investor observed sign forecasts that are +1 for all assets. If the unconditional correlation were highly positive across all asset pairs, then the +1 vector of directional forecasts is quite plausible, whereas if the correlation structure were roughly zero across all assets, all the positive forecasts realizing simultaneously is highly unlikely.

The theoretical framework developed in this paper is compatible with both of the aforementioned interpretations. With the first case (sign forecasts are exogenous, correlation structure of the unconditional returns follows), the probabilities in $\boldsymbol{\zeta}$ and $\boldsymbol{Z}_{sim}$ do not depend on the observed sign forecasts. In the second case, the probabilities can be interpreted as conditional on the observed forecasts (i.e. $\boldsymbol{\zeta}|\boldsymbol{s}$) if we assume the unconditional return correlation structure to be known. It would then be possible to take into account the effect of the observed forecasts and determine the conditional probabilities (see Appendix A for an illustration of this). However, it is equally possible to assume that the unconditional correlation structure is unknown (as is the case in practice), and thus treat the probabilities as unconditional.

Moreover, in the framework presented above, the expected value of the absolute or squared returns were assumed to be independent of the information contained in Ω. This is of course a simplification, as the unconditional correlation structure could provide information on the unconditional absolute returns. Again, the model is compatible with this view as well, but in reality we do not know the true correlation structure of the unconditional returns and hence attempting to condition the expectations of absolute values or squared returns on Ω is likely to be an unnecessary effort in practice if the unconditional returns are near symmetrical around mean zero.

2.6 Summary of theoretical findings

The main findings arising from this theoretical framework are summarized in the following key propositions, where forecasting accuracy (ζ_i) is held constant and the investor follows the directional forecasts.

Proposition 1: Assets with high (idiosyncractic) return volatility are preferred for higher expected portfolio return.

Proposition 2: Low positive or large negative correlation between *absolute* returns of assets is preferred for lower variance of the portfolio if the probability of being simultaneously correct or wrong is higher than 50%. The opposite is true if the probability is lower than 50%.

Proposition 3: High correlation between asset returns, be it positive or negative, is not preferred assuming that it is linked to correlated forecast outcomes.

3 COMPOSITION OF OPTIMAL PORTFOLIOS

3.1 Setup

Analytical examination of the developed framework gives a preliminary understanding to how assets should be picked in simple situations. Next, it is examined how portfolios are formed in more complex scenarios utilizing directional return estimates. The aim is to find out how optimal portfolio weights, on average, are distributed in the asset universe under different optimization criteria. This is carried out as a simulation study, generating assets with different properties regarding their unconditional volatility, probabilistic accuracy of forecasts, and the pairwise correlations of absolute returns – i.e. the factors found to be of importance in Section 2.

Portfolios in the study are optimized using three common criteria: minimum variance, maximum reward-to-risk (Sharpe ratio), and maximum geometric mean. The optimization procedure is repeated for a large number of times, the observations are categorized (by quintiles) and the average optimal weights for each category are calculated. This gives insight into what types of assets are preferred by the different optimization criteria.

The simulation is carried out by generating 40 assets which are different for each simulation run, assigning them random $\boldsymbol{\zeta}, \boldsymbol{Z}_{sim}, \mathbf{M}$, and $\boldsymbol{\Sigma}_{abs}$. Unconditional returns are assumed to be normally distributed for simplicity and hence $\mathbf{M}$ is derived from the unconditional volatilities. The range of forecast accuracies in $\boldsymbol{\zeta}$ is based on studies documenting directional predictability out-of-sample.

For example, for monthly returns, directional accuracies between 55% and 65% (and even higher) have been reported (e.g. Leung, Daouk and Chen (2000), Pesaran and Timmermann (2002), Nyberg (2011), and Chevapatrakul (2013)). The lower bound is set equal to 0.5 to take into account assets with (close to) no directional forecastability[9].

For $\boldsymbol{Z}_{sim}$, the construct in Equation 15 is adopted, selecting a random value to represent the combined effect of the latter two terms in the equation[10]. This rarely results in a positive-semidefinite (PSD) $\boldsymbol{\Sigma}_D$, and hence an algorithm based on Higham (2002) is utilized to find the nearest PSD matrix, keeping diagonal values intact. What this does is, it leaves the actual probabilities ζ_i unchanged, but alters the joint probabilities so that the structure is plausible. After a PSD matrix $\boldsymbol{\Sigma}_D$ has been obtained, new joint probabilities $\boldsymbol{Z}_{sim}$ are recovered from this matrix. For generating $\boldsymbol{\Sigma}_{abs}$, first a random PSD correlation matrix is generated, and it is multiplied by the standard deviations of the absolute returns, which are derived from the unconditional volatilities in the case of normality. Once these components have been generated, $\mathrm{E}[\boldsymbol{r}|\Omega]$ and $\boldsymbol{\Sigma}|\Omega$ can be computed.

At this point it is assumed that the investor does not go against the forecasted views, and thus the forecasted signs need not be taken into account in the simulation. This is a means of simplifying the procedure, as we do not need to separate assets into those that are predicted to go up and those predicted to go down in value. For assets that are sold short, this approach corresponds to a situation where the investor is required to retain the value invested short and count it toward the portfolio weights summing up to one. In the performance analysis that follows in Section 4, this simplification is removed to take the signs of the forecasts explicitly into account.

3.2 Minimum variance portfolios

Directional forecasts not only affect the conditional expected return vector, but also the conditional variance-covariance matrix. Thus, the composition of the global minimum variance portfolio is influenced by the directional forecasts as well. The minimum variance optimization problem is defined as follows:

$$\begin{aligned} \min_{\boldsymbol{w}} \quad & \boldsymbol{w}'\boldsymbol{\Sigma}\boldsymbol{w} \\ \text{s.t.} \quad & \boldsymbol{w}'\mathbf{1} = 1, \end{aligned} \tag{20}$$

where $\boldsymbol{\Sigma}$ is the variance-covariance matrix, and $\boldsymbol{w}$ denotes a column vector of portfolio weights. As mentioned earlier, the portfolio weights are assumed to

[9] The specifications for the simulation are: $\zeta_i \sim \mathrm{U}(0.5, 0.65); \zeta_{ij} = \min\{\zeta_i\zeta_j + (1-\zeta_i)(1-\zeta_j) + e_{ij}, 1\}, e_{ij} \sim \mathrm{U}(0, 0.5); \sigma_i \sim \mathrm{U}(0.02, 0.2); \mathrm{Corr}[|r_i|, |r_j|] \sim \mathrm{U}(-1, 1)$.

[10] This random value is non-negative and hence $\zeta_{ij} \geq 0.5$, which is likely to be the case in practice as well.

be in the direction of the forecasts. This is equivalent to limiting the weights w_i to be non-negative in the optimization and multiplying the conditional estimate for $\mathbf{\Sigma}$, $\mathbf{\Sigma}|\Omega$ according to Equation 19, element-wise by $\boldsymbol{ss'}$. In essence, this simplifies the procedure by restricting that all the weights are in the direction of the forecasts (or zero) and hence, the simulation study does not differentiate between the assets that are predicted to go up and those predicted to go down in value.

After 100,000 simulation runs, the assets are categorized according to their characteristics to form bins based on quintiles. The average weight for the assets in each bin is then computed to find out about the preference relations of assets under the minimum variance criterion.

In Figure 2, the assets are categorized based on the forecast accuracy ζ_i, the unconditional standard deviation, the weighted average joint probability of being correct/wrong simultaneously, ζ_{ij}, and the weighted average absolute correlation between asset pairs. The weighted averages are formed by weighing the values of ζ_{ij} with the corresponding optimal portfolio weights, so that only the joint probabilities across assets *included* in the portfolio matter. The same approach is used in the case of the absolute correlations. The figure displays the average optimal portfolio weight in each of the 25 bins as a circle with a diameter related to the average portfolio weight. The larger the circle is, the more weight, on average, is given to the assets in the corresponding category. The numbers above the columns and next to the rows indicate the upper bound for the corresponding category.

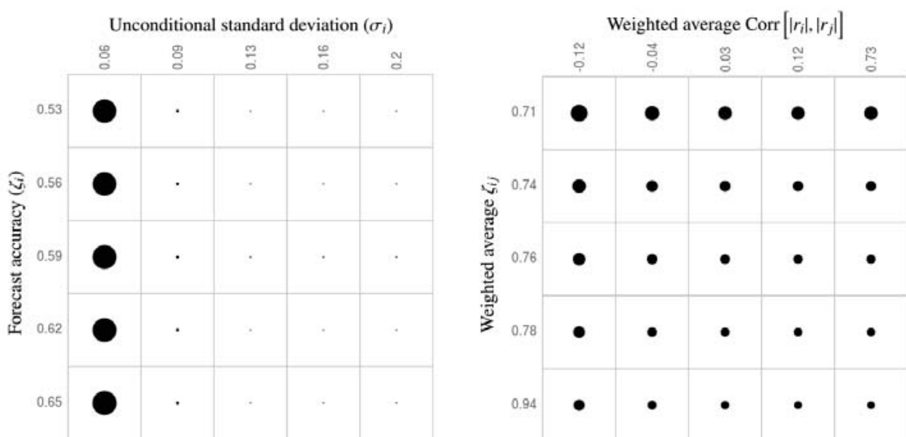

Figure 2: Preference relations under the minimum variance criterion

From the left diagram in Figure 2, it is clear that the minimum variance criterion strongly prefers assets with low unconditional standard deviation, as the circles are the largest on the leftmost column in the diagram, whereas other columns do not get much weight at all. Vertically, there is not much difference between the circle diameters, indicating that forecast accuracy ζ_i does not really matter for minimum portfolio variance compared to the importance of the unconditional volatility.

As is evident from the theoretical construction in Section 2, asset pairs with low ζ_{ij} are valued for lower portfolio variance. The importance of ζ_{ij} can be seen in the right diagram of Figure 2, where the assets are categorized based on the weighted averages of absolute return correlation and ζ_{ij}. The circles get progressively larger as we move up in the diagram, indicating that low ζ_{ij} assets get more weight on average assigned to them. As shown in Figure 1, when the value of $\zeta_{ij} > 0.5$, low (negative) absolute return correlation is preferred. Figure 2 also demonstrates this, as there is a tendency for the circles to get progressively larger as we move left in diagram. From Figure 1, it is evident that ζ_{ij} is the dominant factor affecting conditional covariance over absolute return correlation, and this can be seen in Figure 2 as well, as moving vertically in the right diagram has a larger effect on the average weight than moving horizontally does.

More diagrams like the ones in Figure 2 can be produced for other preference relations, and they are in line with the theoretical findings in the earlier section. Here only two illustrative cases are presented due to limited space. Regarding relations between the other asset characteristics, it can be mentioned that unconditional standard deviation also dominates ζ_{ij}, and it appears to be the most important parameter in selecting assets in the global minimum variance portfolio, as is intuitive to expect.

3.3 Maximum reward-to-risk (Sharpe ratio) portfolios

The maximum expected return-to-risk, or the Sharpe ratio (Sharpe (1966)) omitting a risk-free rate, criterion balances portfolio expected return relative to risk measured by the standard deviation of the portfolio return. In this case, the optimization problem is:

$$\begin{aligned} \max_{\boldsymbol{w}} \quad & \frac{\boldsymbol{w}'\boldsymbol{\mu}}{\sqrt{\boldsymbol{w}'\boldsymbol{\Sigma}\boldsymbol{w}}} \\ \text{s.t.} \quad & \boldsymbol{w}'\mathbf{1} = 1, \end{aligned} \tag{21}$$

where $\boldsymbol{\mu}$ denotes the expected return vector. As in the case of the minimum variance criterion, the investor is not allowed to go against the directional forecasts, which is equivalent to adding a non-negativity constraint for the weights and multiplying the estimate for $\boldsymbol{\mu}$ element-wise by $\boldsymbol{s}$, and multiplying the estimate for $\boldsymbol{\Sigma}$ element-wise by $\boldsymbol{s}\boldsymbol{s}'$.

The optimal weights from 100,000 simulation runs are collected. The assets are sorted into bins according to their characteristics, and the average weights for each bin are calculated. Figure 3 displays the results for assets categorized based on the forecast accuracy ζ_i, unconditional standard deviation, and the weighted average joint probability ζ_{ij}. As opposed to the global minimum variance portfolios, high values for the probability of being correct, ζ_i, are now valued critically. Low standard deviation is still preferred, but forecast accuracy is clearly the dominating factor, which is evident from the left diagram in Figure 3. In the proposed framework, high unconditional standard deviation is favored for higher conditional expected return, but at the same time it increases the conditional variance of the portfolio. For the maximum Sharpe ratio criterion, the gain in conditional expected return appears to be offset by the increase in conditional portfolio return volatility.

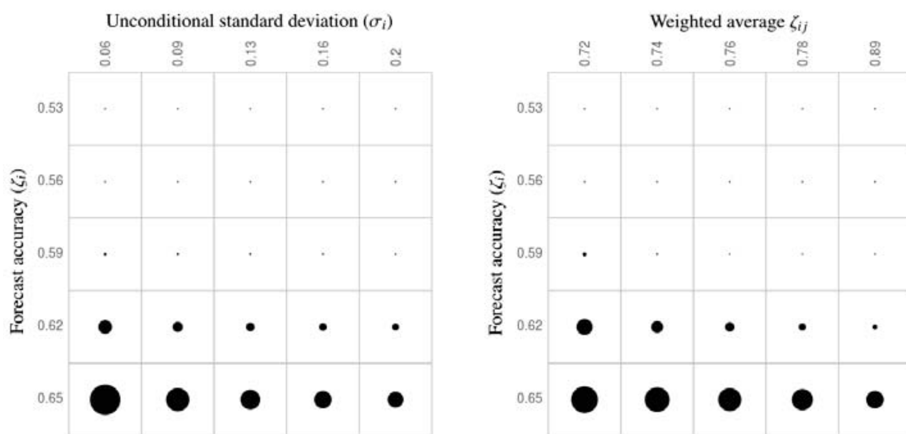

Figure 3: Preference relations under the maximum reward-to-risk criterion

The diagram on the right in Figure 3 shows that forecast accuracy is also more important than the pairwise joint probability ζ_{ij} for maximum Sharpe ratio portfolios, as moving vertically in the diagram has a larger effect on the circle diameters compared to moving horizontally. More diagrams like the ones in Figure 3 can be produced, and they are in line with the findings in the theoretical section of this paper. For example, low or negative absolute correlation is preferred due to ζ_{ij} being greater than 0.5. Overall, the maximum Sharpe ratio criterion appears to pay the most attention to the forecast accuracy ζ_i.

3.4 Maximum geometric mean portfolios

The so-called maximum geometric mean (MGM) criterion has received recent attention by Estrada (2010), who advocates its use. As the name indicates, the criterion seeks to maximize the geometric mean of the portfolio return, which is in line with maximizing the terminal wealth of the portfolio. Following Estrada (2010), this goal can be achieved by solving the following optimization problem:

$$\begin{aligned} \max_{\boldsymbol{w}} \quad & \left[\ln(1+\boldsymbol{w'\mu}) - \frac{\boldsymbol{w'\Sigma w}}{2(1+\boldsymbol{w'\mu})^2} \right] \\ \text{s.t.} \quad & \boldsymbol{w'1} = 1. \end{aligned} \tag{22}$$

As before, the investor is assumed to follow the directional views and hence a non-negativity constraint for the weights is in place and the input estimates are modified accordingly. The assets picked in the MGM portfolios have quite different characteristics compared to those preferred by the previous two optimization criteria. As is evident from Figure 4, a combination of high unconditional standard deviation and forecasting accuracy, ζ_i, is preferred. This is intuitive to expect, since according to the theoretical framework in Section 2, those are the two components affecting the conditional expected return of the portfolio.

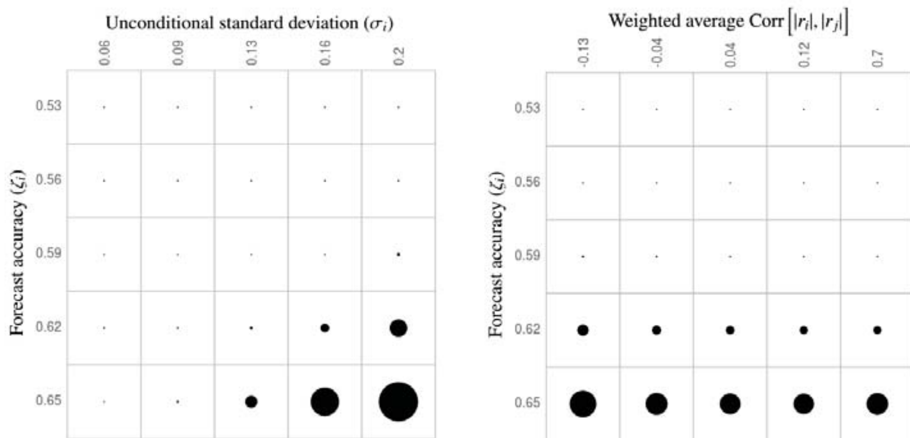

Figure 4: Preference relations under the maximum geometric mean criterion

From the left diagram in Figure 4, it appears that forecasting accuracy ζ_i and the unconditional standard deviation have an equally important role in the MGM portfolios. High standard deviation is not preferred unless the particular asset is accompanied by a high probability of getting the sign of the return correct. In this sense, the MGM criterion is quite selective about the assets it includes in the portfolio.

On the right diagram in Figure 4, assets are sorted based on forecasting accuracy, ζ_i, and the weighted average correlation between the absolute returns. It is clear that ζ_i is the more important factor here, but lower values for the correlation are still clearly preferred to larger ones. This indicates that low portfolio variance is still preferred by the MGM criterion, even though high conditional expected return appears to be the main priority.

4 PORTFOLIO PERFORMANCE AGAINST SIMPLER ALTERNATIVES

4.1 Construct of the analysis

Having found out what types of assets are preferred under different optimization criteria, the next step is to evaluate what kind of a performance increase can be achieved by using the proposed framework compared to simpler alternatives. The investor obtains directional forecasts and is faced with a portfolio selection problem according to the three criteria presented earlier (minimum variance, maximum reward-to-risk, and maximum geometric mean). All of the parameters required in the construction of optimal portfolios are assumed to be known in this section, and thus the analysis conducted is essentially an examination of theoretical performance differences.

The performance of three alternative methods for portfolio construction are compared to that of the proposed framework. These alternatives are simple methods that the investor would likely use if she did not have the proposed framework available. *The first alternative* assumes that the investor constructs the expected return vector by multiplying the directional forecasts by a simple magnitude estimate; the mean of the individual assets' absolute returns, which under normality is simply $\bar{\sigma}\sqrt{2/\pi}$. For the variance-covariance matrix estimate, the (known) unconditional covariance matrix is used in the optimization.

The second alternative constructs the expected return vector by multiplying the directional forecasts with noise generated from independent uniform distributions extending from zero to the mean of all assets' standard deviations, $\bar{\sigma}$. This mimicks a procedure where the investor will try to incorporate noisy estimates of the return magnitudes in order to perform mean-variance optimization. The unconditional covariance matrix is used as an estimate for the covariance matrix in the optimization.

The third alternative seeks to test whether the construct of the covariance matrix estimate matters if the expected return vector is formed as in the proposed framework. The purpose is to test how much worse off the investor will be if she instead opts to use the known *unconditional* covariance matrix together with the proposed framework's *conditional* expected return vector, thereby only utilizing "half" of the framework.

As opposed to the previous section where optimal portfolio composition was examined, here the investor is allowed to go against the directional forecasts. For example, if there is an asset combination that would be favorable for portfolio variance if an opposite position to the directional forecast was taken, the investor may do so by sacrificing expected return for lower variance. A constraint is imposed on the *absolute* values of portfolio weights summing up to one, i.e. no additional funds are obtained by short selling an asset. This prevents the optimization from resulting in combinations where leverage is enormous. Moreover, this type of an approach is applicable to other asset classes as well, not just stocks.

Portfolio optimization is performed utilizing the proposed framework and each of the three alternatives described above. Optimal weights according to the different optimization criteria are solved and they are used to compute the *conditional* expected return and variance of the portfolio according to the proposed framework. That is, even though simpler estimates are used for the alternatives during optimization, the resulting portfolio performance is evaluated assuming that the conditional views of the proposed framework are accurate. In essence, what this does is it evaluates the loss incurring by using a simpler model when the "true" model is available.

4.2 Generating the investment environment

In order for the simulation study to cover as wide and general situations as possible, the characteristics of investable assets are created randomly following a specific set of guidelines that mimicks the characteristics of real world assets. In addition, directional forecasts are generated according to a simple principle explained in what follows.

For each simulation run, all parameters required in the optimization are randomly generated in the same fashion as in the portfolio composition study in chapter 3. In order to test the effect of different scenarios and to make the study more robust, two separate values for the upper limit of ζ_i (55% and 60%, which are in line with a number of recent out-of-sample studies on directional predictability, such as Bekiros and Georgoutsos (2008), Nyberg (2011), and Chevapatrakul (2013)) are included, and two ranges for the unconditional sign correlations are specified ($\text{Corr}[\text{sgn}(r_i), \text{sgn}(r_j)] \in [-1, 1]$ and $\text{Corr}[\text{sgn}(r_i), \text{sgn}(r_j)] \in [0.5, 1]$). Absolute return correlation is set to range between 0 and 1.

The unconditional returns are assumed to have marginal normal distributions with zero means for simplicity. A random unconditional sign correlation matrix $\boldsymbol{R}_{sgn}$ is generated and made PSD by utilizing an algorithm based on Higham (2002). An *absolute* return correlation matrix is generated in a similar fashion, from which the absolute return covariance matrix can be derived under normality. These are then combined to form the unconditional covariance matrix $\boldsymbol{\Sigma}$. This is accomplished by assuming that the zero-mean return signs are independent of the magnitudes[11], in which case

$$\begin{aligned} \mathrm{Cov}[r_i, r_j] &= \mathrm{E}[\mathrm{sgn}(r_i)\mathrm{sgn}(r_j)]\mathrm{E}[|r_i||r_j|] \\ &= \mathrm{Corr}[\mathrm{sgn}(r_i), \mathrm{sgn}(r_j)](\mathrm{Cov}[|r_i|, |r_j|] + \mathrm{E}[|r_i|]\mathrm{E}[|r_j|]). \end{aligned} \tag{23}$$

Under zero-mean normality, the expected absolute values can be expressed in terms of unconditional standard deviations (as demonstrated in the theoretical section). This is easy to generalize to matrix form, and thus the unconditional covariance matrix can be formed.

The joint probabilities ζ_{ij} are generated for each simulation run by assuming the form presented in Equation 15, as was done in Section 3. However, now that there is a generated sign correlation matrix available, the joint probabilities are formed accordingly by assuming that $f(x) = x/2$, which leads to highly correlated assets having their forecast outcomes also correlated. The additional term in Equation 15, $\eta_{ij} \sim \mathrm{U}(-0.1, 0.1)$. These ranges and values should be fairly realistic and cover a wide variety of scenarios.

The directional forecasts are generated by simulating from a multivariate normal distribution with a covariance matrix equal to the generated sign correlation matrix $\boldsymbol{R}_{sgn}$. The directional forecasts, $\boldsymbol{s} \equiv \mathrm{sgn}(\boldsymbol{y})$, where $\boldsymbol{y} \sim \mathrm{N}(\boldsymbol{0}, \boldsymbol{R}_{sgn})$. This procedure has a tendency to generate directional forecasts that are less correlated than indicated by the sign correlation matrix, however, this could also be expected from forecasting models in practice, since their accuracy can be low, and they could be influenced by other factors than just correlations between the asset returns. Moreover, the correlation between the generated sign changes for each simulation run and is not expected to make a noticeable difference in the results.

4.3 Results

The results for the global minimum variance portfolios with 40 investable assets in 1000 different scenarios (simulation runs) are presented in Table 1. The table shows the average portfolio mean, average standard deviation, diversification measure (DI)[12], and weights against the directional forecasts (WA) for the proposed framework and the alternative which uses the unconditional covariance matrix in optimization. Naturally, the results are identical for the other two alternatives in the case of the global minimum variance portfolios and hence they are omitted here.

[11] The assumption is made in order to separate the generation of sign correlation matrix and the absolute correlation matrix. It should not have an effect on the results as the version of the theoretical framework utilized does not take into account possible dependency between the absolute returns and return signs.

[12] The average maximum absolute weight in the portfolio divided by the average number of assets with more than 5% absolute weight. The lower this ratio is, the more diversified the portfolio can be considered to be.

