Dynamic risk management of multi-asset portfolios

Dissertation
an der Fakultät für Mathematik, Informatik und Statistik der
Ludwig-Maximilians-Universität München

eingereicht von

Christian Groll

2017

Dynamic risk management of multi-asset portfolios

Dissertation
an der Fakultät für Mathematik, Informatik und Statistik der
Ludwig-Maximilians-Universität München

eingereicht von

Christian Groll

2017

1. Gutachter: Prof. Stefan Mittnik, PhD
2. Gutachter: Prof. Dr. Martin Missong
Tag der Einreichung: 31.08.2017
Tag der mündlichen Prüfung: 29.05. 2018

Danksagung

Im Folgenden möchte ich die Gelegenheit ergreifen, mich bei all denen zu bedanken, die meine Forschungstätigkeit und diese Dissertation über die Jahre hinweg unterstützt und somit erst ermöglicht haben.

In erster Linie gilt mein Dank meinem Doktorvater Prof. Stefan Mittnik, dem ich meine heutige Leidenschaft für die Forschung überhaupt erst zu verdanken habe, nachdem er mein Interesse für die Finanzökonometrie noch zu Studentenzeiten geweckt hatte. Wie maßgeblich darüber hinaus auch meine weitere Forschung von seiner Betreuung geprägt wurde, sieht man insbesondere daran, dass uns (gemeinsam mit weiteren Kollegen des Lehrstuhls) mittlerweile auch außerhalb der Universität die Forschung beim digitalen Vermögensverwalter Scalable Capital verbindet. Wem derart viele Doktoranden über die Promotion hinaus zu weiteren Forschungsprojekten folgen, der hat offensichtlich vieles richtig gemacht. Insbesondere seine kollegiale und sorgsame Art gegenüber Mitarbeitern ist wohl einzigartig und lässt ihn als Doktorvater auch über die fachlichen Belange hinaus zu einem Vorbild werden.

Des Weiteren gilt mein Dank auch meinem Zweitgutachter Prof. Dr. Martin Missong sowie Prof. Dr. Christian Heumann und Prof. Dr. Helmut Küchenhoff aus meiner Prüfungskommission, die sich allesamt ohne jegliches Zögern zu dem zusätzlichen Zeitaufwand bereit erklärt haben, obgleich der vollgepackte Professorenalltag für gewöhnlich nur wenig Spielraum zulässt.

Mein Dank gilt außerdem meinen aktuellen und ehemaligen Kolleginnen und Kollegen, insbesondere denen, die mich nach gemeinsamer akademischer Zeit nun auch noch weiter bei Scalable Capital begleiten. Die gemeinsame und spannende Forschung sowie die tatsächliche Umsetzung finanzökonometrischer Modelle in Form einer automatisierten Vermögensverwaltung hat mein fachliches Verständnis erweitert und sich auf diese Weise auch deutlich positiv auf meine Dissertation ausgewirkt.

Über das berufliche Umfeld hinaus ist die nötige Unterstützung durch Familie und Freunde sicherlich ebenso unerlässlich. Deshalb möchte ich mich bei meinen Eltern bedanken, die ihrerseits mit einem Doktortitel ausgezeichnet werden sollten für die grenzenlose Unterstützung, die sie mir über all die Jahre haben zukommen lassen. Bedanken möchte ich mich auch bei meinem Bruder, der mir bei Abitur, Studium und Dissertation seit jeher vorgelebt hat, was mit meinem Genmaterial überhaupt alles möglich ist, sowie bei meinen Großeltern und der ganzen Familie, meiner Freundin und meinen Freunden, die allesamt sowohl mein “Laptop-Leben” samt Wochenendarbeit toleriert haben, als auch dafür gesorgt haben, dass mein grübelnder Kopf in regelmäßigen Abständen aus der Welt der Zahlen zurück ins echte Leben geholt wurde. Durch sie alle wurde die Promotion zu einer unvergesslich schönen Zeit für mich.

Summary

Even though the development of mean-variance portfolio selection earned its inventor Harry Markowitz the Nobel Prize, the theory is usually perceived to be flawed in practical applications. Its main deficiency is that it builds on known asset return moments, although inputs have to be estimated in real life. Straightforward application hence tends to maximize errors: assets that appear to be better due to estimation errors will be overweight in portfolios, which leads to suboptimal portfolio choices. In other words, Markowitz portfolio selection is too powerful for the quality of its inputs (Scherer 2002). While increased precision of inputs obviously would lead to better portfolio choices, in this thesis we examine another way to improve results: increasing the resilience to faulty inputs. Therefore, we propose a diversification-aware approach that allows to conduct mean-variance portfolio selection under constraints on diversification. A minimum required level of diversification serves as single parameter to control sensitivity with regards to inputs, and limit cases are plain Markowitz efficient frontier portfolios on the one end, and the equal weights portfolio on the other. The single diversification parameter keeps the amount of manual decisions to a minimum, and the convex nature of the optimization problem allows fast and robust solving.