Table 1: Simulated performance of global minimum variance portfolios

	Proposed framework				*Alternative*			
	Avg.	Avg.			Avg.	Avg.		
SgnCorr[-1,1]	Mean	St.Dev.	DI	WA	Mean	St.Dev.	DI	WA
Accuracy[0.5,0.55]	0.000	**0.005**	0.029	49.8 %	0.000	**0.010**	0.030	50.3 %
Accuracy[0.5,0.60]	0.000	**0.005**	0.029	50.1 %	0.000	**0.010**	0.030	49.8 %
SgnCorr[0.5,1]								
Accuracy[0.5,0.55]	0.000	**0.003**	0.029	49.9 %	0.000	**0.007**	0.029	49.9 %
Accuracy[0.5,0.60]	0.000	**0.003**	0.029	49.9 %	0.000	**0.008**	0.029	50.1 %

The alternative model represents a case where the covariance matrix estimate is the unconditional (known) variance-covariance matrix.

As is to be expected, the average mean returns of the portfolios are virtually zero for all cases. It is evident that the proposed framework produces significantly lower portfolio standard deviations than the simpler alternative which uses the unconditional covariance matrix in optimization. The average portfolio standard deviation produced by the proposed framework is less than half of that produced by the alternative (0.003 vs. 0.007 or 0.008) in the case of high positive unconditional sign correlations. When the unconditional correlations are spread wider (between -1 and 1), the performance difference is still large (0.005 vs. 0.010). Also important to notice is that the range of the forecast accuracy ζ_i does not appear to have any effect on the achieved performance (the small difference that the alternative framework produces in the case of high positive sign correlation is due to randomness included in the simulation). This is in line with the results of Section 3, where it was demonstrated that ζ_i does not have a significant influence on the composition of global minimum variance portfolios.

From Table 1, it can be observed that the values of the diversification measure (DI) are similar for both frameworks in all evaluated scenarios, indicating that the portfolios are in general equally well diversified. The average weight against the forecasted views is close to 50% for both models, indicating that for about half of the assets included in the portfolio, the position taken is against the forecasted direction. For the alternative that uses the unconditional covariance matrix in optimization, the forecasts do not affect the estimate for the covariance matrix, and so it is expected that the WA number be close to 50%. For the proposed framework, the result indicates that when low portfolio variance is first priority, it is not rare to go against the forecasted directions.

Perhaps surprisingly, the average standard deviation for the portfolios is lower in the case of high positive unconditional sign correlations, as opposed to the case where the range is wider for sign correlations. This can be explained by the fact that when going against the forecasted direction, the investor's preference for low ζ_{ij} switches to preference for high ζ_{ij}, which are found with highly correlated assets in the simulation. For the unconditional alternative, having highly positively correlated assets presents the opportunity of shorting a stock for lower covariance. Since the absolute weights of the portfolio are constrained to sum up to one, this does not increase the total variance of the portfolio because no additional funds are obtained by short selling a stock.

For the maximum Sharpe ratio portfolios, the results are presented in Table 2. In addition to reporting the average mean and standard deviation of the portfolios, the average mean-to-standard deviation ratio (Sharpe ratio without a risk-free rate) is presented. For 40 investable assets, the proposed framework produces Sharpe ratios that are nearly twice the amount of those produced by the best alternative in all scenarios, and the difference is in many cases even larger than that. For example, in the case of high positive correlations and ζ_i ranging between 0.5 and 0.6, the average Sharpe ratio of the proposed framework is 0.756 compared to that produced by the best alternative (Alternative 1 in Table 2, where the proposed conditional expected return vector is utilized, but the covariance matrix estimate is unconditional), 0.275.

Table 2: Simulated performance of maximum reward-to-risk portfolios

	Proposed framework					*Alternative 1*				
SgnCorr[-1,1]	Avg. Mean	Avg. St.Dev.	Avg. SR	DI	WA	Avg. Mean	Avg. St.Dev.	Avg. SR	DI	WA
Accuracy[0.5,0.55]	0.002	0.008	**0.260**	0.026	46.9 %	0.003	0.022	**0.131**	0.026	32.6 %
Accuracy[0.5,0.60]	0.004	0.008	**0.509**	0.026	46.9 %	0.006	0.022	**0.263**	0.026	33.4 %
SgnCorr[0.5,1]										
Accuracy[0.5,0.55]	0.002	0.004	**0.397**	0.025	49.7 %	0.002	0.021	**0.140**	0.026	40.6 %
Accuracy[0.5,0.60]	0.003	0.004	**0.756**	0.025	50.1 %	0.004	0.021	**0.275**	0.027	40.6 %

	Alternative 2					*Alternative 3*				
SgnCorr[-1,1]	Avg. Mean	Avg. St.Dev.	Avg. SR	DI	WA	Avg. Mean	Avg. St.Dev.	Avg. SR	DI	WA
Accuracy[0.5,0.55]	0.001	0.015	**0.062**	0.044	40.0 %	0.001	0.016	**0.070**	0.039	36.1 %
Accuracy[0.5,0.60]	0.002	0.015	**0.121**	0.044	40.3 %	0.002	0.016	**0.139**	0.039	36.2 %
SgnCorr[0.5,1]										
Accuracy[0.5,0.55]	0.001	0.014	**0.032**	0.043	45.4 %	0.001	0.015	**0.035**	0.041	43.6 %
Accuracy[0.5,0.60]	0.001	0.014	**0.063**	0.045	45.4 %	0.001	0.015	**0.067**	0.041	43.5 %

Alternative 1 represents the case where the expected return estimates are conditional according to the proposed framework, but the covariance matrix estimate is unconditional. Alternatives 2 and 3 represent the cases where simpler estimates are used for the expected return vector along with the unconditional covariance matrix: Alternative 2 combines the directional forecasts with noise, and Alternative 3 multiplies the forecasts by a constant.

From Table 2, it is evident that alternative 1 (utilizing the "true" conditional expected return vector, but lacking the correct covariance matrix) fairs better than the simpler alternatives 2 and 3, but still notably worse than the proposed framework. This indicates that it is important to use the correct conditional covariance matrix estimate with the proposed expected return vector. Moreover, Table 2 shows that the range of ζ_i matters substantially for performance, as the produced Sharpe ratios are notably higher for the case where the range for ζ_i extends to 60%. This is different from the case of the global minimum variance criterion, where the range of forecast accuracy did not matter for performance. Moreover, the portfolios generated by using the proposed framework appear to be about equally well diversified as those generated by using alternative 1, and notably better diversified as those generated by alternatives 2 and 3.

Table 2 shows that the proposed framework leads to going against the forecasted directions very often (nearly 50% of the portfolio weight for all evaluated scenarios). In the scenario where the unconditional sign correlations have a wider range (between -1 and 1), going against the directional forecasts is not as common as when the correlations are high and positive. This can be explained by the fact that high correlations lead to higher ζ_{ij}, which in turn can make it worth taking a position against the forecasted direction. The proposed framework appears to sacrifice quite a lot of expected return in order to reduce portfolio variance.

Table 3 presents the results for maximum geometric mean (MGM) portfolios. In addition to the average mean and standard deviation, the average geometric mean return is presented and used as the main performance indicator to compare the proposed framework against the different alternatives. The proposed framework produces substantially higher average geometric means than the two simpler alternatives (2 and 3). Against Alternative 1, which uses the "true" conditional expected return vector, the performance increase is relatively small, but still notable especially when the range for forecast accuracy ζ_i extends to just 0.55. This is presumably the case because for the MGM portfolios, generating a high positive expected return is a first priority, and the variance of the portfolio appears to matter less when directional accuracies for the forecasts are high. Thus, the estimate for the covariance matrix does not play as big a role and it may suffice to use the unconditional estimate instead.

From Table 3, it can also be observed that the level of forecast accuracy ζ_i has a large impact on portfolio performance in all presented scenarios. Interestingly, weight against the forecasted directions is close to 0% for all cases, indicating that the optimization criterion very rarely sacrifices expected return of the portfolio for a decrease in portfolio variance by going against the directional forecasts. A slight exception to this is the case when sign correlations are highly positive and the forecast accuracy is relatively low. In this case, the proposed framework leads to about 7.5% of the weights being against the forecasts. This is a situation (lower directional accuracies, higher ζ_{ij}) where portfolio variance is given more attention, and lower variance can be achieved by going against the forecasted directions.

Finally, the average geometric means are higher when the unconditional sign correlations range between -1 and 1 compared to the case of high positive correlation. This is an opposite finding to earlier in the case of minimum variance or maximum Sharpe ratio portfolios, and makes sense because with the MGM portfolios, weights against the forecasted directions are not common, meaning

Table 3: Simulated performance of maximum geometric mean portfolios

	Proposed framework					*Alternative 1*				
SgnCorr[-1,1]	Avg. Mean	Avg. St.Dev.	Avg. Geomean	DI	WA	Avg. Mean	Avg. St.Dev.	Avg. Geomean	DI	WA
Accuracy[0.5,0.55]	0.010	0.078	**0.73 %**	0.036	1.1 %	0.011	0.091	**0.65 %**	0.043	0.0 %
Accuracy[0.5,0.60]	0.023	0.098	**1.82 %**	0.072	0.0 %	0.023	0.102	**1.76 %**	0.072	0.0 %
SgnCorr[0.5,1]										
Accuracy[0.5,0.55]	0.008	0.063	**0.58 %**	0.045	7.5 %	0.010	0.115	**0.27 %**	0.069	2.0 %
Accuracy[0.5,0.60]	0.021	0.109	**1.50 %**	0.097	3.7 %	0.022	0.133	**1.34 %**	0.109	1.1 %
	Alternative 2					*Alternative 3*				
SgnCorr[-1,1]	Avg. Mean	Avg. St.Dev.	Avg. Geomean	DI	WA	Avg. Mean	Avg. St.Dev.	Avg. Geomean	DI	WA
Accuracy[0.5,0.55]	0.003	0.064	**0.10 %**	0.300	0.0 %	0.002	0.024	**0.16 %**	0.044	0.0 %
Accuracy[0.5,0.60]	0.007	0.062	**0.43 %**	0.313	0.0 %	0.004	0.024	**0.36 %**	0.042	0.0 %
SgnCorr[0.5,1]										
Accuracy[0.5,0.55]	0.003	0.075	**0.00 %**	0.407	0.0 %	0.002	0.035	**0.11 %**	0.070	0.0 %
Accuracy[0.5,0.60]	0.006	0.073	**0.33 %**	0.421	0.0 %	0.003	0.035	**0.27 %**	0.075	0.0 %

Alternative 1 represents the case where the expected return estimates are conditional according to the proposed framework, but the covariance matrix estimate is unconditional. Alternatives 2 and 3 represent the cases where simpler estimates are used for the expected return vector along with the unconditional covariance matrix: Alternative 2 combines the directional forecasts with noise, and Alternative 3 multiplies the forecasts by a constant.

that high values for ζ_{ij} are not useful for reducing portfolio variance (by going against the forecasts).

Overall, the results of this section indicate that the performance increase from using the proposed framework can be very significant compared to using simpler alternatives that the investor has available for utilizing directional forecasts in mean-variance optimization. Using the correct form for just the expected return estimate and not the covariance matrix produces better results than the other two alternatives, however, the performance increase from using the whole proposed framework is much larger, especially in the case of maximum Sharpe ratio portfolios.

5 TRADING SIMULATION

To relax the assumptions about normal or symmetric returns with mean zero, a trading simulation with authentic stock market data is conducted. The purpose is to see how well the proposed framework does against a simple alternative when the distributional assumptions behind the model do not accurately hold.

5.1 Construct of the study

In the trading simulation, U.S. stock market data provide an authentic investment environment. Directional return forecasts for the available stocks are generated and portfolio wealth development utilizing the new framework is observed, comparing the results against a simpler alternative. The experiment is conducted using weekly closing levels of total return indices for the S&P100 stocks from December 2002 to December 2014[13]. The weekly interval is chosen as it provides more observations than using monthly returns, and the predictability of daily returns can be assumed to be relatively weak – the weekly interval finds a balance between these two.

A vector of net percentage returns, $\boldsymbol{r}_t$, is computed for each week. The same data are utilized multiple times in the simulation, each time changing the parameters that govern the directional forecasts. This way the scenario changes on each simulation run. Alternatively, one could think of this setting as having a large number of investment managers, each with a unique forecasting model for the return signs. Additionally, only 40 stocks are chosen as investable assets for each simulation run, so that not only the parameters governing the forecasts change, but also the assets themselves are different from one run to the other. The end result is a simulated stochastic investment world based on authentic stock market data.

The generation of directional forecasts is explained in more detail in the next section. As before, we need to generate a probability vector $\boldsymbol{\zeta}$[14], and a matrix of joint probabilities ζ_{ij}, $\boldsymbol{Z}_{sim}$. This is done as previously in the simulation study in Section 4, taking into account the sample-based sign correlations between asset returns and incorporating their effect into ζ_{ij} according to Equation 15[15]. The covariance matrix of $\mathbf{D}$ can then be computed as previously, again utilizing an algorithm based on Higham (2002), keeping the diagonals unchanged, to ensure positive-semidefiniteness (PSD). Once a PSD matrix is obtained, the corresponding probabilities $\boldsymbol{Z}_{sim}$ are recovered from the PSD $\boldsymbol{\Sigma}_{\mathbf{D}}$. These joint probabilities are the values that are then used to generate directional forecasts. All other parameters are obtained directly from the past market data, rolling the estimation window forward as new observations become available to the investor.

Once the directional forecasts have been generated, optimal portfolio weights are computed for each trading period utilizing both, the proposed framework,

[13] Only those stocks in the index that have data available since the beginning of the sample period are included – this results in a total of 92 stocks in the dataset. Survivorship bias is not an issue in this study, as the purpose is to compare portfolio selection models against each other.

[14] For each simulation run, $\zeta_i \sim \mathrm{U}(0.5, 0.6)$ for all i, which can be justified, for example, by the study of Bekiros and Georgoutsos (2008) for weekly market index returns.

[15] The function f is set to equal $x_{ij}/2$ in this simulation. In addition, some randomness is introduced through $\eta_{ij} \sim \mathrm{U}(-0.1, 0.1)$.

and a simpler alternative. For this alternative model, which is something that the investor might resort to when directional forecasts are available, the covariance matrix input is the historical unconditional estimate (based on past data). The expected return vector estimate is formed by multiplying the forecasted directions by the average of the mean historical magnitudes of the available assets. Thus, both models utilize the same directional forecasts – the difference being that the proposed framework utilizes the directional accuracies and deals with the information more efficiently. Optimal portfolio weights are computed utilizing the same three criteria as before (minimum variance, maximum reward-to-risk, and maximum geometric mean), and the investor is allowed to go against the directional forecasts. To avoid extremely large weights, the portfolio weights are constrainted in two alternative ways. The first alternative is to set $-0.2 \leq w_i \leq 0.2, \forall i$. The second alternative is to set the sum of absolute weights equal to one, i.e. $\sum_i |w_i| = 1$. Results are reported separately in the case of each of these constraints.

5.2 Generating directional forecasts

While it would be possible to utilize a few well-known forecasting models to generated directional forecasts, this approach would be too limited for the purpose of this study, which should cover as wide a range of situations as possible. Specifying a set of directional forecasting models would create ambiguity, and therefore a different, more general approach is adopted: The directional forecasts are generated artificially, conforming to the specified directional accuracies so that it appears as if they have been generated by a model defined by these parameters.

A straightforward way to generate directional forecasts in a simulation study such as this is to derive them from the actual stock returns by setting $\boldsymbol{s}_t \equiv \mathbf{D}_t \odot \mathrm{sgn}(\boldsymbol{r}_t)$. This approach makes the directional forecasts endogenous as explained in the theoretical section of this paper. To simulate the outcomes of the forecasts, $\mathbf{D}_t$, governed by the probabilities in $\boldsymbol{\zeta}$ and $\boldsymbol{Z}_{sim}$, the method of Leisch et al. (1998) is modified as explained in Appendix B. This approach allows generating directional forecasts utilizing a multivariate normal distribution.

The convenient property of this approach is that there is no need to concern with what the actual forecasting model is – the directional predictions are simply simulated according to the given probabilities and real stock market data. This is a way to obtain directional forecasts as if they came from a forecasting model governed by the probabilities in $\boldsymbol{\zeta}$ and $\boldsymbol{Z}_{sim}$, which vary between the simulation runs. The trading simulation for portfolio performance evaluation is run out-of-sample (i.e. all market parameters are estimated from past data), but the probabilities with which the directional forecasts are correct are known to the investor. The probabilities do not take into account the additional information in the unconditional correlation matrix (as explained in Appendix A), and so they could be made more accurate, but we elect not to do so in order to keep things more explicit in this simulation study.

5.3 Results

The trading period contains 522 out-of-sample weeks with an estimation window of two years (104 weeks). The simulation was run 100 times with different forecasting accuracies and investable assets each time, yielding a total of 52,200 investment decisions or periods for each optimization criterion. The average performance indicators corresponding to each criterion were calculated and Table 4 presents these figures.

Table 4: Trading simulation results

Optimization criterion	**Portfolio performance**			
	Proposed framework		*Alternative*	
Constraint	1	2	1	2
Min. Variance				
Avg. Standard deviation	**0.017**	**0.009**	**0.019**	**0.004**
Avg. Turnover	0.78	0.83	0.33	0.46
Max. Sharpe ratio				
Avg. Sharpe ratio	**0.242**	**0.342**	**0.163**	**0.150**
Avg. Turnover	4.03	1.01	5.87	1.37
Max. Geometric mean				
Avg. Geometric mean	**1.30 %**	**0.58 %**	**0.93 %**	**0.14 %**
Avg. Turnover	7.62	1.03	7.35	1.46

Table compares the average performance of optimal portfolios over a trading simulation of 522 weeks, repeated 100 times with different assets and forecasting conditions, yielding a total of 52,200 investment decisions. Constraint 1 refers to the setting where the absolute values of individual asset weights are restricted to be less than or equal to 20%. Constraint 2 refers to the case where the sum of absolute weights is restricted to be equal to 1. The proposed framework and the alternative utilize the same directional forecasts.

For the minimum variance portfolios, the average standard deviation using the proposed framework under the first constraint is 1.7%, slightly lower than that produced by the alternative, 1.9%. Notice that for the case of global minimum variance portfolios, the alternative framework uses the unconditional covariance matrix as an input in the optimization and hence does not take into account the predicted signs. Since the estimation window is relatively long (104) weeks, the estimate for the unconditional covariance matrix changes quite little from period to period, and hence the optimal weights tend to change less, leading to less trading being done which shows up in the average turnover amounts in Table 4. On the contrary, with the proposed framework, the directional forecasts are taken into account in the conditional covariance matrix according to Equation 19, and hence, more trading takes place each period.

In the case of the second constraint (i.e. sum of absolute weights is equal to one), the alternative framework produces significantly lower average standard deviation, 0.4% compared to 0.9% produced by the proposed framework. However, it should be noted that when using the proposed framework, the covariance matrix estimate depends on the directional forecasts, and hence the optimal portfolios can change radically from period to period. This also makes it possible for the mean of the portfolio to change from period to period. Therefore, examining the out-of-sample standard deviation of the portfolio may not tell much about the true conditional volatility of the portfolio returns each period.

For the maximum reward-to-risk portfolios, the average Sharpe ratio produced by the proposed framework under the first constraint, 0.242, is substantially higher than the corresponding average Sharpe ratio for the alternative model, 0.163. For the second constraint, this performance difference is even larger (0.342 versus 0.150). It should be noted that in the case of maximum Sharpe ratio portfolios, the alternative framework also takes into account the forecasted directions in the expected return vector. As explained in detail above, the scenario changes for each simulation run: for a concrete illustration, Figure 5 shows simulated wealth development for the maximum Sharpe ratio portfolios under the second constraint (absolute weights summing up to one) in three different scenarios where the forecast accuracies and investable assets are different. In each scenario, the optimization problem is solved using both the proposed framework (solid curves in the figure) and the alternative (dashed gray curves).

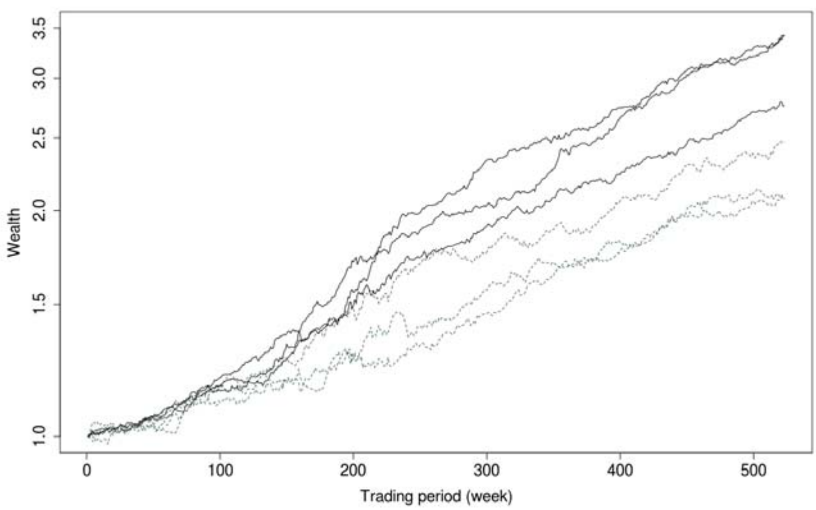

Figure 5: Simulated wealth development of maximum Sharpe ratio portfolios (with absolute portfolio weights summing up to one)

The horizontal axis in Figure 5 denotes the out-of-sample trading period in weeks, and the vertical axis shows the level of wealth relative to starting wealth (equal to one) in log scale. It is evident that in each scenario, the proposed framework produces a more stable wealth development with higher terminal wealth (i.e. higher Sharpe ratio) than what is produced by the alternative. The figure shows just three scenarios for clarity, but the experiment can be repeated for any number of times, and the graph would have a similar appearance.

From Table 4, it can also be observed that the average turnover amount for the Sharpe ratio portfolios is notably lower for the proposed framework than for the alternative under both constraints. This indicates that in addition to providing better performance, the usage of the proposed framework can lead to less trading taking place, and thus, lower transaction costs.

Finally, for the maximum geometric mean portfolios, Table 4 shows that the average geometric mean utilizing the proposed framework under constraint 1 is 1.30% – significantly higher than that produced by the alternative, 0.93%. As is the case with the maximum Sharpe ratio criterion, the performance difference becomes even larger when constraint 2 is in place (0.58% versus 0.14%). This latter performance difference between the two frameworks is illustrated graphically in Figure 6, where the wealth development for three scenarios is plotted utilizing both the proposed framework and the alternative. It should be noted that the terminal wealths produced are unrealistically high due to no trading costs included in the simulation, however, the figure serves the purpose as the objective is to compare the performance of the two frameworks.

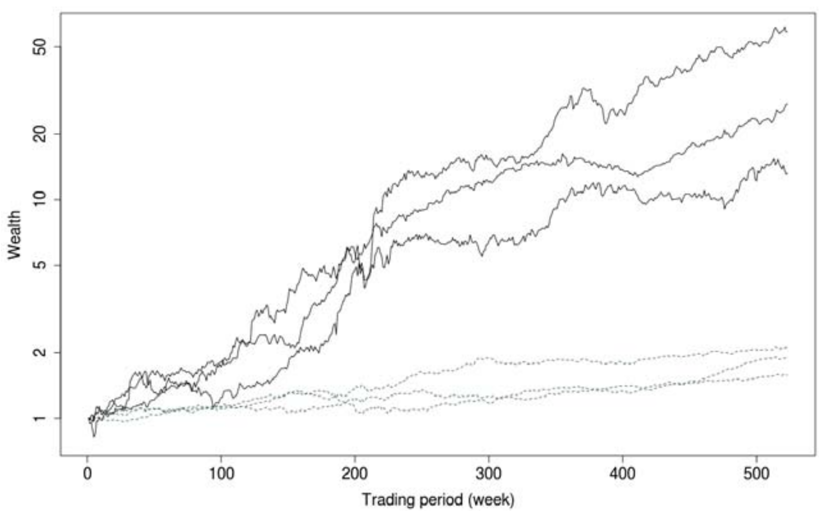

Figure 6: Simulated wealth development of maximum geometric mean portfolios (with absolute portfolio weights summing up to one)

As is evident from Table 4, for the maximum geometric mean portfolios, the average turnover amount for the proposed framework is slightly higher than for the alternative when constraint 1 is in place, however, with the second constraint, the trading incurred in the case of the proposed framework is significantly lower than in the case of the alternative framework.

Overall, the trading simulation shows that notable performance increase for maximum Sharpe ratio and geometric mean portfolios over a simple alternative is possible to achieve by using the proposed framework when directional forecasts are available. Specifically, the simulation has demonstrated that even with the simplifying assumptions behind the proposed framework (most importantly, symmetric unconditional returns with mean zero) not necessarily holding, the performance can still be high compared to a simpler alternative. In other words, applying the proposed model in a trading experiment conducted with authentic stock market data shows that it can be valuable in practice as well.

It should be kept in mind that in this simulation study, the directional accuracies are assumed to be known to the investor. In reality, the performance increase is not expected to be this high as the accuracy estimates contain some error. This is a topic that would be interesting to address in future research by applying specific forecasting models and measuring their accuracy from past data. The approach chosen here is more general and encompasses all directional forecasting models with one setting.

6 CONCLUSION

A novel framework is developed for mean-variance portfolio optimization in order to use directional return estimates as inputs. This makes it possible to bypass the difficult, direct estimation of mean returns. When return directions are forecastable, the analytical results in this paper give rise to three propositions: 1) Assets with high (idiosyncratic) return volatility are preferred for higher expected portfolio returns; 2) Correlation between *absolute* returns affects portfolio variance explicitly; 3) High levels of correlation between asset returns, be it positive or negative, is not preferred assuming that it is linked to correlated outcomes of the directional forecasts.

In simulation studies including a wide variety of scenarios, the developed framework is shown to be capable of producing a substantial performance increase compared to simpler alternatives under three optimization criteria (minimum variance, maximum Sharpe ratio, maximum geometric mean). Moreover, a trading simulation with authentic stock market data demonstrates that even if the simplifying assumptions behind the developed model may not hold accurately, employing the model can still provide substantially higher Sharpe ratios or geometric means compared to a simpler alternative.

It should be noted that due to the generalized setting of the trading simulation, the increase in performance relative to alternative models is not likely to be as high in practice as the directional accuracies of the forecasts need to be estimated. Future research could focus on methods for their estimation and conduct a performance analysis utilizing specific forecasting models that are known to have predictive power out-of-sample. Another topic for future research is the concept of absolute return correlation, which emerges as a parameter explicitly affecting conditional portfolio variance in this setting. Finally, the developed framework does not consider higher moments (asymmetry of the conditional distribution), and a possible extension to the model could be derived in order to take this aspect into account.

REFERENCES

Ait-Sahalia, Y. – Brandt, M. (2001) Variable selection for portfolio choice. *The Journal of Finance*, Vol. 56, 1297–1351.

Anatolyev, S. – Gospodinov, N. (2010) Modeling financial return dynamics via decomposition. *Journal of Business & Economic Statistics*, Vol. 28, 232–245.

Ang, A. – Bekaert, G. (2002) International asset allocation with regime shifts. *Review of Financial Studies*, Vol. 15, 1137–1187.

Barber, B. – Lehavy, R. – McNichols, M. – Trueman, B. (2001) Can investors profit from the prophets? Security analyst recommendations and stock returns. *The Journal of Finance*, Vol. 56, 531–563.

Barberis, N. (2000) Investing for the long run when returns are predictable. *The Journal of Finance*, Vol. 55, 225–264.

Bekiros, S. D. (2010) Heterogeneous trading strategies with adaptive fuzzy actor–critic reinforcement learning: A behavioral approach. *Journal of Economic Dynamics and Control*, Vol. 34, 1153–1170.

Bekiros, S. D. – Georgoutsos, D. A. (2008) Non-linear dynamics in financial asset returns: the predictive power of the cboe volatility index. *The European Journal of Finance*, Vol. 14, 397–408.

Black, F. – Litterman, R. (1992) Global portfolio optimization. *Financial Analysts Journal*, Vol. 48, 28–43.

Campbell, J. Y. – Chan, Y. L. – Viceira, L. M. (2003) A multivariate model of strategic asset allocation. *Journal of Financial Economics*, Vol. 67, 41–80.

Chevapatrakul, T. (2013) Return sign forecasts based on conditional risk: Evidence from the UK stock market index. *Journal of Banking and*

Finance, Vol. 37, 2342–2353.

Christoffersen, P. F. – Diebold, F. X. (2006) Financial asset returns, direction-of-change forecasting, and volatility dynamics. *Management Science*, Vol. 52, 1273–1287.

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

Elton, E. J. – Gruber, M. J. (1987) Portfolio analysis with partial information: The case of grouped data. *Management Science*, Vol. 33, 1238–1246.

Estrada, J. (2010) Geometric mean maximization: An overlooked portfolio approach? *The Journal of Investing*, Vol. 19, 134–147.

Grant, D. (1978) Market timing and portfolio management. *Journal of Finance*, Vol. 33, 1119–1131.

Guidolin, M. – Timmermann, A. (2007) Asset allocation under multivariate regime switching. *Journal of Economic Dynamics and Control*, Vol. 31, 3503–3544.

Hallerbach, W. G. (2014) On the expected performance of market timing strategies. *The Journal of Portfolio Management*, Vol. 40, 42–51.

Higham, N. J. (2002) Computing the nearest correlation matrix - a problem from finance. *IMA Journal of Numerical Analysis*, Vol. 22, 329–343.

Jensen, M. C. (1972) Optimal utilization of market forecasts and the evaluation of investment performance. *Mathematical Methods in Finance*.

Kandel, S. – Stambaugh, R. F. (1996) On the predictability of stock returns: An asset-allocation perspective. *Journal of Finance*, Vol. 51, 385–424.

Leisch, F. – Weingessel, A. – Hornik, K. (1998) On the generation of correlated artificial binary data. *Working Paper Series, SFB "Adaptive Information Systems and Modelling in Economics and Management Science", Vienna University of Economics*.

Leitch, G. – Tanner, J. E. (1991) Economic forecast evaluation: Profits versus the conventional error measures. *The American Economic Review*, Vol. 81, 580–590.

Leone, F. C. – Nelson, L. S. – Nottingham, B. (1961) The folded normal distribution. *Technometrics*, Vol. 3, 543–550.

Leung, M. T. – Daouk, H. – Chen, A. S. (2000) Forecasting stock indices: a comparison of classification and level estimation models. *International Journal of Forecasting*, Vol. 16, 173–190.

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

Markowitz, H. M. (1959) *Portfolio selection: efficient diversification of investments*. New York: John Wiley & Sons, Inc.

Michaud, R. O. (1989) The Markowitz optimization enigma: Is 'optimized' optimal? *Financial Analysts Journal*, Vol. 45, 31–42.

Nyberg, H. (2011) Forecasting the direction of the us stock market with dynamic binary probit models. *International Journal of Forecasting*, Vol. 27, 561–578.

Pesaran, M. H. – Timmermann, A. (2002) Market timing and return prediction under model instability. *Journal of Empirical Finance*, Vol. 9, 495–510.

Pfeifer, P. E. (1985) Market timing and risk reduction. *Journal of Financial and Quantitative Analysis*, Vol. 20, 451–459.

Qian, E. – Gorman, S. (2001) Conditional distribution in portfolio theory. *Financial Analysts Journal*, Vol. 57, 44–51.

Rapach, D. – Zhou, G. (2013) Forecasting stock returns. *Handbook of Economic Forecasting*, Vol. 2, 328–383.

Rydberg, T. H. – Shepard, N. (2003) Dynamics of trade-by-trade price movements: Decomposition and models. *Journal of Financial Econometrics*, Vol. 1, 2–25.

Scott, R. C. – Horvath, P. A. (1980) On the direction of preference for moments of higher order than the variance. *The Journal of Finance*, Vol. 35, 915–919.

Sharpe, W. F. (1966) Mutual fund performance. *Journal of Business*, Vol. 39, 119–138.

Skabar, A. (2013) Direction-of-change financial time series forecasting using a similarity-based classification model. *Journal of Forecasting*, Vol. 32, 409–422.

Tobin, J. (1958) Liquidity preference as behavior towards risk. *The Review of Economic Studies*, Vol. 25, 65–86.

Tu, J. (2010) Is regime switching in stock returns important in portfolio decisions? *Management Science*, Vol. 56, 1198–1215.

Womack, K. L. (1996) Do brokerage analysts' recommendations have investment value? *The Journal of Finance*, Vol. 51, 137–167.

APPENDIX A

The directional forecasts can be viewed either as exogenous, with the unconditional return correlation structure following ($\operatorname{sgn}(\boldsymbol{r}) \equiv \boldsymbol{s} \odot \mathbf{D}$), or as endogenous ($\boldsymbol{s} \equiv \operatorname{sgn}(\boldsymbol{r}) \odot \mathbf{D}$), in which case the unconditional correlation structure contains additional information about the conditional returns, and especially the forecast accuracies. The framework presented in the paper is compatible with both of these interpretations. In the latter case, which is adopted in the trading simulation study in this paper, it would be possible to denote the probabilities ζ_i and ζ_{ij} as conditional on the observed directional forecasts $\boldsymbol{s}$. This is because if $\boldsymbol{s} \equiv \mathrm{sgn}(\boldsymbol{r}) \odot \mathbf{D}$, and we knew the correlation structure of $\mathrm{sgn}(\boldsymbol{r})$, the forecast vector $\boldsymbol{s}$ contains information about the outcome D_i for asset i, affecting its expected value and hence, the probabilities ζ_i and ζ_{ij}.

For example, if every directional forecast in $\boldsymbol{s}$ were $+1$ and the unconditional correlation structure were highly positive for all asset pairs, it is more likely that the forecasts for a pair of assets are simultaneously correct or wrong than in the case where the correlation structure is 0 across the asset pairs. If the forecasted directions are endogenous and the unconditional return correlation structure is assumed to be known, then, in general, $\mathrm{E}[\mathbf{D}|\boldsymbol{s}] \neq \mathrm{E}[\mathbf{D}] = (2\boldsymbol{\zeta} - \mathbf{1})$. The same applies for the joint probabilities in $\boldsymbol{Z}_{sim}$. To take this into account, it would be possible to define the probabilities as being conditional on the directional forecast vector $\boldsymbol{s}$.

For several assets, deriving the explicit form of these conditional probabilities from the unconditional parameters analytically would not be a feasible task, but one could resort to simulating their values. In practice, the true levels of the conditional expected values or probabilities are not known because we cannot know the exact distribution of stock returns. However, one very simple alternative to make the probabilities more accurate is to define two separate cases – one for the case when the asset is predicted to go up in value, and one for the case when the value is forecasted to drop. A similar approach can be adopted for the joint probabilities: one for the case when both assets are predicted to move in the same direction, and one for the case when they are predicted to move in different directions. Instead of determining merely ζ_{ij}, we could find a value separately for ζ_{ij}^{+} (for when predicted directions are the same) and ζ_{ij}^{-} (predicted directions are not the same). This idea can be extended further by making a difference between the case where $s_i = 1, s_j = 1$ and the case where $s_i = -1, s_j = -1$, and so forth. This approach should be relatively easily applicable in practice as well.

APPENDIX B

Generating a vector of discrete random variables, $\mathbf{D}$, with $\mathrm{P}(D_i = 1) = \zeta_i$, $\mathrm{P}(D_i = -1) = (1 - \zeta_i)$ and $\mathrm{P}(D_i = D_j) = \zeta_{ij}$ is done by modifying the method of Leisch et al. (1998), who generate multivariate binary data where $\mathrm{P}(D_i = D_j = 1)$ is known.

The idea is to find a mean vector and a covariance matrix for a random vector $\mathbf{R}$ with a multivariate normal distribution, which will generate values such that

$P(\text{sgn}(R_i) = 1) = \zeta_i$, $P(\text{sgn}(R_i) = -1) = (1 - \zeta_i)$ and $P(\text{sgn}(R_i) = \text{sgn}(R_j)) = \zeta_{ij}$. Simulating from a multivariate normal distribution is then possible with any statistical programming language, and $\mathbf{D} \equiv \text{sgn}(\mathbf{R})$ will have the desired properties.

Following Leisch et al. (1998), the appropriate mean vector $\boldsymbol{\mu}$ is generated by setting $\mu_i = \Phi^{-1}(\zeta_i)$, where Φ is the standard normal cumulative distribution function. The variance of $R_i, \forall i$ is set to equal one, so that the generated covariance matrix is also the correlation matrix of $\mathbf{R}$. By denoting $\tilde{R}_i \equiv R_i - \mu_i$ and $\tilde{R}_j \equiv R_j - \mu_j$, the joint probability

$$\begin{aligned}\zeta_{ij} &= P(R_i > 0, R_j > 0) + P(R_i < 0, R_j < 0) \\ &= P(\tilde{R}_i > -\mu_i, \tilde{R}_j > -\mu_j) + P(\tilde{R}_i < -\mu_i, \tilde{R}_j < -\mu_j).\end{aligned} \tag{24}$$

Since $\tilde{R}_i$ and $\tilde{R}_j$ have a standard bivariate normal distribution, the correlation coefficients ρ_{ij} can be obtained from the following equation

$$\zeta_{ij} = \int_{-\mu_j}^{\infty} \int_{-\mu_i}^{\infty} \phi(r_i, r_j; \rho_{ij}) dr_i dr_j + \int_{-\infty}^{-\mu_j} \int_{-\infty}^{-\mu_i} \phi(r_i, r_j; \rho_{ij}) dr_i dr_j, \tag{25}$$

where $\phi(r_i, r_j; \rho_{ij})$ is the standard bivariate normal pdf with correlation ρ_{ij}.

Solving this equation analytically is not feasible, so what is done instead is that the values of the right hand side of the equation are numerically simulated for specific values of ζ_i, ζ_j, and ρ_{ij}. This creates a table for a range of values. The exact values of ζ_i, ζ_j are then rounded to the nearest value used in the simulation, and the corresponding row is selected from the table of simulated values. Finally, the value for ρ_{ij} is obtained by linear interpolation of the simulated values, so that the exact value of ζ_{ij} is closely matched.

Once the correlation coefficients have been generated for each pair of assets, a covariance matrix is formed. An algorithm based on Higham (2002) is then utilized to find the nearest PSD matrix, keeping diagonal values intact. Now $\mathbf{D}$ governed by ζ_i and ζ_{ij} can be generated by simulating values from a multivariate normal distribution with the corresponding mean vector and covariance matrix, and by obtaining the signs of the resulting values.

ESSAY 2

Hämäläinen, Joonas

Correlation in the magnitude of financial returns

Preprint

Correlation in the Magnitude of Financial Returns

Joonas Hämäläinen*

Abstract

Correlation between the magnitudes of asset returns is an overlooked concept in financial research. It affects portfolio return variance explicitly when the directions of returns are predictable. This paper presents a link from Pearson correlation to *magnitude correlation* and examines its empirical levels in the U.S. stock market. Magnitude correlation is time-varying and has increased marketwide since the recent financial crisis. Abnormally high levels can be interpreted as market anxiety and can be linked to the VIX index measuring implied volatility of stock options.

1 INTRODUCTION

Expected returns of financial assets are notoriously difficult to estimate. However, previous research has provided evidence that the mere *directions* of returns can be forecasted out-of-sample, often with significant accuracy (for recent evidence, see e.g. Bekiros (2010), Nyberg (2011), Chevapatrakul (2013), and Skabar (2013)). Theoretically, even if conditional *mean* returns are independent and hence not forecastable, there can still be dependency in the directions of returns (Christoffersen and Diebold (2006)).

Investment managers and analysts commonly voice directional views such as buy/sell recommendations, which can contain predictive power (e.g. Womack (1996), Barber, Lehavy, McNichols and Trueman (2001)), but it is often not reported by how much an asset's value is going to change. It appears that mean returns are typically too difficult to forecast and the task is often simply avoided. Moreover, errors in mean return estimates can render important applications

* University of Turku, Department of Accounting and Finance, Rehtorinpellonkatu 3, 20500 Turku, Finland. Part of the research has been conducted during the author's visit to Princeton University in 2013/2014. The author is grateful for the comments of Luis Alvarez, Stelios Bekiros, Hannu Kahra, Johan Knif, Shanqui Li, and the participants at the 2014 Paris Financial Management Conference.

such as portfolio optimization unusable in practice (see, e.g. Michaud (1989), DeMiguel, Garlappi and Uppal (2009)).

Essay 1 introduced a portfolio selection framework for utilizing directional return forecasts without having to estimate mean returns directly. This framework adds more dimensions to the traditional mean-variance optimization problem: for example, correlation between the *absolute* values, or magnitudes, of asset returns matters explicitly for conditional portfolio variance.

Although correlation of magnitudes is not an unknown concept in general[1], it appears to be a neglected concept in financial research[2]. Its theoretical or empirical properties have not, to the author's knowledge, been documented to the extent of this paper before. Perhaps closest to the topic at hand comes the notion of tail dependency and certain classes of copula functions. A comprehensive review of dependence measures can be found in, for example, Nelsen (2006).

The contribution of this paper is multifaceted: First, it is explored in detail how magnitude correlation should be interpreted in portfolio selection context. A preliminary part of these results have been obtained in Essay 1 – this paper takes the analysis further. Second, an analytical link between standard (Pearson) correlation and magnitude correlation is derived in the case of bivariate normality. This allows comparisons between observed magnitude correlation and normality-induced magnitude correlation. Third, it is examined what levels of magnitude correlation are found in U.S. stock market data for different intervals and time periods. The empirical link between Pearson return correlation and the correlation in the magnitudes is examined as well, and an interpretation for high levels of magnitude correlation is presented. Finally, for portfolio optimization purposes, methods for out-of-sample forecasting of magnitude correlation are evaluated.

It should be noted that even in the unconditional case (i.e. directions of returns not predictable), magnitude correlation affects portfolio returns implicitly through Pearson correlation under joint normality: in this case, the link between Pearson correlation and magnitude correlation is fixed, as is shown in this paper. However, when returns are not multivariate normal, magnitude correlation is not tied to Pearson correlation, and the setting becomes more interesting.

The empirical results in this paper show that for the S&P500 stocks, average pairwise magnitude correlation has shifted to higher positive levels in 2007–2010 and 2011–2014, compared to 2003–2006. Moreover, pairwise levels of magnitude correlation in the later time periods are, in general, more widely distributed compared to the earlier timeframe (2003–2006). The magnitude corre-

[1] Studies in, for example, the field of medical physics have examined correlation between magnitudes of changes (see e.g. Ashkenazy, Ivanov, Havlin, Peng, Goldberger and Stanley (2001) and Ashkenazy, Havlin, Ivanov, Peng, Schulte-Frohlinde and Stanley (2003)).

[2] Only one recent study about financial returns that mentions correlation of absolute values was found by the author, namely, Ivanov, Yuen, Podobnik and Lee (2004).

lation coefficients are on average lower for monthly returns compared to weekly or daily returns. For daily returns, pairwise magnitude correlation levels are considerably higher than that implied by their Pearson correlation under bivariate normality. For monthly returns, magnitude correlation is more in line with the implied, normality-induced correlation.

The average pairwise level of magnitude correlation for the S&P500 stocks does not appear to move in sync with the normality-induced Pearson correlation, and leaves a gap between the two, which we call the *abnormal* level of magnitude correlation in this paper. The marketwide abnormal level for weekly returns is found to be loosely linked to the levels of the VIX index measuring implied volatility of stock options. In general, high levels of magnitude correlation can be interpreted as market anxiety.

The paper proceeds as follows: Section 2 explains what magnitude correlation is and examines how it affects portfolio return distributions. A link from Pearson correlation to magnitude correlation is presented under bivariate normality. Section 3 conducts a comprehensive examination on the empirical levels of magnitude correlation in the U.S. stock market. Section 4 evaluates methods for estimating magnitude correlation out-of-sample for portfolio management purposes. Section 5 summarizes the findings and suggests topics for future research.

2 THEORETICAL PROPERTIES

2.1 Definition of magnitude correlation

In the traditional mean-variance approach of Markowitz (1952, 1959), correlation between asset returns is a crucial element determining the diversification gains for a portfolio of assets. Specifically, correlation affects portfolio return variance directly. If returns are multivariate normal, (Pearson) correlation completely describes the dependency relation between the individual asset returns. When directional forecasts are available for asset returns, Essay 1 shows that conditional portfolio variance is explicitly affected by the *absolute* return covariance between each asset pair,

$$\mathrm{Cov}[|r_i|,|r_j|] = \mathrm{Corr}[|r_i|,|r_j|]\sqrt{\mathrm{Var}[|r_i|]\mathrm{Var}[|r_j|]}. \tag{1}$$

The variances of absolute values are proportional to the variances of the returns; the *magnitude correlation* term, $\mathrm{Corr}[|r_i|,|r_j|]$, is what this paper focuses on. When the directions of asset returns are predictable, this correlation either smooths the return series of the portfolio, or makes it vary more, depending on the value of Corr$[|r_i|, |r_j|]$ and the joint pairwise accuracy of the directional forecasts[3], denoted here by ζ_{ij}.

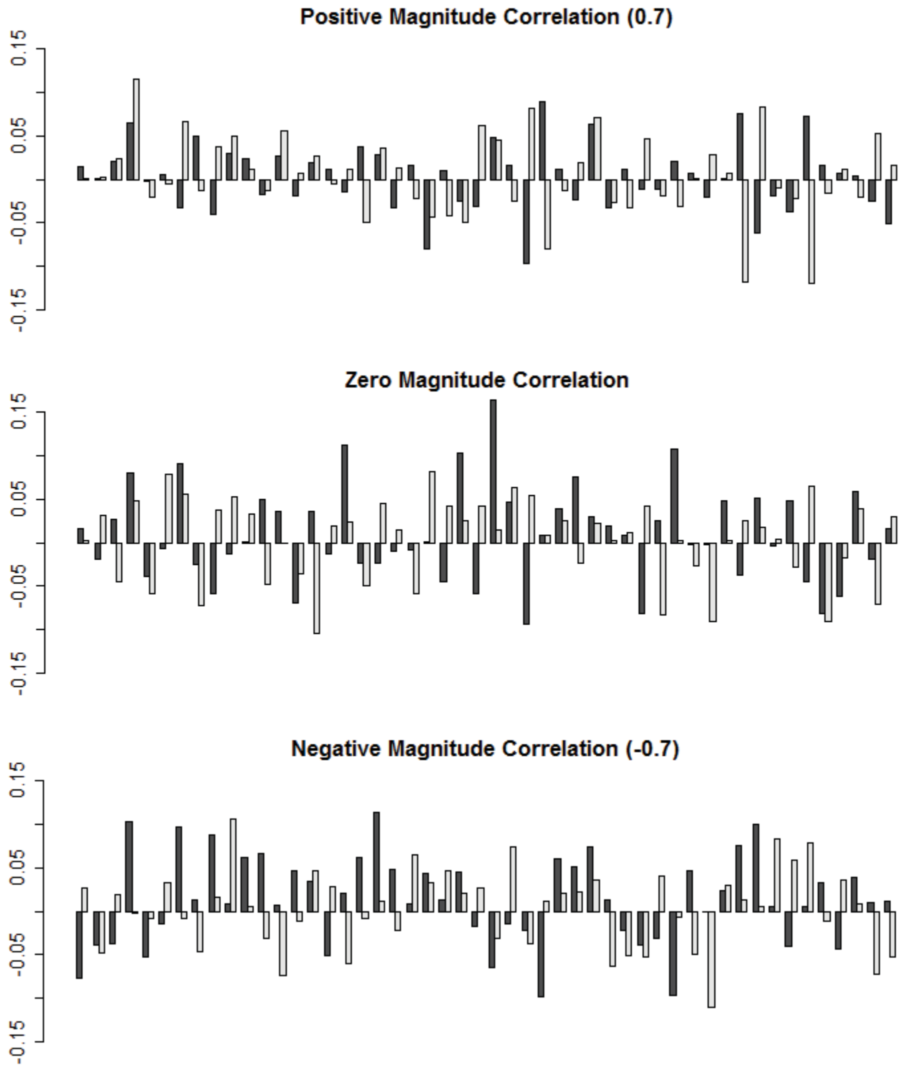

Figure 1: Simulated asset returns exhibiting different levels of magnitude correlation

In practice, *negative* magnitude correlation between two assets implies that when the relative magnitude of the return on one asset is large, the other asset's return tends to be close to zero, and vice versa. By *relative* magnitude, we mean the deviation from the expected magnitude $E[|r_i|]$, which is proportional to the volatility of the asset return. *Positive* magnitude correlation, on the other hand, implies that when the return on one asset deviates largely from its mean

[3] The joint accuracy ζ_{ij} is the probability of being simultaneously either correct or wrong on both assets' directions.

magnitude, the return on the other asset is likely to be large in relative terms as well. Figure 1 illustrates this by showing simulated returns on two normally distributed zero-mean assets with a) highly positive magnitude correlation, b) zero magnitude correlation, and c) highly negative magnitude correlation.

The directions of returns (positive or negative) in Figure 1 are irrelevant; the standard (Pearson) correlation, and hence the sign correlation, for these simulated returns is zero (see Appendix A for details on the simulated return generation). However, it is clear that in the uppermost panel, both assets tend to have large returns in either direction simultaneously. Similarly, small returns tend to go hand in hand. This is a simplified example of positive magnitude correlation. In the middle panel, the return magnitudes are not correlated, and hence appear random. In the lowermost panel, the tendency is that when one asset exhibits a large return in either direction, the other asset has a small return (positive or negative), and vice versa – an example of negative magnitude correlation.