In an empirical application we show that this diversification-aware approach produces promising allocation decisions when used as part of a dynamic risk management strategy. In a backtest period of approximately 17 years, it achieves out-of-sample risk-return profiles close to the true, in-sample efficient frontier. Even more remarkable, results hold although asset return moments are obtained from an unreflected application of exponentially weighted sample moments to assets of multiple asset classes. This estimator is predestined to generate poor portfolio optimization inputs, as it is backward-looking only and treats highly heterogeneous assets equally. Still, backtest results are promising and hence indicate that diversification-aware portfolio choices have the potential to effectively deal with faulty inputs. In addition, dynamic risk management helps to considerably reduce drawdowns as compared to static portfolio weights, and hence results are promising even beyond a perspective on mean-variance utility only.

Dynamic strategies are set up with particular focus on real life feasibility, and hence should reflect some aspects that go beyond the usual academic requirements. In order to address a broad range of clients, the strategy is built to be scalable with regards to a spectrum of risk aversion levels. Furthermore, within each risk category clients are managed individually, in order to allow customization with regards to client-specific tax situations in a subsequent layer. The dynamic strategy hence has an additional optimization step based on relative tracking errors, designed to keep both trading costs and dispersion of client performances within bounds. This particularly cost-efficient implementation provides the dynamic risk management approach with enough scope for action. Again, tracking error optimization is formulated as convex optimization problem, such that it allows fast and robust solving.

Zusammenfassung

Auch wenn die Entwicklung des Erwartungswert-Varianz-Ansatzes zur Portfoliooptimierung seinem Erfinder Harry Markowitz den Nobelpreis eingebracht hat, so gilt der Ansatz in praktischen Anwendungen im Allgemeinen trotzdem als ungeeignet. Als wesentliche Schwäche gilt, dass die benötigten Renditemomente als vollständig bekannt vorausgesetzt werden, obgleich diese in realen Anwendungen geschätzt werden müssen. Auf diese Weise werden Schätzfehler der Inputs unvermindert aufgegriffen und Anlagen, die aufgrund von Schätzfehlern vorteilhaft erscheinen, erlangen einen übermäßig hohen Anteil am Gesamtportfolio. Dies führt zu suboptimalen Anlageentscheidungen - angesichts der mangelhaften Qualität der Inputs ist Markowitz Portfoliooptimierung zu drastisch (Scherer 2002). Während Verbesserungen in der Qualität der Inputs offensichtlich zu besseren Allokationsentscheidungen führen, verfolgt die vorliegende Arbeit einen anderen Weg: Ziel ist, die Robustheit bezüglich fehlerhafter Inputs zu steigern. Dies geschieht mit Hilfe eines diversifikationsfördernden Ansatzes, der die Erwartungswert-Varianz-Optimierung um ein gefordertes Mindestmaß an Diversifikation ergänzt. Dieses Mindestmaß dient als Stellschraube, mit deren Hilfe die Sensitivität bezüglich der geschätzten Inputs festgelegt werden kann. Je nach Umfang der geforderten Diversifikation lassen sich als Grenzfälle entweder der unveränderte effiziente Rand der Erwartungswert-Varianz-Analyse erzielen oder die Portfolioallokation in ein vollständig gleichgewichtetes Portfolio überführen. Die Steuerbarkeit anhand eines einzigen Parameters reduziert manuelle Eingriffe auf ein Minimum. Zusätzlich erlaubt die konvexe Natur des Optimierungsproblems eine schnelle und robuste Berechnung.

In einer empirischen Anwendung wird gezeigt, dass der diversifikationsfördernde Ansatz vielversprechende Allokationsentscheidungen trifft, wenn er als eine Komponente einer dynamischen Risikomanagement-Strategie verwendet wird. In einer Rückrechnung mit annähernd 17-jähriger Laufzeit erzielt die Strategie out-of-sample Risiko-Rendite-Charakteristika nahe dem echten in-sample effizienten Rand. Dies ist umso beachtlicher, weil die für die Strategie notwendigen Inputs als empirische Momente mit exponentiell gewichteten Beobachtungen geschätzt wurden. Ein derartiger rückwärtsgewandter Schätzer erzielt für gewöhnlich mangelhafte Ergebnisse in praktischen Anwendungen, umso mehr, da er in der vorliegenden Anwendung identisch auf heterogene Anlagen aus unterschiedlichen Anlageklassen angewendet wird. Die vielversprechenden Resultate der Rückrechnung lassen demnach darauf schließen, dass ein diversifikationsfördernder Ansatz das Potenzial hat, auch mangelhafte Inputs nutzbringend einzusetzen. Zusätzlich erreicht das dynamische Risikomanagement eine deutliche Reduzierung anhaltender Verluste (Drawdowns) gegenüber einer statischen Allokationsentscheidung. Die Resultate werden somit auch einer über die Erwartungswert-Varianz-Betrachtung hinausgehenden Nutzenfunktion gerecht.