It should be noted that in the univariate case, autocorrelation in absolute returns can be caused by an autoregressive conditional heteroskedastic (ARCH or GARCH) process and time-varying volatility (for details about ARCH/GARCH processes, see Engle (1982) and Bollerslev (1986)). In a similar manner, magnitude correlation in the multivariate case can arise from a multivariate ARCH/GARCH process. For more information on the multivariate processes, see e.g. Tsay (2010). The focus of this paper is on the properties of the correlation itself, and hence possible connections to these econometric models is outside its scope – this would be a potential area for future research.

To construct an analytical framework for portfolio optimization when directional forecasts are available, Essay 1 makes the assumption that unconditional asset returns are symmetrically distributed around mean zero. This assumption/approximation fairly well describes empirical asset returns, especially when the investment horizon is short. In order to keep the analysis comparable to the earlier work, the same assumption is adopted in this paper. This is also convenient for analytical tractability.

2.2 Link to Pearson correlation

An interesting question to follow is: What kind of levels of magnitude correlation can normally distributed (marginal) returns exhibit, i.e. does the joint distribution set limitations?[4] For bivariate normal random variables, magnitude correlation is defined by the Pearson correlation between the random variables alone. However, multivariate normal returns is a strong assumption[5] for finan-

[4] Ultimately, this question could be answered by deriving a suitable class of copula functions, however, this approach is not in the scope of the paper.

[5] In fact, there is only one copula function that produces a jointly normal multivariate distribution, namely, the Gaussian copula.

cial returns, and in practice, wide variation can be expected. Nevertheless, the bivariate normal case serves as a good reference point.

The covariance between the absolute values of two bivariate normal random variables with zero means can be determined by first computing the expectation

$$\mathrm{E}[|r_i r_j|] = \iint\limits_{r_i r_j \geq 0} r_i r_j f(r_i, r_j) dr_i dr_j - \iint\limits_{r_i r_j < 0} r_i r_j f(r_i, r_j) dr_i dr_j, \tag{2}$$

where

$$f(r_i, r_j) = \frac{1}{2\pi\sigma_i\sigma_j\sqrt{1-\rho^2}} \exp\left(\frac{r_i^2/\sigma_i^2 - 2\rho r_i r_j/\sigma_i\sigma_j + r_j^2/\sigma_j^2}{2(\rho^2-1)}\right),$$

i.e. the probability density function of a bivariate normal distribution with zero means; ρ denotes the correlation coefficient between r_i and r_j. Solving the integrals in Equation 2 yields

$$\mathrm{E}[|r_i r_j|] = \frac{2\sigma_i\sigma_j\left(\sqrt{1-\rho^2} + \rho \mathrm{ArcSin}(\rho)\right)}{\pi}. \tag{3}$$

From Appendix B, we know that $\mathrm{E}[|r_i|] = \sigma_i\sqrt{2/\pi}$, and therefore,

$$\mathrm{Cov}[|r_i|, |r_j|] = \frac{2\sigma_i\sigma_j\left(\sqrt{1-\rho^2} + \rho \mathrm{ArcSin}(\rho) - 1\right)}{\pi}. \tag{4}$$

Finally, since $\mathrm{Corr}[|r_i|, |r_j|] = \mathrm{Cov}[|r_i|, |r_j|]/(\sigma_i\sigma_j(1-2/\pi))$, it follows that

$$\mathrm{Corr}[|r_i|, |r_j|] = \frac{2\left(\sqrt{1-\rho^2} + \rho \mathrm{ArcSin}(\rho) - 1\right)}{\pi - 2} \tag{5}$$

for bivariate normal variables with zero means. Figure 2 illustrates this dependency graphically. It is evident that in the bivariate normal case, magnitude correlation can never be negative.

The probability distribution function for the absolute value of a normal variable is known as the folded normal distribution (Leone, Nelson and Nottingham (1961)). The *bivariate* folded normal distribution has been explored in Psarakis and Panaretos (2000). However, the above type of link between magnitude correlation and Pearson correlation has not, to the author's knowledge, been expressed elsewhere.

When we relax the assumption of bivariate normality, magnitude correlation becomes more interesting. Specifically, negative values become plausible. The values depend on the copula function joining the two marginal distributions together. To the best of the author's knowledge, copula functions creating this type of magnitude dependence explicitly have not been developed. A close alternative to this could be the concept of tail dependence (for a review of dependence

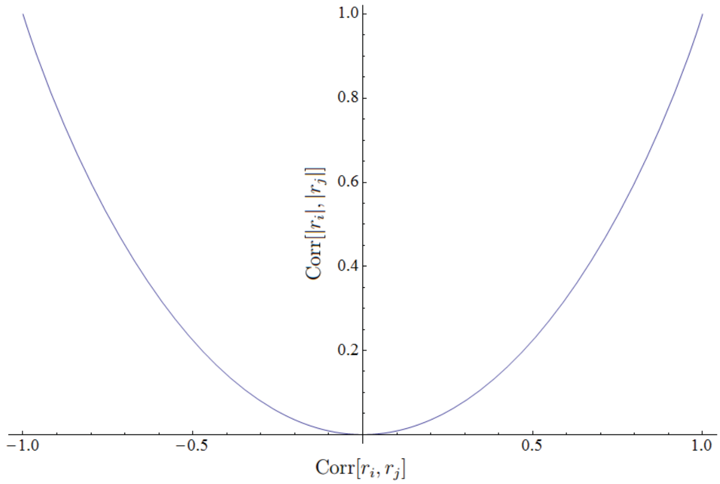

Figure 2: Relationship between Pearson correlation and magnitude correlation for bivariate normal variables with zero means

measures, the reader is referred to Nelsen (2006)). In what follows, a few simplistic examples of some extreme cases of magnitude correlation are given when marginal distributions are normal.

Consider a return for asset i, $r_i \sim \mathrm{N}(0,1)$, and define $r_j \equiv \mathrm{sgn}(r_i)|e|, e \sim \mathrm{N}(0,1)$, where e is independent of r_i (standard deviations are set to equal 1 here for simplicity, even though they are not realistic values for asset returns). The Pearson correlation between r_i and r_j is high, $2/\pi$ to be exact[6]. However, their magnitude correlation is zero because the magnitudes are drawn from normal distributions that are independent. To examine another extreme situation in the form of an example, let $r_k \equiv \mathrm{sgn}(e)|r_i|$. Now the Pearson (and sign) correlation between r_i and r_k is clearly zero, however, the magnitude correlation is clearly equal to one. This and the above case are both clear examples of bivariate non-normality with normal margins.

Finally, can marginally normally distributed returns exhibit negative magnitude correlation? At first, it would seem that one distribution would need to have more mass in the tails of the distribution for this to be possible. On the other hand, observed financial returns are exactly like that, i.e. leptokurtic. However, this feature is also plausible for normal returns. While it is difficult to come up with a simplistic example as in the above cases, consider $r_i \sim \mathrm{N}(0,1)$ and $r_j \sim \mathrm{N}(0,1)$, indepedent of each other. Then arrange the values of r_i by their absolute values, ascending, and do the same for the values of r_j descending. If we now compute the Pearson correlation between these two returns, it should be approximately zero, while the magnitude correlation is highly negative. Fur-

[6] $\mathrm{Corr}[r_i, r_j] = \mathrm{E}[r_i \mathrm{Sgn}(r_i)|e|] = \mathrm{E}[\mathrm{sgn}(r_i)^2|r_i|]\mathrm{E}[|e|] = \sqrt{2/\pi}\,\sqrt{2/\pi} = 2/\pi.$

thermore, if the signs of r_j were transformed to match those of r_i, we would have sign correlation equal to one, and thus a high Pearson correlation coefficient between the returns, yet, the magnitude correlation would stay the same (i.e. highly negative). Therefore, normal distributions as margins do not set strict limitations on the values of magnitude correlation when the form of the joint distribution is not restricted.

As a final note, it is presumed that the correlation between absolute values cannot be perfectly negative (i.e. equal to -1), because the absolute values of normally distributed variables with (approximately) zero mean have skewed distributions to the right. If there is more mass on the left sides of the means (values close to zero) compared to the right sides, it would appear that magnitude correlation cannot be perfectly negative. The rank correlation, however, can clearly be equal to -1. In simulations, the Pearson correlation coefficient for absolute values can reach values smaller than -0.8, so this limitation is likely to not have much practical importance.

2.3 Effect on portfolio variance and higher moments

When returns have a multivariate normal distribution, Pearson correlation gives a complete description of the dependency between the individual returns. The shape of a portfolio return distribution is normal as well in this case. However, when the assumption of joint normality is relaxed, and the fixed link between Pearson correlation and magnitude correlation is broken, the shape of the portfolio return distribution is affected explicitly by correlation in the magnitudes of returns. To see this in effect, consider simple equal-weighted two-asset portfolio return distributions depicted in Figure 3 (each asset has a 5% standard deviation, and zero Pearson or sign correlation between the returns).

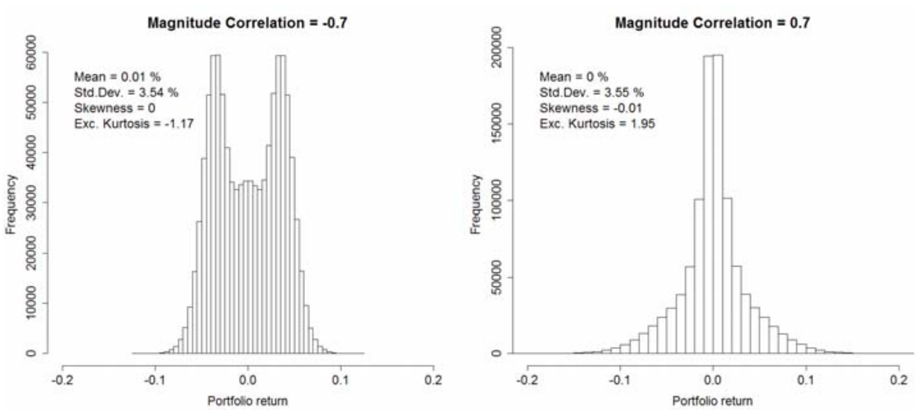

Figure 3: Unconditional two-asset (equal-weighted) portfolio return distributions

In Figure 3, the effect of magnitude correlation on the shape of the portfolio return distribution can be seen. In the histogram on the left, magnitude correlation is strongly negative (-0.7), and portfolio return distribution has two peaks, with its excess kurtosis being negative (-1.17). On the contrary, when magnitude correlation is large and positive (0.7), the portfolio return distribution is clearly leptokurtic with excess kurtosis value of 1.95. However, the variance of portfolio returns is not affected by magnitude correlation.

Also noteworthy is that the distribution on the left has two peaks, i.e. it is bimodal. The distribution is more "uniformly" shaped than the distribution on the right where magnitude correlation is high and positive. This implies that the density of portfolio returns on the left distribution is relatively high for a wide variety of returns, whereas in the case of the latter distribution, the density is concentrated at zero but at the same time the range appears to be wider. This means that the majority of returns tend to be close to zero, but extreme events are more likely than in the case of negative magnitude correlation.

When directional forecasts for asset returns are available, it can be the case that magnitude correlation affects portfolio variance explicitly: Essay 1 presents a framework for portfolio selection to use directional forecasts as inputs in portfolio optimization. In this framework, conditional portfolio variance is affected by the probability of the directional forecasts for each asset pair being simultaneously either correct or wrong, denoted by ζ_{ij}. If this probability is high (close to one) and the investor follows the views generated by the forecasting model, then *negative* magnitude correlation between assets is to be preferred, and vice versa[7]. To be exact, the conditional covariance[8] of returns on a pair of assets i and j,

$$\mathrm{Cov}_{ij|\Omega} = (2\zeta_{ij} - 1)\mathrm{Cov}[|r_i|, |r_j|] + [(2\zeta_{ij} - 1) - (2\zeta_i - 1)(2\zeta_j - 1)]\mu_{abs,i}\mu_{abs,j}, \quad (6)$$

where ζ_i denotes the probability of the directional prediction being correct for asset i, and $\mu_{abs,i} \equiv \mathrm{E}[|r_i|]$.

In essence, for two assets that have a high joint probability ζ_{ij} indicating that the investor is simultaneously correct or wrong on the directions of these assets' returns, *negative* magnitude correlation "smooths" the portfolio return, thus decreasing its variance. In this case, for minimal variance, one would prefer highly negatively correlated absolute returns since going from being wrong on both assets directions to the case where the investor is correct on both directions is smoothed by negative magnitude correlation. If the correlation were large and positive, then the returns for the portfolio would vary more wildly from large and negative to large and positive. In practice, ζ_{ij} can be assumed to have values greater than 0.5 due to asset returns being correlated, and hence negative mag-

[7] See Essay 1, Figure 1, for illustration of the nonlinear relationship that magnitude correlation and joint probability have in affecting conditional pairwise covariance.

[8] Conditional on forecasted directions and directional accuracies, contained in the information set Ω.

nitude correlation is preferred for lower portfolio variance. Figure 4 captures conditional two-asset portfolio return distributions for high values of ζ_{ij}.

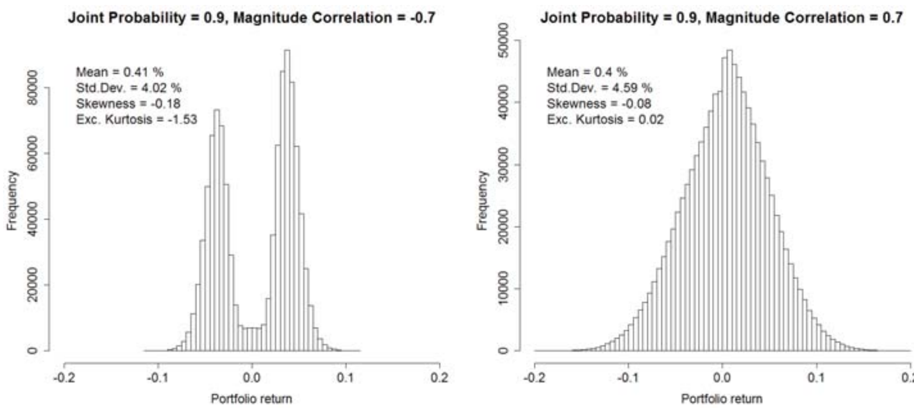

Figure 4: Conditional two-asset (equal-weighted) portfolio return distributions

The portfolio return distributions in Figure 4 are for equal-weighted portfolios of two assets (each with a 5% standard deviation, zero Pearson and sign correlation between the unconditional returns), investing 50% in asset 1 and 50% in asset 2 according to directional predictions, which are simulated for each period (see Essay 1, Appendix B for the method of generating directional forecasts governed by the probabilities ζ_i and ζ_{ij}). That is, the investor goes either long or short on the asset based on the directional view. In this simulation, the probability of being correct, ζ_i, is set equal to 55%, which is a relatively modest and plausible accuracy for a directional forecasting model.

The left histogram in Figure 4 depicts a case of highly negative (-0.7) magnitude correlation combined with a high probability of being simultaneously correct on both assets' directions (0.9). First, from the figure it is apparent that the shape of the distribution is affected by the negative magnitude correlation – the distribution has negative excess kurtosis. The distribution is also slightly negatively skewed. More importantly, the standard deviation of this portfolio return is about 4.0%. On the right histogram in Figure 4, the scenario with high positive magnitude correlation (0.7) is illustrated. This produces a higher standard deviation for the portfolio return (about 4.6%), and there appears to be close to no excess kurtosis for this distribution. As was the case with unconditional returns in Figure 3, the distribution on the left is bimodal, whereas the distribution on the right is not.

Negative magnitude correlation produces asset returns that can smooth the conditional portfolio return, producing lower portfolio standard deviation compared to the case where magnitude correlation is highly positive. The portfolio optimization framework introduced in Essay 1 is simplistic in the sense that it does not take into account higher moments. Without delving into utility theory at this point, it is useful to mention that previous research has shown that positive skewness in portfolio returns is preferred to negative skewness, and low kurtosis is preferred to high kurtosis (see e.g. Scott and Horvath (1980)). The scenario producing the lowest standard deviation, i.e. negative magnitude correlation, also produces negative excess kurtosis. Notice that in both of the cases, the mean portfolio return stays the same, as is to be expected since the mean depends only on the accuracy of the individual forecasts, set equal to 55% in this simulation.

As was seen in Figure 3, even in the unconditional case (i.e. returns not forecastable) the shape of the portfolio return distribution changes according to the magnitude correlation level. If returns are bivariate normal, then the magnitude correlation is linked to Pearson correlation and there is no deviation from a normal distribution for the unconditional portfolio returns. However, when joint normality is not present, magnitude correlation becomes an explicit factor. Therefore, it is important to characterize its empirical properties even if returns do not exhibit directional predictability.

3 EMPIRICAL LEVELS OF MAGNITUDE CORRELATION

3.1 Data and setup

In this section, it is examined what kind of levels of magnitude correlation are found in the U.S. stock market returns. Naturally, we cannot know the true disributions of asset returns at any given point in time, and hence they have to be estimated. As a starting point, we examine time series of data and compute the sample absolute correlation as follows:

$$\widehat{\mathrm{Corr}[|r_i|,|r_j|]} = \frac{\sum_{k=1}^{T}(|r_{i,k}| - \widehat{\mu}_{abs,i})(|r_{j,k}| - \widehat{\mu}_{abs,j})}{\sqrt{\sum_{k=1}^{T}(|r_{i,k}| - \widehat{\mu}_{abs,i})^2 \sum_{k=1}^{T}(|r_{j,k}| - \widehat{\mu}_{abs,j})^2}} \tag{7}$$

where T denotes the length of the time series, and $\widehat{\mu}_{abs,i}$ denotes the sample mean of $|r_i|$.

The data are closing levels for total return indices of S&P500 stocks from December 31, 2002 to December 31, 2014. Only those stocks in the index that have a history extending throughout this timeframe are included in the analysis – this results in a total of 443 individual stocks. The daily, weekly, and monthly closing levels are each examined separately. In addition, magnitude correlation levels are examined for three indices of stock and commodity markets (the S&P500 index, the FTSE100 index, and the London Metal Exchange index).

From the closing levels, simple net returns for daily, weekly, and monthly data are computed.

To begin with, a simple analysis of magnitude correlation is performed by dividing the total sample into three subsamples of equal length (2003–2006, 2007–2010, and 2011–2014). This division is motivated by the fact that levels of magnitude correlation may differ in different economic conditions, and especially the middle subsample is known for turbulent times in the stock market. Furthermore, we compare the magnitude correlations in the total sample to the observed levels of standard (Pearson) return correlation. This is motivated by the proposition made in Essay 1, that high positive or negative sign correlation in asset returns is likely to give rise to higher levels of joint probability ζ_{ij}. Hence, it would be interesting to see if asset pairs with high standard correlation with close to zero or negative magnitude correlation can be found. Moreover, a look at rolling (simple moving average and exponentially weighted) correlations is offered to examine the evolvement of magnitude correlation over time. Finally, the gap between the observed and normality-induced magnitude correlation is examined.

3.2 Sample magnitude correlations

Sample magnitude correlations are estimated according to Equation 7 for the three subsamples in the case of daily, weekly, and monthly returns. The different subsamples represent different times and economic conditions in the market, and hence, different levels of pairwise magnitude correlations can be expected. Figure 5 presents the histograms of all pairwise magnitude correlations along with simple descriptive statistics.

From Figure 5, it is evident that the average pairwise correlation between the absolute values of returns varies between the different subsamples and time intervals. Most notably, the mean level of magnitude correlation appears to be larger for daily returns compared to the weekly or monthly return intervals. Moreover, the first subsample (2003–2006) features lower average pairwise magnitude correlations than the two later subsamples. It appears that during the turbulent times in the market, in 2007–2010, mean level of magnitude correlation was the highest out of the three subsamples. Still, after this turbulent time, in the latest subsample (2011–2014), average magnitude correlation has remained in relatively high levels compared to the first subsample (2003–2006).

It can also be observed from Figure 5 that the distributions are significantly wider in the lower frequency cases, which is quantified by the higher standard deviation of the distributions. However, one must notice that monthly returns contain less observations, therefore presenting less accurate estimates of these correlations, so part of the increased variation may be due to that.

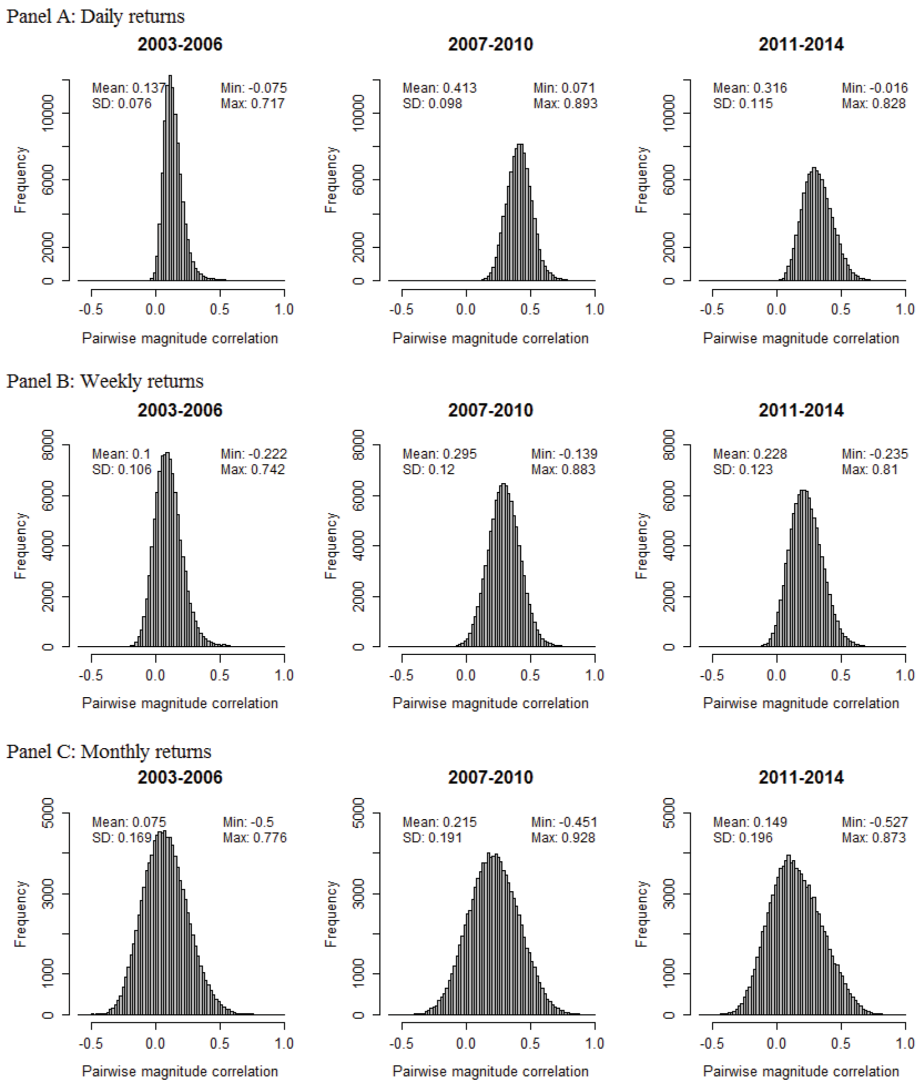

Figure 5: Histograms of pairwise magnitude correlations for S&P500 stocks

From a portfolio management perspective, it is interesting to note that monthly returns present a minimum level of magnitude correlation at -0.527 (in the latest subsample), considerably lower than that exhibited for the lowest case for daily returns, -0.075 (in the first subsample), or weekly returns, -0.235 (in the latest subsample). The significant negative magnitude correlations may present some useful opportunities to a portfolio manager with directional return forecasts, as described earlier in Section 2.

3.3 Relationship to Pearson correlation of returns

In the theoretical section of this paper, a formula for the magnitude correlation value as a function of standard (Pearson) correlation in the bivariate normal case was derived (Equation 5). It is interesting to see what the relationships are like empirically. Additionally, a portfolio manager utilizing the directional optimization framework developed in Essay 1 is in general interested in finding assets that exhibit near zero Pearson correlation and/or negative values for magnitude correlation. With jointly normally distributed variables, this case would not be plausible, as zero Pearson correlation implies zero magnitude correlation as well.

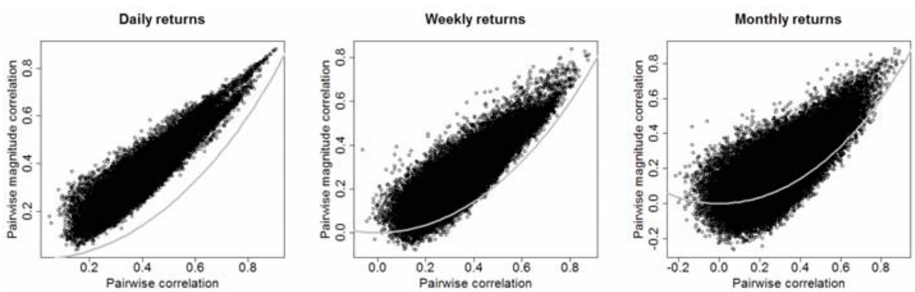

Figure 6: Relationship between standard correlation and magnitude correlation for the S&P500 stocks in the full sample

Figure 6 plots all pairwise magnitude correlations and standard return (Pearson) correlations for the 443 stocks' daily, weekly, and monthly returns. In each plot, the gray curve depicts the *theoretical* values implied by the bivariate normal distribution (i.e. it plots the function in Equation 5). A clear dependency structure can be observed as is to be expected – higher standard correlation implies higher magnitude correlation in general. Interestingly, however, for daily returns, all values of magnitude correlation lie above the reference curve. Based on this, we could reject bivariate normality and focus on picking assets with desirable properties from a portfolio management perspective, as described earlier.

For weekly returns, there is a bit more dispersion and part of the observations are actually below the curve depicting the bivariate normal case. This presents more possibilities for portfolio optimization, and even more so in the case of monthly returns, where the lowest values of magnitude correlation are below -0.2. For subsamples, there are even lower negative values as is evident from Figure 5 above. For monthly returns, variation around the reference curve is large, giving more reason to examine magnitude correlation from a portfolio management perspective. However, also notice that the monthly correlations appear to be more in line with bivariate normality (the shape of the plotted values resembles the reference curve).

3.4 Time-varying correlations

Since the sample magnitude correlation values appear to be significantly different between the subsamples, it is reasonable to assume that they are time-varying. As a simple way to examine the evolvement of magnitude correlation in time, we compute rolling correlations for the entire time series between each asset pair, and then average these out to get a marketwide representation. Additionally, the same procedure is applied to Pearson correlations to see if they move in sync with the magnitude correlations. For daily returns, two estimation window lengths are considered: 60 days and 250 days (corresponding roughly to one quarter and one year in trading days). For weekly returns, the estimation windows are 52 weeks and 208 weeks. For monthly returns, only one estimation window length is considered, 48 months. Figure 7 presents the resulting graphs.

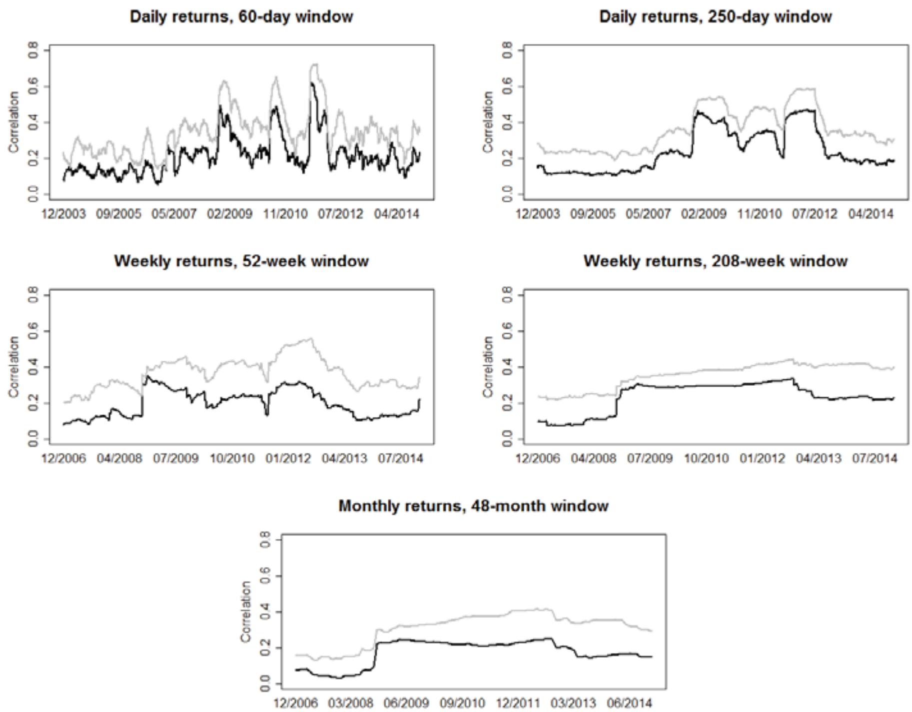

Figure 7: Average pairwise rolling correlations for S&P500 stocks in the full sample (black curve = magnitude correlation, gray curve = standard correlation)

First, it is evident that magnitude correlation, like standard correlation, does not appear to be time invariant. Second, from Figure 7 it can be seen that the average (marketwide) pairwise magnitude correlation moves closely in sync with the average pairwise standard correlation – this is the case for daily, weekly, and monthly returns with all tested estimation windows. In general, the level for magnitude correlation and standard correlation has been lower in the beginning of the sample, then increasing in the mid-part of the sample and appearing to fluctuate more wildly, and then again decreasing in the late sample. This could of course be partly due to a few extreme observations during the recent financial crisis.

It is important to notice that the curves presented in Figure 7 are formed of average pairwise correlations – they do not tell the whole story about individual assets. As an example, a few assets are selected for comparison. Figure 8 presents examples of rolling correlations for unique asset pairs and market indices.

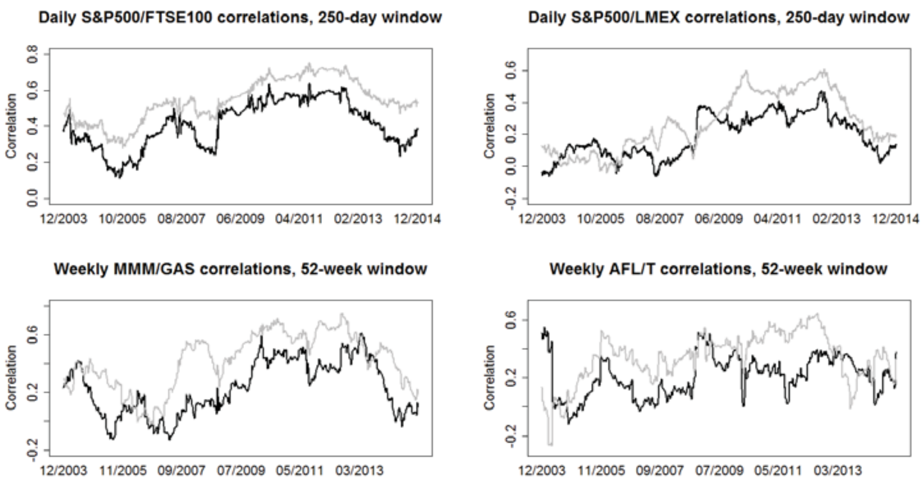

Figure 8: Pairwise rolling correlations for selected stocks and indices (black curve = magnitude correlation, gray curve = standard correlation)

In the upper left hand corner of Figure 8, the rolling magnitude correlation and standard correlation are presented for the daily returns of the S&P500 index and the FTSE100 index. It is interesting to note that there are points in time when the standard correlation and magnitude correlation seem to be at almost equal levels, but also time periods when there is a notable gap between the two. Another interesting case in Figure 8 is the lower left hand corner, where rolling correlations for two individual stocks' (New York Stock Exchange ticker symbols MMM and GAS) weekly returns are presented. As with the previous example, there are times when the magnitude correlation and standard return correlation appear to be nearly equal, and other times when there is a large gap between the two. The variation over time seems to be particularly large, with magnitude correlation being close to zero in the early part of the sample, and then rising to levels close to 0.5 after mid-sample, only to drop back to zero at the end of the sample.

The upper right hand corner in Figure 8 presents a slightly more stable, yet still fluctuating case with two indices (S&P500 and LMEX). Interestingly, the magnitude correlation between the two appears to exceed the standard return correlation level in several points in time during the first half of the sample. The two correlations do not really appear to move closely in sync, however, this may be due to the relatively long estimation window (250 days). These examples in Figure 8 illustrate that it is likely that different asset pairs present different relationships between their standard correlation and magnitude correlation, thus giving rise to more options from a portfolio management perspective.

Finally, as the simple rolling correlation measure, or the moving average (MA) model, gives equal weight to all observations in the estimation window, a slightly more sophisticated way of estimating the correlations is applied by using an exponentially weighted moving average (EWMA) model, defined as

$$\mathrm{Cor}_{EW}\widehat{(|r_i|, |r_j|)} = \frac{\sum_{k=1}^{T} \gamma(1-\gamma)^{k-1}(|r_{i,k}| - \widehat{\mu}_{abs,i})(|r_{j,k}| - \widehat{\mu}_{abs,j})}{\sqrt{\sum_{k=1}^{T} \gamma(1-\gamma)^{k-1}(|r_{i,k}| - \widehat{\mu}_{abs,i})^2 \sum_{k=1}^{T} \gamma(1-\gamma)^{k-1}(|r_{j,k}| - \widehat{\mu}_{abs,j})^2}}, \tag{8}$$

where γ denotes the smoothing weight, set equal to 0.1 (slightly higher than the 0.06 of RiskMetrics, in order to facilitate the use of a shorter estimation window) in this section. We perform an analogous EWMA estimation for the standard correlation between asset returns, and plot the two in Figure 9 for average pairwise correlations for the weekly and monthly S&P500 stock returns.

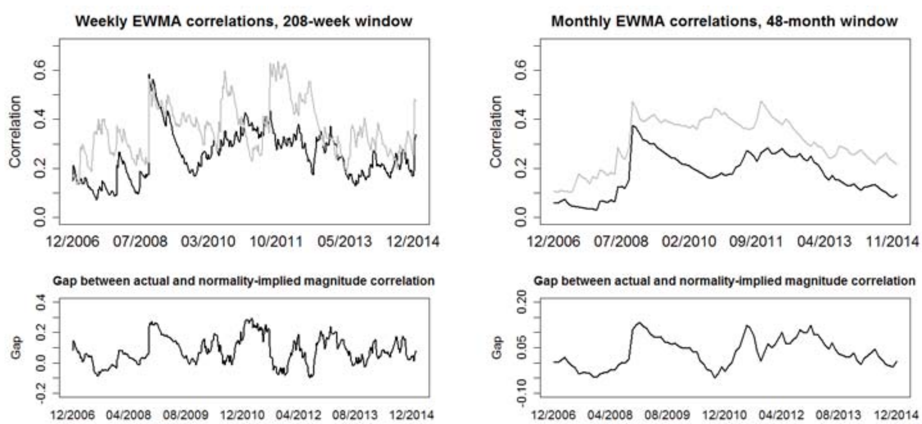

Figure 9: Average pairwise EWMA correlations for S&P500 stocks (black curve = magnitude correlation, gray curve = standard correlation in the top two graphs)

Figure 9 also includes two additional graphs depicting a measure called "Gap"; this is the gap between the observed and implied theoretical magnitude correlation in the bivariate normal case. It follows straightforwardly from Equation 5, and is computed as follows:

$$\text{Gap} \equiv \widehat{\text{Corr}[|r_i|, |r_j|]} - \frac{2\left(\sqrt{1-(\widehat{\rho})^2} + \widehat{\rho}\text{ArcSin}(\widehat{\rho}) - 1\right)}{\pi - 2}, \tag{9}$$

where $\widehat{\rho}$ is the sample estimate of the Pearson correlation coefficient. In Figure 9, the gap measures are averaged out for all stocks to get a marketwide representation.

From Figure 9 it is evident that the spikes or shocks in the measured correlations do not persist as long as in the simple MA case, which is natural to expect due to the EWMA model emphasizing recent observations more. The separate gap graphs in Figure 9 highlight "odd" times in the market, when the levels of magnitude correlation have been out of sync compared to normality-implied magnitude correlation. For example, for the weekly returns, there is a clear jump up near the end of 2008, which is also evident in the upper graph depicting the measured correlations – in this case, magnitude correlation jumped up much more than expected by the theoretical relationship implied by bivariate normality. A different scenario takes place close to the end of the year 2011 when the gap drops significantly – this is explained in the upper graph by the fact that standard correlation jumps up while magnitude correlation remains at the same level or even declines slightly. Naturally, the length of the estimation window and the level of γ in the EWMA model affect the measured levels of correlation and the raggedness of the graphs.

For monthly returns, Figure 9 shows a clear jump in the gap near the beginning of the sample (close to the end of the year 2008), and the upper graph indicates that this is caused by the magnitude correlation level rising exceptionally high relative to the corresponding standard correlation. The gap then slowly returns close to zero, as the magnitude correlation declines faster than the standard correlation does. It is important to notice that the gap measure does not depict the actual gap between the observed magnitude correlation and standard correlation, but the gap between the observed magnitude correlation and the level *implied* by the observed standard (Pearson) correlation.

3.5 What do abnormal levels of magnitude correlation signal?

So far nothing has been said about the possible interpretations of high magnitude correlation in asset returns. Here an explanation is offered and we see whether it can be backed up by the data. In the case of high positive magnitude correlation, the returns of two assets tend to be simultaneously either large or small in magnitude relative to their corresponding mean magnitudes. If the average magnitude correlation in the market is abnormally high, it resembles a time when markets are edgy, even panicking, indicated by prices that react with larger than normal movements – there is strong marketwide reaction to any kind of news, but markets are relatively quiet on no news. In essence, high marketwide magnitude correlation would appear to signal that the market is paying close attention to big news that affect the market as a whole, and do not react strongly to small news about individual sectors or stocks. Based on this reasoning, one could draw the conclusion that high abnormal magnitude correlation might be a sign of anxiety in the market.

As mentioned earlier, magnitude correlation could be linked to multivariate ARCH/GARCH processes and time-varying volatility. Return volatilities could be dependent across assets: there can be an economic force that drives the volatility of several assets' returns up simultaneously, which would manifest as positive magnitude correlation. This force could be, for example, a specific risk factor of a certain sector of companies. If the marketwide average magnitude correlation is high, this could imply that several assets are experiencing an increase or decrease in their return volatility simultaneously.

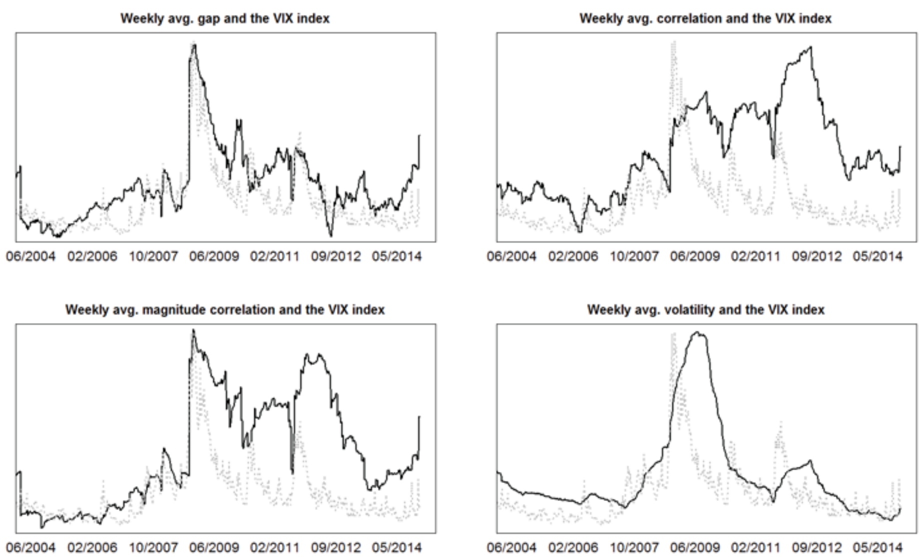

Figure 10: Weekly VIX index compared to different correlation measures and average volatility (dashed gray curve depicts the weekly VIX index values)

A common measure for market anxiety is the VIX index, which tracks the implied volatility of the S&P500 index options. It seems reasonable to compare the movements of the VIX index and that of the gap between observed and theoretically implied magnitude correlation, i.e. abormal magnitude correlation. Figure 10 plots the weekly VIX values and the marketwide abnormal magnitude correlation estimated using a MA52 model (i.e. computing the average gap over the past 52 weeks). For comparison, depicted are also the historical average pairwise standard correlation, the historical average pairwise magnitude correlation, and the average asset return volatility based on the past 52 weeks.

From Figure 10, it appears that the VIX index and abnormal magnitude correlation (i.e. the gap) tend to be high simultaneously and vice versa, low levels also tend to coincide, which can be thought of as giving support to the earlier interpretation of market anxiety. Especially the spike coinciding with the VIX near mid-sample and then quickly declining fits the picture well. Although the other measures also follow the levels of the VIX to some extent, they appear to be less accurately tracking the VIX compared to the gap measure. For example, average historical volatility appears to track the VIX with a lag, as is to be expected. The standard correlation and magnitude correlation appear to remain too high after the mid-sample spike compared to the VIX or the gap. However, it should be noted that the choice of estimation window length has an effect on the fit, and, for example, a longer estimation window does not produce graphs which line up with the VIX this well.

The augmented Dickey-Fuller test indicates that neither the VIX index nor the abnormal magnitude correlation (i.e. the gap) is stationary, so the first differences of the VIX and the abnormal magnitude correlation are taken. The correlation coefficient of the first differences is 0.102 between the VIX and the abnormal magnitude correlation, indicating that the changes do not appear to be strongly correlated. Overall, based on Figure 10, it appears that magnitude correlation raises to exceptionally high levels with anxiety in the market as proxied by the VIX. However, the two appear to be only loosely linked, and there might be other factors at play. Interestingly, the abnormal level of magnitude correlation appears to be rising steeply at the end of the sample, while VIX stays at a relatively low level. This could tell about another type of market anxiety, to which the expected (implied) volatility in the market does not react. This is a potential topic for further research.

4 EVALUATION OF OUT-OF-SAMPLE FORECASTS

As illustrated above, magnitude correlation can have an effect on the shape of the portfolio return distribution in the unconditional case, and explicitly affect the conditional portfolio return variance when the directions of returns are predictable. For portfolio optimization, more accurate estimates of parameter values are important. If one can obtain information about the future values of the parameters required as inputs, the resulting portfolio performance is expected to be enhanced. In this section, the aim is to evaluate a few simple methods for estimating magnitude correlation out-of-sample.

Since it is impossible to know the true levels of stock market correlations at any given time period, direct evaluation of forecast quality is not a plausible method. However, we know from above that under the framework introduced in Essay 1, magnitude correlation affects portfolio variance directly. For this reason, we can test different estimation models for magnitude correlation and compare resulting *out-of-sample* optimal portfolio variances to find out which of the methods produces the most accurate forecasts. This kind of an indirect evaluation, as described in Patton and Sheppard (2009), is useful for practical portfolio management purposes.

The earlier introduced dataset of S&P500 stocks is used in this study as well. The idea is to estimate the weekly[9] one-period ahead covariance matrix, solve the global minimum variance portfolio weights, and then measure the "realized" standard deviation of the portfolio,

$$\widehat{\sigma}_{p,t} = \sqrt{\boldsymbol{w}_t^{*\prime} \boldsymbol{r}_t \boldsymbol{r}_t' \boldsymbol{w}_t^*}, \tag{10}$$

where $\boldsymbol{w}_t^*$ denotes the vector of optimal portfolio weights, and $\boldsymbol{r}_t$ is the return vector at time t. If one estimation method consistently produces lower values for Equation 10, it can be considered the more accurate estimation method (Patton and Sheppard (2009)). Notice that we cannot simply compute the realized portfolio standard deviation over the entire time frame, because the mean return of the optimal portfolios changes at each time period due to (likely) time-variation in return distributions. The realized standard deviation must be computed at each time period separately, and then the mean of these values can be used as an indicator of the performance of the evaluated estimation model.

In order to evaluate magnitude correlation forecasts, a trading simulation study is conducted where the directions of stock returns are predictable as in the directional optimization framework of Essay 1. In this case, magnitude correlation (i.e. the correlation of absolute values of returns) affects the conditional covariance matrix of the returns. Specifically, Equation 6 can be generalized into matrix form, resulting in the conditional covariance matrix (in the case of marginal normal distributions with zero means for the unconditional returns)[10]:

$$\hat{\boldsymbol{\Sigma}}_t|\Omega = \boldsymbol{s}_t \boldsymbol{s}_t' \odot \left((2\boldsymbol{Z}_{sim} - \mathbf{J}) \odot (1 - (2/\pi)) \hat{\mathbf{S}} \hat{\mathbf{R}}_{abs,t} \hat{\mathbf{S}} + \frac{2}{\pi} \hat{\mathbf{S}} \boldsymbol{\Sigma}_{\mathbf{D}} \hat{\mathbf{S}} \right), \tag{11}$$

where $\boldsymbol{s}_t$ denotes the vector of predicted signs, $\boldsymbol{Z}_{sim}$ is a matrix of the joint probabilities ζ_{ij}, $\mathbf{J}$ denotes a matrix of ones, $\hat{\mathbf{S}}$ denotes a diagonal matrix containing the sample-based standard deviaton estimates, and $\hat{\mathbf{R}}_{abs,t}$ is the estimated magnitude correlation matrix. $\boldsymbol{\Sigma}_{\mathbf{D}}$ is the covariance matrix of the forecast outcomes, and it can be expressed with the probabilities ζ_i and ζ_{ij}.

[9] We choose the weekly interval here as it represents a middle-ground between the daily and monthly returns.

[10] For a detailed explanation of the model, see Essay 1.

Notice that the factors affecting the covariance matrix, and hence the global minimum variance portfolio, are the magnitude correlations, standard deviations of asset returns, and the directional forecast probabilities. To give as much weight to magnitude correlation as possible, we need to eliminate the effect of standard deviations and forecast probabilities. Fortunately, this is relatively easily achieved by i) setting the probabilities $\zeta_i = 0.5$ for all i (implying no predictability) and setting ζ_{ij} as high as possible, i.e. equal to one, ii) examining assets in groups exhibiting near equal sample standard deviations and setting the values exactly equal for the estimates.

Part i above implies that all forecast outcomes are perfectly correlated, but there is no real power in the forecasts. The framework can accomodate this, and the term $2\boldsymbol{Z}_{sim} - \mathbf{J}$ now becomes simply a matrix of ones, as does $\boldsymbol{\Sigma}_{\mathbf{D}}$. Hence, the covariance matrix in Equation 11 is simplified to

$$\hat{\boldsymbol{\Sigma}}_t|\Omega = \boldsymbol{s}_t\boldsymbol{s}_t' \odot \left((1-(2/\pi))\hat{\mathbf{S}}\hat{\mathbf{R}}_{abs,t}\hat{\mathbf{S}} + \frac{2}{\pi}\hat{\mathbf{S}}\mathbf{J}\hat{\mathbf{S}}\right). \tag{12}$$

Since $\zeta_{ij} = 1$ for each asset pair, this means that the directional forecasts $\boldsymbol{s}_t$ are simultaneously either correct or wrong for all assets. For this reason, in the simulation, the directional forecasts can be generated by setting $\boldsymbol{s}_t = \mathrm{sgn}(X_t)\mathrm{sgn}(\boldsymbol{r}_t)$, where $X_t \sim \mathrm{U}(-1, 1)$. Finally, to eliminate the effect of return standard deviations, and to give emphasis on magnitude correlation estimation, the stocks are grouped into 88 categories based on their sample standard deviations (the full sample contains 443 stocks of which the remainder in the division by 5 is left out, i.e. three stocks are not included in the study). Each group consisting of 5 stocks is treated separately in the simulation, and the values for standard deviation estimates are set equal to the group's mean. This way, the standard deviations play no part in optimization, and the actual volatility of the asset returns should be relatively similar.

A start-up estimation window of 200 weeks is chosen, leaving over 400 weeks for the out-of-sample evaluation. This is repeated for the 88 groups of 5 stocks, so in total there are over 35,000 observations for global minimum variance portfolios that are used for evaluating magnitude correlation forecasts. The global minimum variance portfolio is solved for each time period using past data, rolling the estimation window forward when progressing through the dataset. The portfolio optimization problem utilizes a weight constraint $-1 \leq w_i \leq 1$ for each asset i in order to avoid irrationally large weights.

For each group of stocks, we average out the realized volatilities which are computed as in Equation 10. Finally, these average values are averaged out for the 88 different volatility groups. The tested models for the magnitude correlation estimation are the MA and EWMA models with three different window lengths and values for γ, as listed in Table 1. In addition, since it is known that shrinkage estimation methods can provide improved portfolio performance (see, e.g. Ledoit and Wolf (2003)), two simple shrinkage methods are utilized for each of the MA and EWMA estimation models: one where 50% weight is given to the actual magnitude correlation estimate, and 50% weight to the mean of these estimates; essentially, this is a shrinkage estimation toward the mean. The extreme case of using purely the average value as an estimate for each pairwise correlation is also examined (100% shrinkage).

Table 1: Evaluation of out-of-sample magnitude correlation estimation

Estimation model	Shrinkage 0%	Shrinkage 50%	Shrinkage 100%
MA50	1.997 %	2.003 %	2.015 %
MA100	1.993 %	2.002 %	2.014 %
MA200	1.992 %	2.002 %	2.014 %
EWMA0.06	2.000 %	2.001 %	2.012 %
EWMA0.04	1.995 %	2.000 %	2.012 %
EWMA0.02	1.991 %	2.000 %	2.013 %

These figures represent the average realized portfolio standard deviation according to Equation 10. The average is first computed across time for each asset group, after which the average across the 88 asset groups is calculated.

The results are reported in Table 1. It is evident that longer estimation windows for the MA models produce better portfolio performance, and hence can be considered to be providing better forecasts of magnitude correlation. For the EWMA models, a lower value of gamma is better (implying that less weight is given on more recent observations). With gamma equal to 0.02, the EWMA model produces, on average, a realized standard deviation of 1.991%, which is slightly better than that produced by the MA200 model. The differences are small in absolute terms, however, in relative terms they can be considered meaningful as they are averaged out over a large pool of observations, and the only parameter affecting the results is magnitude correlation (by the design of the simulation).

From Table 1, it is evident that the shrinkage models produced, on average, the worst performance of all the evaluated models. This is perhaps slightly surprising, as for traditional (unconditional) portfolio optimization, shrinkage models can provide improved performance. It would be intriguing to try out more sophisticated models for the estimation of magnitude correlation and compare the results to the ones obtained above. For example, the dynamic conditional correlation (DCC) framework of Engle (2002) is one potential candidate. On the other hand, the innovations in absolute return generating process can be assumed to be strongly non-normal, possibly requiring modifications to the standard DCC model. This is something that could be considered in future research.

5 CONCLUSION

This paper examines a largely neglected parameter in financial literature – the correlation between absolute values, or magnitudes, of asset returns. Previous research has shown that this correlation plays an explicit part in forming optimal portfolios when return directions are predictable. Moreover, magnitude correlation plays an indirect part in portfolio returns through joint distributions of asset returns even without return predictability. If the joint distribution of returns is not normal, magnitude correlation affects the shape of the portfolio return distribution explicitly.

The paper established a link relating magnitude correlation to standard Pearson correlation in the bivariate normal case with zero means. When the joint distribution form is not restricted, it appears that nearly all types of combinations of magnitude correlation and standard (Pearson) correlation are possible when the marginal distributions are normal. Future research in this area could focus on determining specific copula functions for expressing the link between magnitude correlation and standard return correlation in a more rigorous fashion. In addition, magnitude correlation can be thought of arising from multivariate ARCH/GARCH processes and this could be explored in detail in future research.

Empirical properties and observed market levels of magnitude correlation are intriguing and vary widely between subsamples and return intervals. For daily S& P500 stock returns, pairwise sample magnitude correlations have been relatively high and positive during the 21st century, with clear differences in the sample before the recent financial crisis (2007) and during/after the crisis. For weekly returns, the values of magnitude correlation are in general lower than for daily returns, being closer to zero with a few negative values observed. For monthly returns, the distribution of pairwise magnitude correlations is wider, with its mean close to zero and the distribution containing negative values.

Additionally, the empirical values of magnitude correlation deviate widely from the theoretical values implied by the observed Pearson correlation under bivariate normality. In essence, asset returns do not appear to follow a multivariate normal distribution, and this presents more opportunities from a portfolio management perspective. High and abnormal levels of magnitude correlation, i.e. the difference between the observed and theoretically implied values, can be interpreted as market anxiety, and the empirical levels appear to be loosely linked to the VIX index.

Finally, the accuracy of out-of-sample forecasts of magnitude correlations was evaluated by employing an indirect method measuring realized optimal portfolio variances. The results indicate that longer estimation windows are better for estimating future levels of magnitude correlation. Simple shrinkage models, on the other hand, did not fair well in the evaluation. Future research could evaluate more sophisticated models for estimating magnitude correlation.

REFERENCES

Ashkenazy, Y. – Havlin, S. – Ivanov, P. C. – Peng, C. K. – Schulte-Frohlinde, V. – Stanley, H. E. (2003) Magnitude and sign scaling in power-law correlated time series. *Physica A*, Vol. 323, 19–41.

Ashkenazy, Y. – Ivanov, P. C. – Havlin, S. – Peng, C. K. – Goldberger, A. L. – Stanley, H. E. (2001) Magnitude and sign correlations in heartbeat fluctuations. *Physical Review Letters*, Vol. 86, 1900–1903.

Barber, B. – Lehavy, R. – McNichols, M. – Trueman, B. (2001) Can investors profit from the prophets? security analyst recommendations and stock returns. *The Journal of Finance*, Vol. 56, 531–563.

Bekiros, S. D. (2010) Heterogeneous trading strategies with adaptive fuzzy actor–critic reinforcement learning: A behavioral approach. *Journal of Economic Dynamics and Control*, Vol. 34, 1153–1170.

Bollerslev, T. (1986) Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, Vol. 31, 307–327.

Chevapatrakul, T. (2013) Return sign forecasts based on conditional risk: Evidence from the UK stock market index. *Journal of Banking and Finance*, Vol. 37, 2342–2353.

Christoffersen, P. F. – Diebold, F. X. (2006) Financial asset returns, direction-of-change forecasting, and volatility dynamics. *Management Science*, Vol. 52, 1273–1287.

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

Engle, R. (1982) Autoregressive conditional heteroscedasticity with estimates of variance of United Kingdom inflation. *Econometrica*, Vol. 50, 987–1008.

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

Ivanov, P. C. – Yuen, A. – Podobnik, B. – Lee, Y. (2004) Common scaling patterns in intertrade times of U.S. stocks. *Physical Review E*, Vol. 69.

Ledoit, O. – Wolf, M. (2003) Improved estimation of the covariance matrix of stock returns with an application to portfolio selection. *Journal of Empirical Finance*, Vol. 10, 603–621.

Leone, F. C. – Nelson, L. S. – Nottingham, B. (1961) The folded normal distribution. *Technometrics*, Vol. 3, 543–550.

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

Markowitz, H. M. (1959) *Portfolio selection: efficient diversification of investments*. New York: John Wiley & Sons, Inc.

Michaud, R. O. (1989) The Markowitz optimization enigma: Is 'optimized' optimal? *Financial Analysts Journal*, Vol. 45, 31–42.

Nelsen, R. B. (2006) *An Introduction to Copulas, 2nd ed.* Statistics, Springer.

Nyberg, H. (2011) Forecasting the direction of the us stock market with dynamic binary probit models. *International Journal of Forecasting*, Vol. 27, 561–578.

Patton, A. J. – Sheppard, K. (2009) Evaluating volatility and correlation forecasts. *Handbook of Financial Time Series*, 801–838.

Psarakis, S. – Panaretos, J. (2000) On some bivariate extensions of the folded normal and the folded t distributions. *Journal of Applied Statistical Science*, Vol. 10, 119–136.

Scott, R. C. – Horvath, P. A. (1980) On the direction of preference for moments of higher order than the variance. *The Journal of Finance*, Vol. 35, 915–919.

Skabar, A. (2013) Direction-of-change financial time series forecasting using a similarity-based classification model. *Journal of Forecasting*, Vol. 32, 409–422.

Tsay, R. S. (2010) *Analysis of Financial Time Series, 3rd ed.* New York: John Wiley & Sons, Inc.

Womack, K. L. (1996) Do brokerage analysts' recommendations have investment value? *The Journal of Finance*, Vol. 51, 137–167.

APPENDIX A

The asset returns in Figures 1, 3, and 4 are generated as follows. Let $r_1 \equiv \mathrm{sgn}(n_1) \cdot m_1$ and $r_2 \equiv \mathrm{sgn}(n_2) \cdot m_2$, where $n_1 \sim \mathrm{N}(0,1)$, $n_2 \sim \mathrm{N}(0,1)$ are independent (this implies zero sign correlation, and hence, zero Pearson correlation coefficient between the generated returns). $m_1 \equiv d_{abs}(p_N(e_1))$, $m_2 \equiv d_{abs}(p_N(e_2))$, where $p_N(\cdot)$ is a function returning the percentile in the standard normal distribution, $d_{abs}(\cdot)$ is a function returning the value in the folded normal distribution (Leone et al. (1961)) with mean zero, and e_1 and e_2 are standard normal variables whose correlation coefficient determines the correlation between the magnitudes m_1 and m_2. The correlation of the magnitudes is not directly equal to the correlation of e_1 and e_2, but closely mimics it. For the illustrative examples, the values of the function d_{abs} are simulated.

APPENDIX B

For normally distributed return r_i with zero mean, $\mathrm{Var}[|r_i|] = (1 - 2/\pi)\mathrm{Var}[r_i]$. This is because the probability distribution function of the absolute value of a normal variable with mean μ and variance σ^2 is:

$$f(x) = \frac{1}{\sigma\sqrt{2\pi}}\exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right) + \frac{1}{\sigma\sqrt{2\pi}}\exp\left(-\frac{(x+\mu)^2}{2\sigma^2}\right), \qquad (x \geq 0). \tag{13}$$

This is known as the *folded normal distribution* (Leone et al. (1961)). Subsequently, the expected value of $|r_i|$ with $\mu_i = 0$,

$$\begin{aligned} \mathrm{E}[|r_i|] &= 2\int_0^{\infty} \frac{|r_i|}{\sigma_i\sqrt{2\pi}} e^{-\frac{|r_i|^2}{2\sigma_i^2}}\, d|r_i| \\ &= \sigma_i\sqrt{2/\pi}. \end{aligned}$$

Now, $\mathrm{Var}[|r_i|] = \mathrm{E}[r_i^2] - \mathrm{E}[|r_i|]^2$, which for a normally distributed variable with mean zero is equal to $\mathrm{Var}[r_i] - (2/\pi)\mathrm{Var}[r_i] = (1 - 2/\pi)\mathrm{Var}[r_i]$.

ESSAY 3

Hämäläinen, Joonas

Predictable returns and portfolio optimization:
Directional versus whole return forecasts as inputs

Preprint

Predictable Returns and Portfolio Optimization: Directional versus Whole Return Forecasts as Inputs

Joonas Hämäläinen*

Abstract

Estimation error in expected returns can render mean-variance optimization unusable. Investors typically have views only on the directions of returns, and lack accurate estimates for whole returns. We show that when return estimates are noisy but contain directional forecasting power, notable performance increase can be achieved out-of-sample by using mere directional estimates as inputs in portfolio optimization. Even with accurate information about the magnitudes of returns available, the investor can still achieve competitive portfolio Sharpe ratios by extracting the signs of return forecasts for use in portfolio optimization.

1 INTRODUCTION

Investment portfolio optimization is traditionally built on the concept of minimizing variance and maximizing expected return, as introduced in Markowitz (1952, 1959). The separation theorem of Tobin (1958) tells that investors should hold risky assets in same proportions in the portfolio regardless of their preferences. However, for decades, academic research has tried to provide solutions to the problems that mean-variance optimization suffers from in practice, most notably parameter estimation error. This often renders practical application infeasible because estimation error in the optimization inputs can result in portfolio weights that are far from the true optimal weights (see, e.g. Michaud (1989)). In fact, a naive investment strategy that divides wealth evenly across assets can

* University of Turku, Department of Accounting and Finance, Rehtorinpellonkatu 3, 20500 Turku, Finland. Part of the research has been conducted during the author's visit to Princeton University in 2013/2014. The author is grateful for the comments of Matthijs Lof, Mika Hannula, Mika Vaihekoski, Luis Alvarez, Hannu Schadewitz, Hannu Kahra, Johan Knif, and the seminar participants at GSF Winter Workshop (Helsinki) and at University of Turku.

often perform just as well out-of-sample as optimized portfolios (DeMiguel, Garlappi and Uppal (2009)).

Especially the estimation of expected returns is known to be difficult in practice. Investment managers or analysts often do not express views on the mean returns of assets, but instead disclose information about the mere directions of future price movements. A prominent example is buy/sell recommendations, which are known to have some predictive power on future stock returns (Womack (1996), Barber, Lehavy, McNichols and Trueman (2001)). These types of *directional* forecasts appear to be exceedingly common in practice. Previous research has also acknowledged the significance of separating mean return estimation and return direction, or sign, estimation: Christoffersen and Diebold (2006) show that even if conditional mean returns are independent (i.e. unforecastable), the directions of returns can still exhibit dependence, thus making them forecastable. Directional forecasting has also accumulated a considerable amount of empirical evidence, the out-of-sample accuracy of these forecasts often being surprisingly high (see, e.g. Pesaran and Timmermann (2002), Bekiros and Georgoutsos (2007), Bekiros and Georgoutsos (2008a), Bekiros (2010), Nyberg (2011), and Chevapatrakul (2013)).

Essay 1 in this dissertation introduced a framework for portfolio optimization where directional forecasts are used as inputs (hereafter, *DF framework*). This leaves out the direct, difficult estimation of mean returns and could potentially provide enhanced portfolio performance relative to using simpler solutions to utilize directional forecasts. It is not clear what the value of mean or whole return forecasts (on a continuous scale) is compared to mere *directional* forecasts in portfolio optimization. Naturally, whole return forecasts contain more information, but if this information is noisy, it may not be of much practical value. For example, Leitch and Tanner (1991) have found that professional forecasts may not be worth using when evaluated by traditional measures such as the root-mean-squared error, while at the same time the profitability of these forecasts, connected to directional accuracy, can be notable. This would seem to imply that the value of professional forecasts may be mostly in the directional component of the estimates. Especially in portfolio optimization, where even small errors in inputs are magnified in the results, using noisy return estimates can be detrimental.

While leaving out the direct estimation of mean returns, the DF framework requires the estimation of directional forecast accuracy in the form of probabilities. In a simulated environment, the model is shown to produce a substantial performance increase over simpler alternatives (see Essay 1). However, the previous study does not take into account estimation error in the directional forecast accuracies. Since estimation error plays a large role in portfolio optimization, the estimation of directional accuracies needs to be paid more attention to. By addressing this issue, this paper aims to find out whether using mere directional forecasts in portfolio optimization is feasible in practical asset management.

More specifically, the paper aims to answer the following two main research questions: 1) When directional return forecasts are available, does it matter for out-of-sample portfolio performance which optimization framework is utilized? 2) If whole return forecasts are available, is the investor better off by extracting only the signs of these forecasts and using them as inputs in portfolio optimization, as opposed to using noisy return estimates directly as expected return inputs. To answer these questions, an out-of-sample trading simulation is conducted to examine the performance of optimized portfolios.

For the out-of-sample study, specific forecasting models are not used, but instead a more general and powerful method is employed. The return forecast models used in the study are calibrated in-sample, however, parameter estimation is done using only past data, so that the use of these models resembles an out-of-sample investment setting. This way we can be sure from the outset that the models exhibit true predictive power, but the investor will not know any of the parameters of the model. Instead, they are estimated from past data, and the performance of the constructed portfolios is then evaluated on out-of-sample data.

The results of the trading simulation show that when only directional estimates are available (or when the magnitude components of the return estimates are pure noise) the investor can achieve significantly higher Sharpe ratios and geometric means by utilizing only the directional forecasts and the estimated directional accuracies. The choice of estimation method for the forecast accuracies matters for portfolio performance, and this should be taken into account when utilizing directional forecasts in portfolio optimization. Interestingly, a constant value for the individual accuracies appears to be the all-around best solution, which simplifies the task of the investor. A simpler alternative which does not take into account the accuracy of the forecasts does not fair as well as the DF framework.

When return forecasts contain accurate information about the true magnitude of returns, the investor is still equally well or better off in many cases by extracting the signs of these estimates and using only those as inputs in optimization. Especially for maximum Sharpe ratio portfolios, the directional framework produces better performance than nearly all evaluated forecasting models that produce whole return estimates. For maximum geometric mean portfolios, accuracy in the magnitude components of return estimates can make them more profitable to use than merely extracting the signs of these forecasts.

Overall, the results indicate that using mere directional forecasts as inputs in portfolio optimization can provide a significant performance increase as opposed to using traditional alternatives, including the popular model of Black and Litterman (1992), when whole return forecasts are too noisy. Even if valuable information about future returns in the magnitude estimates is available, the directional approach can still be a competitive choice compared to the alternatives, especially in the case of maximum Sharpe ratio portfolios.

The paper proceeds as follows: Section 2 briefly reviews the framework for using directional return estimates as inputs in portfolio optimization. Section 3 goes through the construct of the empirical study, and evaluates estimation methods for directional accuracy. Section 4 presents the results for the out-of-sample trading simulation. Section 5 concludes the paper.

2 PORTFOLIO CHOICE WITH DIRECTIONAL RETURN FORECASTS

Traditional mean-variance portfolio selection requires the mean return vector and variance-covariance matrix estimates as inputs in the optimization procedure. When mean return estimates are not available, but directional forecasts are provided, Essay 1 establishes a theoretical framework for mean-variance optimization. The framework is based on an idea that the investor's return on an asset i can be decomposed into the forecasted direction, an outcome component for the forecast and a magnitude component: $r_i = s_i D_i |r_i|$, where $s_i \in \{-1, 1\}$ denotes the predicted sign, $D_i \in \{-1, 1\}$ is a random variable denoting the outcome of the directional forecast (+1 if it is correct, -1 if it is incorrect), and $|r_i|$ is the absolute value of the asset's unconditional return. Under a few simplifying assumptions, it is shown in Essay 1 that the conditional[1] expected return vector

$$\mathrm{E}[\boldsymbol{r}|\Omega] = \boldsymbol{s} \odot (2\boldsymbol{\zeta} - \mathbf{1})\,\mathbf{M}, \tag{1}$$

where $\boldsymbol{s}$ denotes a vector of predicted return signs, $\boldsymbol{\zeta}$ denotes a vector of probabilities with which the directional predictions are correct for each asset, and $\mathbf{M}$ is a diagonal matrix containing the mean magnitudes of the asset returns, i.e. $\mathbf{M} \equiv \mathrm{diag}(\boldsymbol{\mu}_{abs})$. The operator $\odot$ denotes the element-wise (Hadamard) product.

The conditional variance-covariance matrix in this situation,

$$\boldsymbol{\Sigma}|\Omega = \boldsymbol{s}\boldsymbol{s}' \odot ((2\boldsymbol{Z}_{sim} - \mathbf{J}) \odot \boldsymbol{\Sigma}_{abs} + \mathbf{M}\boldsymbol{\Sigma}_{\mathbf{D}}\mathbf{M}), \tag{2}$$

where $\mathbf{J}$ denotes a matrix of ones, $\boldsymbol{Z}_{sim}$ is a matrix of joint probabilities of being simultaneously correct or wrong on a pair of assets[2], $\boldsymbol{\Sigma}_{abs}$ is the covariance matrix of the *absolute* values of unconditional returns, and $\boldsymbol{\Sigma}_{\mathbf{D}}$ is the covariance matrix of the forecast outcomes $\mathbf{D}$, which can also expressed using the probabilities in $\boldsymbol{\zeta}$ and $\boldsymbol{Z}_{sim}$.

[1] Conditional on the information set Ω, which is understood to contain the directional forecasts and probabilities governing the outcomes.

[2] The elements of this matrix, $\zeta_{ij} = \mathrm{P}(D_i = D_j)$.

The framework implies that when return directions are forecastable, the investor should pick volatile assets for higher expected returns, and that the correlation of *absolute* values of returns matters directly for portfolio variance. Moreover, if the forecast outcomes are linked to unconditional return correlations, the investor should prefer asset pairs with zero correlation; negative correlation is disliked because it presents the possibility of the investor being simultaneously correct or wrong on a pair of assets (see Essay 1 for more details).

In this framework, the inputs for portfolio optimization do not require direct estimation of mean returns. As opposed to traditional mean-variance optimization, a new parameter class that needs to be estimated is the accuracy of the forecasts in the form of probabilities ζ_i and ζ_{ij}. Since expected returns are notoriously difficult to estimate in practice, it is plausible that the reduction of estimation into prediction accuracies might produce better results in portfolio optimization. In the empirical section of this paper, it is examined whether this is actually the case in an out-of-sample trading simulation.

Essay 1 shows, in a simulated environment, that the usage of the proposed framework can make a substantial difference in portfolio performance compared to simpler alternatives when directional forecasts are available. However, the study does not include estimation error in the directional accuracy parameters and its effect on the end result could be substantial (for the effect of parameter estimation error on portfolio optimization, see, e.g. Michaud (1989)). In this paper, the issue is addressed by assuming a situation where the investor has to estimate all parameters from past data, and the performance of the portfolios is measured on out-of-sample data.

The probabilities in $\boldsymbol{\zeta}$ and $\boldsymbol{Z}_{sim}$ can be estimated, for example, by evaluating the directional accuracy of the forecasting model in the past, computing how often it produces the correct sign. However, it is not clear what the best method for achieving a good accuracy for these estimates would be – this is another area that will be focused on in the empirical section of this paper. Additionally, Essay 1 presents a hypothesis about the theoretical construct of the pairwise joint probability ζ_{ij}:

$$\zeta_{ij} \equiv \zeta_i\zeta_j + (1-\zeta_i)(1-\zeta_j) + f\Big(|\mathrm{Corr}(\mathrm{sgn}(r_i), \mathrm{sgn}(r_j))|\Big) + \eta_{ij}, \tag{3}$$

where $f(\cdot)$ denotes a monotonic function that gives the absolute value of sign correlation a proper weight. By examining empirical forecasts, this hypothesis can be tested and the empirical part of the paper examines whether it is possible to establish such a link between the sign correlation of asset returns and the joint probability ζ_{ij}.

3 EMPIRICAL STUDY

The empirical study attempts to answer the two main research questions: 1) When directional return forecasts are available, does it matter for out-of-sample portfolio performance which optimization framework is utilized? 2) If whole return forecasts are available, is the investor better off by extracting only the signs of these forecasts and using them as inputs in portfolio optimization, as opposed to using noisy return estimates directly as expected return inputs. In addition, the link between the joint probability ζ_{ij} and the sign correlation of returns (presented in Equation 3) is examined, and estimation methods for the forecast accuracies are evaluated.

3.1 Data and methodology

The study is conducted using the closing levels of weekly total return indices for the S&P100 stocks[3] from 12/31/2002 to 12/30/2014. The weekly interval is selected for a number of reasons: First, the number of observations is greater than when using monthly returns. Second, while daily returns would provide even more observations, their forecastability can be considered weak. Third, weekly returns are in general closer to the assumption of zero-mean returns, which is made in the model of Essay 1 for analytical tractability, than monthly returns. Finally, in practical asset management, daily portfolio rebalancing could be considered too costly due to transaction costs. In essence, weekly returns feature characteristics that are a compromise between the good sides and drawbacks of using monthly or daily returns. Simple net returns are calculated for each stock based on the weekly closing levels of the corresponding total return indices.

There is considerable evidence of out-of-sample directional predictability in recent academic research, such as the studies of Nyberg (2011) and Chevapatrakul (2013). While it would be ideal to use specific forecasting models, such as binary logit or dynamic probit models employed in these studies, that kind of a direct approach suffers from a significant problem: the number of forecasting models with known or permanent out-of-sample predictability is limited and would be based on previous research with no guarantee that the model works with another set of data. This makes a general level portfolio performance evaluation with forecasts infeasible. The scope of this paper is not to create or find profitable forecasting models or trading strategies, but simply to test if and how much portfolio optimization can be enhanced in case returns *are* forecastable to some extent.

[3] Only those stocks in the index with data available from the beginning of the period are included; this results in a total of 92 stocks in the dataset. Survivorship bias is not considered an issue here, as the study compares different portfolios against each other.

Instead of choosing a few specific forecasting models, a much more general and powerful method is adopted: a variety of directional forecasts are generated from autoregressive models that exhibit directional predictability *in sample*, but we utilize these as if the investor does not know beforehand how well the models will fair, i.e. the investor does not know the accuracy of the models nor the probabilities required as inputs in the DF framework. This way, we have a realistic set of forecasting models whose accuracy we do not know in advance, but we can at the same time be sure that on a general level, forecastability is there, making the study valid.

To generate the forecasts, a simple autoregressive modelling approach, described in the following section, is utilized. The models need to be calibrated before the trading simulation takes place, and is done in the full sample in advance, for each stock. In the trading simulation, these calibrated models are used to generate one period ahead forecasts, which can then be used as inputs in portfolio optimization. The generation of the forecasts is broken down into two steps: first, the direction of the forecast is generated, then a magnitude estimate is combined with the directional estimate. The DF framework will simply extract the sign of each forecast, whereas the alternatives will use the whole produced forecast as an estimate of the expected return.

The following two assumptions guide the generation of the forecasts: 1) A forecasting model cannot predict the direction of *large* returns better than *small* returns, or vice versa, i.e. the directional accuracy does not depend on the magnitude of the return. 2) The forecasted magnitude does not contain information about the directional accuracy, i.e. the magnitude value does not correlate with the probability of being correct on the direction. These assumptions can be considered fairly realistic, and considerably simplify the task at hand.

3.2 Properties of the forecasting models

A simple way to conform to the assumptions above is to deal with the directions and the magnitudes of return forecasts separately. To generate the sign component of a forecast, return signs are regressed on the one-period lagged return signs of 5 randomly selected stocks[4] in the dataset:

$$\operatorname{sgn}(r_{i,t}) = \boldsymbol{\beta}_i' \operatorname{sgn}(\boldsymbol{x}_{t-1}) + \epsilon_{i,t}, \tag{4}$$

where $\boldsymbol{x}_{t-1}$ denotes a vector of 5 predictors (lagged returns). The regression does not include an intercept to avoid introducing a baseline for the directional forecasts. By changing the predictor variables $\boldsymbol{x}_{t-1}$, i.e. choosing a different set of stocks as predictors, a different forecasting model is adopted. Repeating this

[4] Using any other economic variables as predictors would suffice as well, but since applicable information is contained in the present dataset, there is no reason to seek further.

procedure can create as many different forecasting models as there are groups of 5 stocks in the dataset. Alternatively, a logistic regression could be used, however, as the purpose is only to generate directional forecasts, the method described above is sufficient.

For the magnitude part of the return forecasts, two different models are used. The first, simpler method is to construct[5] the magnitude estimates as follows:

$$\widehat{m}_{i,t} = \alpha|r_{i,t}| + (1-\alpha)|\epsilon_{i,t}|. \tag{5}$$

where $\alpha \in \{0, 0.25, 0.50\}$ is the intensity and $\epsilon_{i,t} \sim N(0, \bar{\sigma})$. Thus, the estimate contains information about the true magnitude, but also simulated estimation error, which is generated from a normal distribution[6] with standard deviation equal to the average of all sample standard deviations, $\bar{\sigma}$. The standard deviation of the error term is set equal for all stocks in order to avoid including information about the relative magnitudes between the different stocks. When $\alpha = 0$, the magnitude estimate is pure noise. The sign and magnitude estimates are then combined to generate return forecasts for each period:

$$\widehat{r}_{i,t} = \mathrm{sgn}(\widehat{\boldsymbol{\beta}}_i' \mathrm{sgn}(\boldsymbol{x}_{t-1}))\widehat{m}_{i,t}. \tag{6}$$

The above type of magnitude estimate is very elementary in the sense that more accurate magnitude estimates are not necessarily closer to the true value of the return if the directional accuracy is low. For example, even if the magnitude estimate would contain no noise, a low accuracy for the directional component would make this model to produce forecasts with high mean squared error. For this reason, a more complex model is also considered:

$$\widehat{m}_{i,t} = \left| o_{i,t} + d_{i,t}\mathrm{sgn}(r_{i,t} - o_{i,t}) g_{i,t}|o_{i,t}| \right|, \tag{7}$$

where $o_{i,t} \equiv \mathrm{sgn}(\widehat{\boldsymbol{\beta}}_i' \mathrm{sgn}(\boldsymbol{x}_{t-1}))|\epsilon_{i,t}|$ denotes a starting point for the forecast, which is then given a nudge to the direction of $d_{i,t}\mathrm{sgn}(r_{i,t} - o_{i,t})$, where $d_{i,t} = \mathrm{sgn}(u_{i,t})$, and $u_{i,t} \sim \mathrm{U}(-\gamma, 1)$. The distance of the nudge is determined by the term $g_{i,t}|o_{i,t}|$, where $g_{i,t} \sim \mathrm{U}(0, 1)$.

What the above type of model for the magnitude estimate does is, it gives the original return estimate a nudge toward the true return $1/(1+\gamma)$ of the time. The smaller the value of γ, the more often the nudge is given toward the true value. Otherwise, the nudge is given in the opposite direction, i.e. further away from the true value. In the trading simulation, $\gamma \in \{0.8, 0.9\}$. This type of a magnitude estimation model should produce more accuracy for the whole return forecasts, however, it can violate the second assumption made at the beginning (i.e. the magnitude estimate does not contain information about the directional accuracy)

[5] It is important to notice that this model is not estimated from the data, but is used to generate artificial, noisy, magnitude forecasts based on the data.

[6] The shape of the distribution is debatable, however, historical sample means for stock returns, for example, tend to follow roughly a normal distribution.

because if the directional forecast is wrong, it is more likely that the magnitude estimate is given a nudge toward zero. In practice, some forecasting models may possess this kind of a feature, and hence it is also included in this study.

The values for the parameters α and γ are rather arbitrary, since unfortunately there are no guidelines for calibrating these kinds of magnitude models. It is hoped that experimenting with these different values will give a rough idea on how different portfolio selection models will perform in different situations. Future research could attempt to include a few specific forecasting models for mean returns. However, as mentioned before, that approach may quickly run into other problems, namely, we cannot be sure if the forecasts possess any real power, and furthermore, whether the power only exists in one particular dataset.

To examine what kind of directional accuracy the generated forecasts exhibit in sample, we compute how often, on average, they produce the correct sign. Naturally, this can be done without taking into account the magnitude estimates, and thus only sign forecasts are generated to determine the average directional accuracy. The results for 1000 different forecasting models are shown in the form of a historgram in Figure 1.

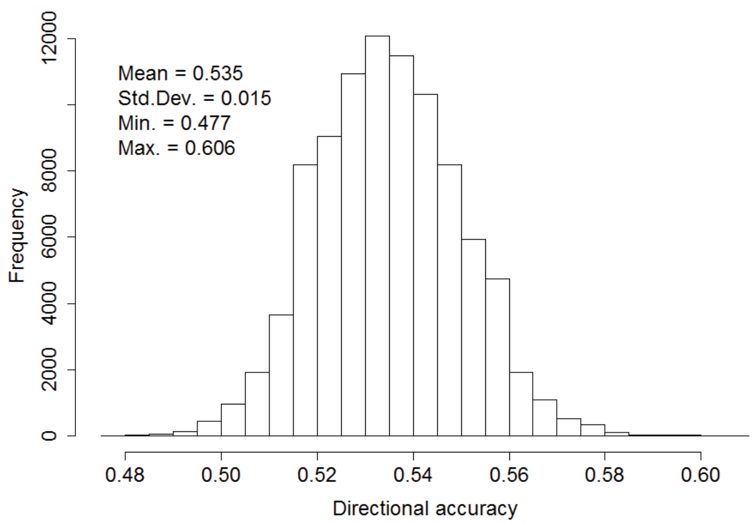

Figure 1: Histogram of directional forecast accuracies for S&P100 stocks under 1000 different forecasting models

From Figure 1, it is evident that the forecasting models for each asset produce reasonable, modest directional accuracy ranging between 47.7% and 60.6%, on average 53.5%. This indicates that for some assets, there is no real forecasting power, while for some assets there is a reasonably high accuracy of being correct. On average, the accuracy percentage is in line with out-of-sample empirical studies conducted about return sign predictability (for example, Bekiros and

Georgoutsos (2008b) forecast the direction of weekly market index returns with different models with directional accuracies ranging between 52% and 59%). In the trading simulation that follows, the investor will not know these accuracies. The calibrated models are used to generate the forecasts, but the accuracy will have to be estimated from past performance, which introduces estimation error in the trading simulation experiment.

To examine the directional accuracy of the forecasts in more detail, the sample is divided into two parts of equal length; the mean directional accuracy is approximately the same for both subsamples. However, there are notable differences between the two periods for each asset under the same model. The average absolute difference between the first subsample directional accuracy and second subsample accuracy for each asset is 0.031, indicating that for one period, the forecasting model can provide a reasonably high accuracy, which then disappears in the other period, or vice versa. This is exactly how forecasting models can often behave in the real world, exhibiting power during a stretch of time and then losing that power in the following time period, possibly regaining it again in the next stretch.

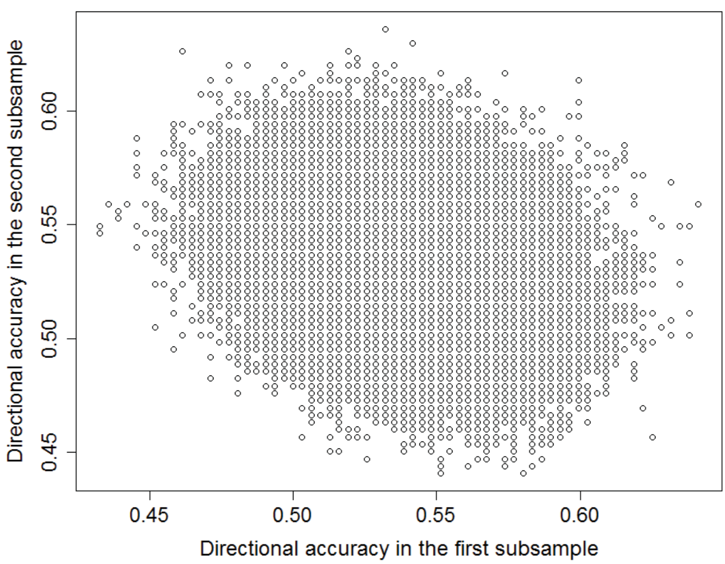

Figure 2: Directional forecast accuracy for the S&P100 stocks in the two subsamples under 1000 different forecasting models

Figure 2 shows the directional accuracy in the two subsamples. As is evident from the graph, there appears to be a tendency for the directional accuracy to decline if it had been high to begin with, and increase if it had been low in the first subsample. For example, a forecasting model that produced a 60% directional accuracy for a stock in the first subsample can have a directional accuracy of 50% in the second subsample. This kind of a tendency makes out-of-sample estimation of directional accuracy very difficult if the investor does not know about the pattern.

The accuracy of the magnitude estimates can be measured by the average correlation between the magnitude estimate and the true absolute return. For one random forecasting model, this average correlation for $\alpha = 0.25$ is 0.35, a reasonably high value. In the second case, when $\alpha = 0.50$, the mean correlation level is 0.73, which is very high. It is doubtful that this kind of accuracy can be achieved with any real forecasting model. However, for research purposes, it is interesting to include this extreme case in the trading simulation experiment as well.

It is important to notice that in portfolio optimization, mean-squared-error type of accuracy of estimates does not necessarily matter if the ranking produced for assets is accurate, resulting in the right assets being picked in the optimal portfolio. Moreover, even if the magnitude estimates are accurate or rank assets fairly well, they are of not much use if the directional forecasts are not accurate, because the investor would be facing large estimates in the wrong direction, misleading them to place a large weight in the wrong direction. The "nudge" approach for the magnitude forecasts described above should address this issue better.

3.3 Return correlation and ζ_{ij}

Equation 3 presents a hypothetical link between the pairwise joint probability of being simultaneously right or wrong on the forecasted directions, ζ_{ij}, and the sign correlation between the asset returns. Since we have now calibrated forecasting models that generate directional predictions, and we know the true return signs for each asset, it is possible to compute the joint probabilities for each asset pair. For the 1000 different forecasting models, this renders a total of 4 186 000 values, which are then averaged over the models, yielding 4 186 joint probability estimates. It is then examined how these joint probabilities are related to the actual sample-based sign correlation coefficients between the returns.

Figure 3 shows that there is a clear dependency structure between the average measured joint probabilites and the level of pairwise sign correlations: When correlation is high, the measured level of joint accuracy also tends to be high. When correlation is close to zero, the measured joint probabilities are close to 0.5, which is expected based on Equation 3, as the individual probabilities ζ_i for each asset are not large.

To further examine this relationship, simple regressions are run for the joint probability values on absolute values of sign correlation and its powers up to

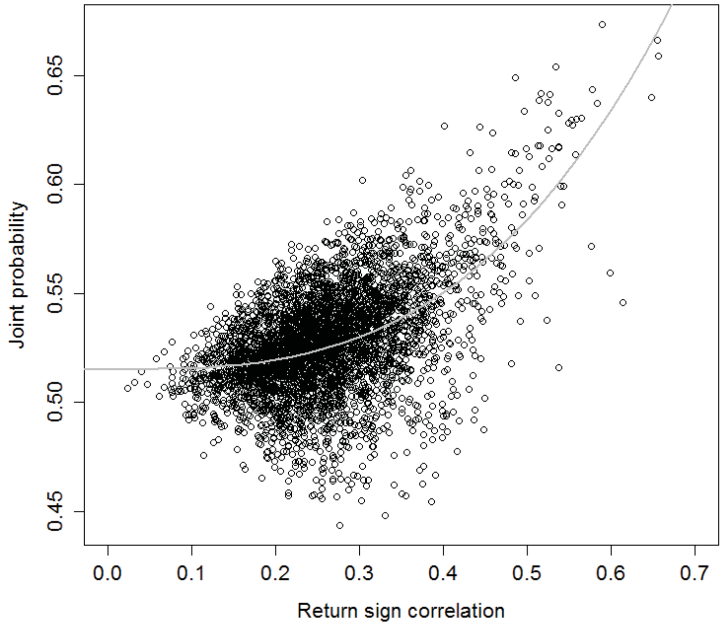

Figure 3: Average pairwise joint probabilities (ζ_{ij}) for S&P100 stocks under 1000 different forecasting models

the 4th power. The model that produces the highest R^2 is the third power (R^2 value equal to 28.1%). The t-values for the intercept (0.515) and the regression coefficient (0.551) are extremely high (>40). This fitted curve is depicted in Figure 3 in gray color.

Thus, there appears to be a strong link between the joint probabilities and asset return (sign) correlations, as was stated in the hypothesis earlier. This opens up a possibility for using sign correlations as an indirect method for estimating the joint probabilities by utilizing the correlation estimates multiplied by a constant coefficient such as the one obtained in the regression above.

3.4 Optimization criteria

The DF framework and alternative models employed in this study can be used with any portfolio optimization criterion that utilizes a mean vector and a covariance matrix. In this study, two common criteria are utilized: maximum reward-to-risk (i.e. Sharpe ratio, Sharpe (1966)) and maximum geometric mean

(MGM). In the first case, the following optimization problem is solved:

$$\begin{aligned} \max_{\boldsymbol{w}} \quad & \frac{\boldsymbol{w}'\boldsymbol{\mu}_t}{\sqrt{\boldsymbol{w}'\boldsymbol{\Sigma}_t\boldsymbol{w}}} \\ \text{s.t.} \quad & \boldsymbol{w}'\mathbf{1} = 1, \end{aligned} \tag{8}$$

where $\boldsymbol{w}$ denotes a column vector of portfolio weights, $\boldsymbol{\mu}_t$ is an estimate for the mean return vector for time period t, and $\boldsymbol{\Sigma}_t$ is the variance-covariance matrix estimate for time t. For the directional framework of Essay 1, the covariance matrix estimate is formed as in Equation 2, with the individual estimates based on past data. For alternative frameworks used in this study, the covariance matrix is the sample-based estimate based on past data.

The geometric mean maximization problem is defined as (see, e.g. Estrada (2010)):

$$\begin{aligned} \max_{\boldsymbol{w}} \quad & \left[\ln(1+\boldsymbol{w}'\boldsymbol{\mu}_t) - \frac{\boldsymbol{w}'\boldsymbol{\Sigma}_t\boldsymbol{w}}{2(1+\boldsymbol{w}'\boldsymbol{\mu}_t)^2} \right] \\ \text{s.t.} \quad & \boldsymbol{w}'\mathbf{1} = 1. \end{aligned} \tag{9}$$

This portfolio criterion is consistent with maximizing terminal wealth. To limit absurdly large weights, it is customary to include an inequality constraint in the optimization problems. For both of the above criteria, we apply the constraint $|w_i| \leq 0.2$, i.e. no asset can have a larger absolute weight than 20% of initial wealth. In the empirical study, another type of a constraint is also tested for part of the models: the absolute values of the portfolio weights summing up to one, i.e. $\sum_i |w_i| = 1$. In this case, no upper limit constraint for individual weights is set.

3.5 Estimating out-of-sample directional accuracy

The probabilities ζ_i and ζ_{ij}, which the DF framework requires as inputs, will be estimated using only previous data known to the investor before the current investment decision. The author is not aware of previous research that delves deeply into this kind of estimation, so the aim here is to proceed carefully and utilize simple models that can be easily applied in practice.

To begin with, estimates of directional accuracy from past data are likely to be noisy, especially if the forecasting model's power fluctuates in the sample, as is the case with this study. Especially the tendency for directional accuracy to revert as time progresses (see Figure 2) makes the estimation difficult. The estimation window length used in the study is 208 weeks (corresponding to 4 years). Based on a few test runs, the directional accuracy estimates vary considerably and would likely benefit from a longer estimation window. However, since data is limited and the performance of the forecasting models is subject to change at any given period, we instead opt to use simple shrinkage models, where the idea is to "shrink" the estimate toward a more structured one, hence reducing estimation error. Shrinkage estimators have been widely used in financial applications, such as in the estimation of the covariance matrix (see, e.g. Ledoit and Wolf (2003)).

The actual estimation of the probabilities ζ_i and ζ_{ij} can be unconditional, or conditional on the forecasted signs. The latter approach means that the probabilities are estimated from past data by taking into account only those observations where the forecasted sign matches the sign of the current forecast. For example, $\widehat{\zeta_i}|(s_i = 1)$ denotes an estimate of the probability when the predicted direction is up (or +1). Similarly, an estimate for $\widehat{\zeta_i}|(s_i = -1)$ can be obtained. These two estimates are likely to differ if the forecasting model does a better job at forecasting positive signs than negative signs, or vice versa.

For the joint probability ζ_{ij}, a similar approach can be used: The estimate can either be unconditional, or take the forecasted directions into account as follows: $\zeta_{ij}|(s_i = s_j)$ and $\zeta_{ij}|(s_i \neq s_j)$. This partitioning idea could be taken even further by specifying conditional probabilities $\zeta_{ij}|(s_i = 1, s_j = 1)$, $\zeta_{ij}|(s_i = 1, s_j = -1)$, and so on, not only taking into account the cases where the forecasted directions are equal, but actually separating cases where asset i is forecasted to go up and asset j forecasted to go down in value, and so forth.

In order to find out which methods are feasible to use in the estimation of directional accuracies, we conduct a study where the performance of optimal portfolios are evaluated using different estimation procedures for ζ_i and ζ_{ij}. The two methods used for ζ_i are the unconditional and the conditional case described above. The three methods for the joint probability ζ_{ij} are the unconditional, the conditional for equal/not equal signs, and the conditional method which takes into account the forecasted signs explicitly. To reduce noise, all of these estimates are average (or shrinked) toward the mean of the forecasts, i.e. $\widehat{\zeta_i^{shrink}} = (\widehat{\zeta_i} + \bar{\zeta})/2$, where $\bar{\zeta} = (1/N)\sum_i \widehat{\zeta_i}$. The same procedure is utilized for the conditional case as well, shrinking the estimates toward the mean of the unconditional estimates $\bar{\zeta}$. Finally, for simplicity, the case where all estimates are set equal to the mean estimate is included in the evaluation as well.

For ζ_{ij}, exactly the same kind of shrinkage procedure is used, i.e. the unconditional estimates are shrinked toward the mean of all unconditional ζ_{ij}'s. The second method shrinks the conditional ζ_{ij} estimates toward the mean of unconditional estimates. The third method takes into account the forecasted signs explicitly when conditioning ζ_{ij}, and once again shrinks the estimates toward the mean of the unconditional estimates. Naturally, a variety of more sophisticated methods could be developed, but this is not in the scope of this paper, and hence is left for future research.

The earlier S&P100 dataset is used in this study as well. The dataset is run through multiple times, each time using a different forecasing model, with directional forecasts generated as explained previously. In addition, for each simulation run, only 40 stocks are selected as the investable assets as this greatly reduces the time required for the simulation to run through without sacrificing too much generality. In addition, this way the available stocks are different for each simulation run to cover different scenarios. The estimation window length is 208 weeks (corresponding to four years), which leaves well over 400 weeks for out-of-sample evaluation. The procedure is repeated 50 times (for 50 different forecasting models), creating in total over 20,000 investment periods. The Sharpe ratios and geometric means are calculated for each simulation run, and finally, these numbers are averaged out over the 50 forecasting models. Table 1 presents the average Sharpe ratios and geometric means achieved by using each of the estimation methods for the directional probabilities.

Table 1: Evaluation of estimation methods for directional accuracy

Maximum Sharpe ratio portfolios

	Directional accuracy		
Joint prob.	Unconditional	Conditional	Constant
Unconditional	0.171	0.128	0.212
Conditional 1	0.195	0.166	0.238
Conditional 2	0.192	0.170	0.233
Constant	0.159	0.117	0.208

Maximum geometric mean portfolios

	Directional accuracy		
Joint prob.	Unconditional	Conditional	Constant
Unconditional	0.748 %	0.690 %	0.896 %
Conditional 1	0.761 %	0.707 %	0.898 %
Conditional 2	0.762 %	0.709 %	0.904 %
Constant	0.745 %	0.687 %	0.911 %

The upper panel shows the average Sharpe ratios achieved, and the lower panel shows the average geometric means achieved. Unconditional estimates do not condition on the forecasted directions. Conditional 1 conditions on the forecasted directions being equal/not equal, and Conditional 2 conditions on the forecasted directions being equal to signs expressed explicitly (either +1 or -1). Constant refers to the cases where the estimates are set to equal the mean of the unconditional estimates.

From Table 1, it is evident that the choice of estimation method has a notable effect on portfolio performance. For both portfolio optimization criteria, the usage of constant ζ_i estimates, i.e. setting the estimates equal to the mean of the individual (unconditional) estimates, provides the best performance. Compared to, for example, the conditional case for ζ_i, the difference in Sharpe ratios can be substantial. Using the constant values for ζ_i's appears to guarantee good performance no matter which method is used for estimating the joint probabilities ζ_{ij}.

The choice for the estimation method of ζ_{ij} affects portfolio performance as well, especially in the case of maximum Sharpe ratio portfolios. From Table

1, it can be seen that the highest Sharpe ratio, on average, is produced by using the Conditional 1 estimate, which means conditioning the estimation on the forecasted pairwise directions being equal or not equal. Conditional 2 produces nearly the same level of performance, but since the latter one is more complicated to estimate, resorting to use Conditional 1 appears to be the best choice.

For the maximum geometric mean portfolios, the same kind of performance effect can be observed as in the case of maximum Sharpe ratio portfolios, however, the choice of the estimation method for the joint probabilities does not appear to make a big difference. This can be explained by the fact that the MGM criterion places less emphasis on the variance of the portfolio than does the maximum Sharpe ratio criterion, leading the estimates for ζ_{ij} not mattering as much in portfolio optimization.

Overall, the most sensible combination that takes into account both criteria, based on Table 1, is to combine constant ζ_i estimates with the conditional ζ_{ij} estimates, shrinking the latter one toward the mean of the unconditional estimates. This combination of estimation methods will be utilized in the trading simulation that follows.

It should be emphasized that these methods of estimating directional accuracy are very elementary in nature. While it is possible to develop these ideas further, the methods presented here serve as a starting point as no previous literature seems to examine the subject on a deep level. Developing the estimation of directional accuracies further after obtaining the results from this study is a potential topic for future research.

3.6 Alternative frameworks

Equipped with return forecasts for the next period, the investor has a number of ways to utilize this information. The most simple method would be to use the forecasted returns as inputs for the mean return vector, and this serves as a natural benchmark choice. A more sophisticated and widely used method is the framework of Black and Litterman (1992), which blends active views with the market equilibrium expected returns, and is also employed in this study. The traditional approach to the Black-Litterman framework is to use the Capital Asset Pricing Model (CAPM, see e.g. Sharpe (1964)) to provide baseline views for the investor, with which the tactical allocation views are then combined. In this study, a vector of zeros is used as the baseline views, since each asset has a forecast for its return in any case, and the resulting mean vector will always include a view on each asset. In addition, the link between the CAPM beta and cross-sectional returns has been shown to be weak, if not non-existent (e.g. Fama and French (1992)), which also speaks for not using the CAPM as the baseline model.

In this case, the conditional mean vector in the Black-Litterman framework becomes:

$$\boldsymbol{\mu}_{BL} = ((\tau\boldsymbol{\Sigma})^{-1} + \mathbf{P}'\boldsymbol{\Omega}^{-1}\mathbf{P})^{-1}(\mathbf{P}'\boldsymbol{\Omega}^{-1}\mathbf{q}). \tag{10}$$

where τ is a sensitivity parameter, $\mathbf{P}$ is a matrix signaling which assets the views concern, $\boldsymbol{\Omega}$ is a matrix containing estimates for the error term variance, and $\mathbf{q}$ is a vector containing the forecasted returns (i.e. active views). The covariance matrix used in portfolio optimization will be the same as in the case of using a traditional mean vector, i.e. the sample-based covariance matrix.

In the literature, there appears to be no clear consensus on what the value of τ should be; generally, it ranges between 0 and 1. In the trading simulation, we explored the performance by setting $\tau = 1/208$, i.e. one divided by the sample size. This produced poor performance likely due to the fact that the Black-Litterman estimates were extremely close to zero in that case, leading portfolio optimization to be conducted mostly based on the variance of the portfolio. Setting $\tau = 1$ produces significantly better performance, and therefore this setting is adopted in this simulation study.

Since the forecast models provide an estimate for each stock's future return, the matrix $\mathbf{P}$ is simply the identity matrix, and the vector $\mathbf{q}$ contains the forecasted values. Thus, the Black-Litterman expected return vector is simplified into

$$\boldsymbol{\mu}_{BL} = ((\tau\boldsymbol{\Sigma})^{-1} + \boldsymbol{\Omega}^{-1})^{-1}(\boldsymbol{\Omega}^{-1}\mathbf{q}). \tag{11}$$

The $\boldsymbol{\Omega}$ matrix contains the variances of the error terms, often expressed as the confidence levels on the forecasts – these are relatively easy to estimate by computing them from historical forecasts and the realized returns. Since the expected value of the error term should be zero, the variance is simply equal to the mean squared error of the forecasts. This is a nice way of solving the estimation of the matrix $\boldsymbol{\Omega}$.

The framework of Essay 1 is much different in that it completely disregards noisy information in whole return forecasts, and extracts only the predicted signs from them. Alternatively, one does not even need whole return forecasts, but simply directional views expressed by, for example, an investment manager. In addition, the framework includes a conditional covariance matrix for optimization, as introduced in Equation 2.

To examine if the actual framework makes a difference, or just the inputs (i.e. using directional estimates versus whole return estimates), we use a simplistic model to incorporate directional forecasts into mean-variance optimization by defining the mean return vector as:

$$\boldsymbol{\mu}_{simplistic} = \boldsymbol{s}\widehat{\mu}_{abs} \tag{12}$$

where $\boldsymbol{s}$ is a vector containing the forecasted directions, and $\widehat{\mu}_{abs}$ is the average absolute return of all available assets estimated using the data in the estimation window. This simplistic model is used together with the sample-based covariance matrix.

The comparison of these four methods (traditional mean-variance optimization, the Black-Litterman framework, the DF framework, and the simplistic method for incorporating directional forecasts) should help answer the research questions posed in the introduction and offer guidance in choosing the best model in the event that returns exhibit predictability.

In the trading simulation, the same dataset of S&P100 stocks is used repeatedly, each time randomly selecting 40 stocks as investable assets and generating return forecasts for them. The forecasting model is also different for each time (based on randomly selected 5 stocks in the full sample, as described earlier), and thus the simulation process can be repeated an arbitrary number of times to achieve robust results. In the trading simulation results that follow, the performance is evaluated over 50 different forecasting models. The estimation window used to compute all parameter estimates is 208 weeks (corresponding to 4 years).

4 RESULTS OF OUT-OF-SAMPLE TRADING SIMULATION

4.1 Directional inputs in optimization

To answer the first research question (i.e. When directional return forecasts are available, does it matter how they are utilized for portfolio optimization?), the trading simulation in this section looks at the out-of-sample performance of optimal portfolios under the maximum Sharpe ratio and maximum geometric mean criterion. In addition to answering the question above, the results also shed light to the question whether the investor is better off by extracting only the signs of return forecasts when the magnitude estimates are pure noise.

The return forecasts are generated as described in Section 3, choosing $\alpha = 0$ in Equation 5. Thus, the directional forecasts are combined with a completely random value drawn from a normal distribution. This resembles a situation where the investor's model has power in forecasting the directions, but the whole return estimates are very noisy. While one cannot know for certain whether forecasting models in practice exhibit this kind of a feature, there is evidence that can be considered to support this kind of a setup: forecasting models in practice may be profitable because of their directional accuracy, whereas by traditional measures (such as the mean squared error), their performance can be poor (see, e.g. Taylor (1980), Leitch and Tanner (1991)).

Table 2 presents the results for the average performance of 40 asset portfolios, the standard deviations of the performance indicator, and the average turnover of the portfolio in order to monitor how much trading needs to take place to keep the portfolio updated. The constraint $|w_i| \leq 0.2$ is in place for the portfolio weights. It needs to be emphasized that these are out-of-sample results, i.e. parameter estimation is done with data available prior to the investment decision for each period. Thus, these results include parameter estimation error, which is known to be a crucial factor affecting portfolio performance (see, e.g. Michaud (1989)).

Table 2: Average performance of portfolios when the directions of returns are predictable, with pure noise as magnitude estimates

Maximum Sharpe ratio portfolios

	Framework			
	DF	Simplistic DF	Black-Litterman	Traditional
Avg. Sharpe ratio	**0.234**	0.201	0.191	0.183
Std. dev. of Sharpe ratios	0.037	0.046	0.044	0.042
Avg. turnover	5.27	5.51	5.75	5.71

Maximum geometric mean portfolios

	Framework			
	DF	Simplistic DF	Black-Litterman	Traditional
Avg. Geometric mean	**0.88 %**	0.72 %	0.78 %	0.78 %
Std. dev. of Geometric means	0.21 %	0.19 %	0.27 %	0.24 %
Avg. turnover	7.47	6.50	7.91	7.91

40-asset portfolios with 50 different forecasting models for the sample. The DF column reports the average performance of portfolios formed according to the directional framework. The simplistic DF column has mean return estimated according to Equation 12 and uses the sample-based covariance matrix estimate. The traditional column includes portfolios where whole return forecasts are used as direct inputs for the expected returns. The Black-Litterman column includes portfolios with mean estimates according to the model of Black and Litterman (1992).

The average Sharpe ratios for optimal portfolios in Table 2 show that the best performance is achieved using the DF framework, utilizing only the signs of the return forecasts, which yields an average Sharpe ratio of 0.234. The second best alternative is the simplistic framework, which combines the directional inputs with a constant value for the magnitudes, producing an average Sharpe ratio of 0.201. This is notably lower than the value produced by the DF framework, implying that it pays off to utilize the more sophisticated framework. However, it should be kept in mind that the estimation method for ζ_i and ζ_{ij} (discussed in Section 3) needs to be paid close attention to, or otherwise the performance of the DF framework would likely be considerably worse.

From Table 2, it can be observed that the Black-Litterman model, using the whole return forecasts to form the conditional expected return vector, produced on average a Sharpe ratio of 0.191, which is lower than that produced by either the DF or the simplistic directional framework. The traditional model produced the worst result, an average Sharpe ratio of 0.183, which demonstrates that combining directional estimates with noise as magnitude estimates to utilize traditional mean-variance optimization is not a useful strategy.

In addition to providing superior performance in case of the maximum Sharpe ratio portfolios, the DF framework also produces a lower variability of performance compared to the simplistic directional framework, as the standard deviation of Sharpe ratios is 0.037 compared to 0.046 for the simplistic framework. Moreover, the average turnover for the DF framework was slightly lower (5.27) compared to that of the simplistic framework (5.51), or either of the models using whole return forecasts as inputs (5.75 and 5.71).

The average geometric means for the optimal portfolios in Table 2 indicate that when directional forecasts are available, the investor is better off by utilizing the DF framework also in the case of the MGM criterion, as it produces an average geometric mean of 0.88%. The second best result, 0.78%, is produced by both the Black-Litterman and the traditional framework. The difference between these two values is notable, as the terminal wealth over a long period of time (e.g. 10 years) provided by the DF framework can be nearly twice that produced by the alternatives. The worst alternative in the case of MGM portfolios is the simplistic framework with an average geometric mean of 0.72%, however, it produces the least variable performance and also has the lowest turnover amount of the four evaluated models.

Overall, from the results presented in Table 2, we can infer that if the magnitudes for return estimates are pure noise (or close to being pure noise) as can be the case in practice, it can be advisable to extract the signs of these estimates and use them as inputs, either utilizing the simplistic framework, or better yet, the DF framework. However, it should be kept in mind that in order to achieve the performance illustrated in Table 2, the directional accuracy estimation needs to be paid close attention to, as described earlier.

4.2 More accurate return forecasts

To answer the second question presented in the introduction (i.e. If accurate whole return forecasts are available, is the investor still better off by extracting the signs of these forecasts and using them as inputs in portfolio optimization?), the trading simulation looks at the out-of-sample performance of portfolios when return forecasts contain accurate information about the magnitudes of the future returns, as described in Section 3.

The four alternatives for the magnitude forecasts are the 25:75 and 50:50

ratios of true information vs. noise (setting $\alpha = 0.25$ or $\alpha = 0.5$ in Equation 5) and the nudge forecasts (according to equation 7, where the starting point for the magnitude estimate is pure noise) with γ-parameters of 0.9 and 0.8. This kind of accuracy for magnitude estimates may not be possible to achieve in practice considering how difficult it is to forecast mean returns, but they will serve the purpose as it is useful to include slightly extreme scenarios. In essence, if the investor is not sure about the accuracy of the return forecasts, he/she could resort to following the guidance provided in the previous section when the magnitude components of the return estimates are pure noise. If, on the other hand, the investor is confident that there is true information in the magnitudes, then the results in this section can act as a guideline.

Table 3: Average performance of optimal portfolios when the directions of returns are predictable, with true information contained in magnitude estimates

Maximum Sharpe ratio portfolios

		Framework			
	DF	Trad25	Trad50	Tradn08	Tradn09
Avg. Sharpe ratio	*0.234*	0.188	0.185	0.217	0.190
Std. dev. of Sharpe ratios	*0.036*	0.041	0.040	0.043	0.041
Avg. turnover	*5.19*	5.78	5.81	5.41	5.43
		BL25	BL50	BLn08	BLn09
Avg. Sharpe ratio		0.184	0.178	**0.234**	0.205
Std. dev. of Sharpe ratios		0.040	0.040	0.037	0.039
Avg. turnover		5.82	5.84	5.40	5.41

Maximum geometric mean portfolios

		Framework			
	DF	Trad25	Trad50	Tradn08	Tradn09
Avg. Geometric mean	*0.85 %*	0.88 %	0.94 %	**0.97 %**	0.89 %
Std. dev. of Geometric means	*0.23 %*	0.27 %	0.28 %	0.28 %	0.23 %
Avg. turnover	*7.51*	7.90	7.90	7.90	7.90
		BL25	BL50	BLn08	BLn09
Avg. Geometric mean		0.82 %	0.85 %	0.96 %	0.83 %
Std. dev. of Geometric means		0.28 %	0.28 %	0.28 %	0.27 %
Avg. turnover		7.88	7.89	7.87	7.88

40-asset portfolios with 50 different forecasting models for the sample. The DF column reports the average performance of portfolios formed according to the directional framework, merely extracting the signs of the return forecasts. The trad. columns include portfolios where whole return forecasts are used as direct inputs for the expected returns, and the suffix refers to the information content in the magnitude estimates. The BL columns include portfolios with mean estimates according to the model of Black and Litterman (1992), with the suffix denoting how the magnitude estimates are formed.

Table 3 shows the average performance of 40-asset portfolios when magnitude estimates are accurate to the extent described earlier. The DF framework simply extracts the signs of the return forecasts and performs optimization as before, whereas the alternative frameworks utilize the whole forecasts as inputs. For the case of maximum Sharpe ratio portfolios, even with accurate information about future magnitudes, the DF framework still produces a very competitive average Sharpe ratio of 0.234, which is matched by the Black-Litterman model using the nudge magnitude estimate with $\gamma = 0.8$. Moreover, the standard deviation of the Sharpe ratios produced by the DF framework is reasonably low (0.036) compared to the alternatives. In addition, the average turnover of the DF framework is lower than that produced by any of the other models, indicating that less trading is needed in order to keep the portfolio updated.

In the case of maximum geometric mean portfolios, the more accurate magnitude estimates help producing, on average, higher geometric means than simply extracting the signs and using the DF framework. The latter produced an average geometric mean of 0.85%, whereas the best alternative is the traditional framework when magnitude estimates are generated with the nudge method with $\gamma = 0.8$, producing an average geometric mean of 0.97%. The difference is notable, however, it should be noted that this kind of accuracy in magnitude estimates may not be possible to achieve in practice. If the accuracy of the information is lower, there may not be much difference in the performance between the DF framework and the alternatives.

The figures in the above tables are not to be taken too literally, as this is a simulation study with certain assumptions in place and randomness involved. The results vary depending on which 40 assets are sampled as investable assets in each simulation run. The most important take-away is the relative performance between the different models: For Sharpe ratios, the directional framework appears to work well against all the alternative models tested, even with accurate whole return forecasts. For the maximum geometric mean criterion, the DF framework can produce competitive performance if the whole return forecasts are not very accurate.

In the study above, the value for noisy magnitude estimates is drawn from a normal distribution with a standard deviation equal to the average standard deviation in the sample of weekly returns. This may produce rather large values for the whole return forecasts (this is of course debatable, as some forecasting models may produce very fluctuating values, whereas others constantly produce a value close to zero). For this reason, we also experimented with smaller, scaled values for the whole return forecasts, but this did not appear to have a meaningful effect on the results.

Finally, one more test is run to see if the portfolio weight constraint plays a part in the results. In the above results, the weight constraint $|w_i| \leq 0.2$ was in place. Table 4 reports the average Sharpe ratios and geometric means when there is no constraint set for the individual asset weights, but instead the sum of absolute weights is set equal to one, i.e. $\sum_i |w_i| = 1$.

Table 4 compares the performance of the overall best model, the Black-Litterman method with the nudge approach with $\gamma = 0.8$, and that of the DF framework. Under this portfolio weight constraint, there is no clear difference

Table 4: Average performance of optimal portfolios with the constraint $\sum_i |w_i| = 1$

Maximum Sharpe ratio portfolios

	Framework	
	DF	BLn08
Avg. Sharpe ratio	**0.242**	0.238
Std. dev. of Sharpe ratios	0.051	0.046
Avg. turnover	1.22	1.51

Maximum geometric mean portfolios

	Framework	
	DF	BLn08
Avg. Geometric mean	0.25 %	**0.35 %**
Std. dev. of Geometric means	0.12 %	0.27 %
Avg. turnover	1.17	1.95

The DF column presents the average performance of portfolios optimized using the directional framework. The BLn08 column includes portfolios where the model of Black and Litterman (1992) is used to for estimates of expected returns, and the magnitude estimates are generated by the nudge method with gamma equal to 0.8.

between the performance of the two models under the Sharpe ratio criterion. However, for the maximum geometric mean portfolios, the Black-Litterman model utilizing the nudge forecasts still produces notably better performance (0.35% versus 0.25%). It should be noted that the standard deviation of this performance is considerably higher for the Black-Litterman model, or 0.27% compared to that of the DF framework, 0.12%. Thus, utilizing this looser constraint causes the performance of the model using the whole return forecasts to be very variable. In addition, the average turnover is considerably lower for the DF framework.

Overall, based on the results presented in this section, the investor would be approximately equally well or better off when forming maximum Sharpe ratio portfolios by utilizing the framework of Essay 1 instead of using whole return forecasts as inputs in optimization even when they contain some amount of true information about the magnitudes. However, in order to reach this level of performance, the estimation of directional accuracies in the DF framework needs to be paid close attention to. If magnitude estimates contain enough useful information, using the whole return forecasts can provide better performance in the case of maximum geometric mean portfolios. However, it should be noted that the level of accuracy in the magnitude estimates in practice can be questioned.

5 CONCLUSION

Forecasting future asset returns is a notoriously difficult task and using noisy return estimates in portfolio optimization often magnifies the problem. This paper approaches the issue by asking whether it is more beneficial to use only part of the information in these estimates, namely the signs, to achieve improved portfolio performance out-of-sample.

Instead of choosing specific forecasting models, the empirical study is conducted by generating directional forecasts that mimick the performance of comparable forecasting models in practice. For example, the generated models present a level of directional accuracy that varies through time, which can be expected to be the case in practice as well. The directional estimates are combined with magnitude estimates, which can be either pure noise, or contain true information about the returns. This type of an approach allows generating a vast number of different forecasting models and evaluating the performance of portfolios in a robust setting by providing a large number of observations. In addition, the approach is not tied to a specific forecasting model, but the results can apply on a general level to all models that share these qualities.

In the out-of-sample trading simulation conducted, it was found that when the magnitude components of return forecasts are purely noise, the investor can achieve a notable performance increase by using only the directions, or signs, of the return forecasts instead of using the whole return forecasts as estimates for expected returns. The directional inputs alone are not sufficient if they are not utilized correctly: the framework of Essay 1 provides a solution to this. The estimation of directional accuracies in this framework needs to be undertaken carefully: this paper also evaluated methods for estimating directional accuracy, and it appears that using constant or shrinkage estimates can produce good results.

When accurate information about the magnitudes is included in the return forecasts, the directional framework using the mere signs as inputs still produces high Sharpe ratios, on average, compared to the alternative models employed in the study. For maximum geometric mean portfolios, the alternatives utilizing the whole return forecasts fair better, on average, than extracting the mere sign of these forecasts and using them as inputs in optimization.

Overall, extracting the signs of return forecasts and using them as inputs instead of whole return forecasts appears to produce good performance in a wide variety of settings. However, to achieve this performance, the estimation of directional accuracy needs to be paid close attention to. It is possible to develop the idea further by incorporating more sophisticated estimates, which could possibly lead to better performance.

The power of this study, namely that return forecasts are generated artifically and they can be applied in a wide variety of settings, can also be seen as a limitation: there is no way of knowing for certain how well the generated forecasts represent reality. The directional accuracy generated in the study is in line with some recent empirical out-of-sample performance. The accuracy of the magnitude estimates is debatable, however, it can be assumed that greater accuracy than that provided in this study would be hard to obtain in practice, and therefore, the results can be viewed as supporting the directional approach.

Future research could sacrifice some generality, and instead focus on a few specific forecasting model and aim to find out whether the results in this paper are supported by using forecasting models known to exhibit power in practice. Moreover, the estimation of directional accuracy could be explored further by including more sophisticated models and different lengths of estimation window. If a general level estimation method that produces good performance in most situations can be found, it would be helpful for applying the directional framework in practice.

REFERENCES

Barber, B. – Lehavy, R. – McNichols, M. – Trueman, B. (2001) Can investors profit from the prophets? security analyst recommendations and stock returns. *The Journal of Finance*, Vol. 56, 531–563.

Bekiros, S. D. (2010) Heterogeneous trading strategies with adaptive fuzzy actor–critic reinforcement learning: A behavioral approach. *Journal of Economic Dynamics and Control*, Vol. 34, 1153–1170.

Bekiros, S. D. – Georgoutsos, D. A. (2007) Evaluating direction-of-change forecasting: Neurofuzzy models vs. neural networks. *Mathematical and Computer Modelling*, Vol. 46, 38–46.

Bekiros, S. D. – Georgoutsos, D. A. (2008a) Direction-of-change forecasting using a volatility-based recurrent neural network. *Journal of Forecasting*, Vol. 27, 407–417.

Bekiros, S. D. – Georgoutsos, D. A. (2008b) Non-linear dynamics in financial asset returns: the predictive power of the cboe volatility index. *The European Journal of Finance*, Vol. 14, 397–408.

Black, F. – Litterman, R. (1992) Global portfolio optimization. *Financial Analysts Journal*, Vol. 48, 28–43.

Chevapatrakul, T. (2013) Return sign forecasts based on conditional risk: Evidence from the UK stock market index. *Journal of Banking and Finance*, Vol. 37, 2342–2353.

Christoffersen, P. F. – Diebold, F. X. (2006) Financial asset returns, direction-of-change forecasting, and volatility dynamics. *Management Science*, Vol. 52, 1273–1287.

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

Estrada, J. (2010) Geometric mean maximization: An overlooked portfolio approach? *The Journal of Investing*, Vol. 19, 134–147.

Fama, E. F. – French, K. R. (1992) The cross section of expected stock returns. *Journal of Finance*, Vol. 47, 427–465.

Ledoit, O. – Wolf, M. (2003) Improved estimation of the covariance matrix of stock returns with an application to portfolio selection. *Journal of Empirical Finance*, Vol. 10, 603–621.

Leitch, G. – Tanner, J. E. (1991) Economic forecast evaluation: Profits versus the conventional error measures. *The American Economic Review*, Vol. 81, 580–590.

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

Markowitz, H. M. (1959) *Portfolio selection: efficient diversification of investments*. New York: John Wiley & Sons, Inc.

Michaud, R. O. (1989) The markowitz optimization enigma: Is 'optimized' optimal? *Financial Analysts Journal*, Vol. 45, 31–42.

Nyberg, H. (2011) Forecasting the direction of the us stock market with dynamic binary probit models. *International Journal of Forecasting*, Vol. 27, 561–578.

Pesaran, M. H. – Timmermann, A. (2002) Market timing and return prediction under model instability. *Journal of Empirical Finance*, Vol. 9, 495–510.

Sharpe, W. F. (1964) Capital asset prices: A theory of market equilibrium under conditions of risk. *Journal of Finance*, Vol. 19, 425–442.

Sharpe, W. F. (1966) Mutual fund performance. *Journal of Business*, Vol. 39, 119–138.

Taylor, S. J. (1980) Conjectured models for trends in financial prices, tests and forecasts. *Journal of the Royal Statistical Society. Series A (General)*, Vol. 143, 338–362.

Tobin, J. (1958) Liquidity preference as behavior towards risk. *The Review of Economic Studies*, Vol. 25, 65–86.

Womack, K. L. (1996) Do brokerage analysts' recommendations have investment value? *The Journal of Finance*, Vol. 51, 137–167.

THE FOLLOWING PUBLICATIONS HAVE BEEN RELEASED SINCE 2014 IN TURKU SCHOOL OF ECONOMICS PUBLICATION SERIES A

A-1:2014	Kirsi-Mari Kallio ”Ketä kiinnostaa tuottaa tutkintoja ja julkaisuja liukuhihnaperiaatteella…?” – Suoritusmittauksen vaikutukset tulosohjattujen yliopistojen tutkimus- ja opetushenkilökunnan työhön
A-2:2014	Marika Parvinen Taiteen ja liiketoiminnan välinen jännite ja sen vaikutus organisaation ohjaukseen – Case-tutkimus taiteellisen organisaation kokonaisohjauksesta
A-3:2014	Terhi Tevameri Matriisirakenteen omaksuminen sairaalaorganisaatioissa – Rakenteeseen päätyminen, organisaatiosuunnittelu ja toimintalogiikan hyväksyminen
A-4:2014	Tomi Solakivi The connection between supply chain practices and firm performance – Evidence from multiple surveys and financial reporting data
A-5:2014	Salla-Tuulia Siivonen “Holding all the cards” The associations between management accounting, strategy and strategic change
A-6:2014	Sirpa Hänti Markkinointi arvon muodostamisen prosessina ja sen yhteys yrittäjyyden mahdollisuusprosessiin – Tapaustutkimus kuuden yrityksen alkutaipaleelta
A-7:2014	Kimmo Laakso Management of major accidents – Communication challenges and solutions in the preparedness and response phases for both authorities and companies
A-8:2014	Piia Haavisto Discussion forums – From idea creation to incremental innovations. Focus on heart-rate monitors
A-9:2014	Sini Jokiniemi "Once again I gained so much" – Understanding the value of business-to-business sales interactions from an individual viewpoint
A-10:2014	Xiaoyu Xu Understanding online game players’ post-adoption behavior: an investigation of social network games in

A-11:2014 Helena Rusanen
Resource access and creation in networks for service innovation
A-12:2014 Joni Salminen
Startup dilemmas
– Strategic problems of early-stage platforms on the Internet
A-13:2014 Juulia Räikkönen
Enabling experiences – The role of tour operators and tour leaders in creating and managing package tourism experiences
A-14:2014 Natalie S. Mikhaylov
New school ties: Social capital and cultural knowledge creation in multicultural learning environments

A-1:2015 Hanne-Mari Hälinen
Understanding the concept of logistics cost in manufacturing
A-2:2015 Arto Ryömä
Mielelliset ja keholliset johtajuusprosessit yksilöllisyyttä ja yhteisöllisyyttä kietomassa
– Empiirinen tarkastelu jääkiekkojoukkueen kontekstissa
A-3:2015 Kati Suomi
Managing brand identity and reputation
– A case study from Finnish higher education
A-4:2015 Jenni Jaakkola
Essays on the decision-making in representative democracy
A-5:2015 Frederick Ameyaw Ahen
Strategic corporate responsibility orientation for sustainable global health governance: Pharmaceutical value co-protection in transitioning economies
A-6:2015 Janne Ohtonen
Business process management capabilities
A-7:2015 Johanna Raitis
Identity and multinational corporations
– Building unity within diversity
A-8:2015 Hannu Vahtera
Impact of disruptions on collaborative business models and interorganizational networks in supply management
– Longitudinal contextual analysis of the business networks of two focal companies
A-9:2015 Päivikki Kuoppakangas
Decision-making and choice in the adoption of a municipal enterprise form in public healthcare organisations
– Reasoning, goals, legitimacy and core dilemmas

A-10:2015 Joonas Hämäläinen
Portfolio selection under directional return predictability

All the publications can be ordered from

KY-Dealing Oy
Rehtorinpellonkatu 3
20500 Turku, Finland
Phone +358-2-333 9422
E-mail: info@ky-dealing.fi

Turun kauppakorkeakoulu
Turku School of Economics

ISBN 978-952-249-445-0