Ein spezieller Fokus bei der empirischen Überprüfung der dynamischen Anlagestrategien liegt auf der Umsetzbarkeit in realen Situationen, sodass der Rahmen der Anwendung teils über die üblichen akademischen Anforderungen hinausgeht. Um einem möglichst breiten Kundensegment gerecht zu werden, ist die Strategie derart konstruiert, dass sie ein großes Spektrum an Risikokategorien abbilden kann. Zusätzlich sollen einzelne Kundenportfolios individuell umsetzbar sein, um eine maßgeschneiderte Lösung für Kunden mit unterschiedlichen

Steuermerkmalen zu ermöglichen. Die dynamische Strategie beinhaltet deshalb eine zusätzliche Optimierungskomponente basierend auf sogenannten relativen Tracking-Fehlern, die sowohl Transaktionskosten als auch Dispersion von Kundenperformances im Rahmen halten soll. Diese kosteneffiziente Umsetzung ermöglicht dem dynamischen Risikomanagement den nötigen Handlungsspielraum. Auch diese Optimierung ist als konvexes Optimierungsproblem formuliert und erlaubt deshalb eine schnelle und robuste Berechnung.

Contents

1 Introduction **1**
1.1 Literature overview 4
1.2 Introduction to quantitative portfolio management 8

2 Single-period portfolio selection **16**
2.1 Markowitz portfolio selection 16
2.2 Fully uninformed market view 32
2.3 Diversification-aware asset allocation 37

3 Dynamic risk management **49**
3.1 Objectives 49
3.2 Multi-step asset management approach 52
3.3 Target portfolio weight filters 54

4 Empirical application **63**
4.1 Backtest settings 63
4.2 Missing values 67
4.3 Moment estimation 69
4.4 Currency risk 72
4.5 Results 87

5 Conclusion **120**

6 Appendix **125**

References **135**

Chapter 1

Introduction

Due to historically low interest rates in the aftermath of the global financial crisis of 2008, investors constantly have to search for yield and new investment opportunities these days. But yield itself is only one perspective on an investment, the other one being the risks that come along. Higher expected yields usually require higher risks. In other words, risk is the currency that needs to be paid for return. Only when both dimensions are taken into account, return and risk, the challenges of the current financial environment fully reveal.

With years of accommodative monetary policies in almost all major economies, bond yields diminished to extremely low levels. Hence, investors increasingly have to shift to more risky asset classes in order to meet their return targets. While a bond portfolio might have been sufficient to hit a 7.5% annual return target in 1995, a more realistic present asset allocation for the same target return would almost quadruple the portfolio risk (Rieder 2016).

Against this backdrop, investments into additional asset classes other than stock and bond markets gain attraction. While most portfolios have been stock-and-bond portfolios some decades ago, asset classes like real estate, commodities or private equity increasingly gain access to portfolios. At the same time, portfolios also have become more globally diversified. While U.S. markets, together with domestic markets for non-U.S. investors, previously have made up almost 100% of portfolios, exposure to Europe, China and other emerging markets is on the rise. Such globally diversified portfolios with investments into multiple asset classes further on will be referred to as multi-asset portfolios.

The potential benefits of multi-asset portfolios are manifold. First, diversification effects inherent to multi-asset portfolios can improve existing risk-return tradeoffs. This way target returns might become achievable with less risk than with traditional stock-and-bond portfolios. And second, additional asset classes and regions can extend the spectrum of achievable risk-return characteristics. For example, high yield portfolios only can be achieved if at least some portfolio components are associated with a high return potential. In that respect, emerging markets or some even more exotic asset classes might do the trick.

These potential benefits become even more pronounced in a setting of dynamic and risk-aware portfolio management. Financial risks are usually not evenly distributed over time, but they tend to be more or less emphasized at certain points in time. Even assets of generally higher risk can experience periods of low price fluctuations, and low-risk assets can become risky at certain states of the financial environment. Provided adequate detection of time-varying risks, risks can be bypassed by moving money away from assets of temporarily elevated volatility. For example, even though the global financial crisis did impact markets all around the world, some emerging markets and some asset classes were affected considerably less than major developed markets like for example the U.S. stock market. Shifting exposure into comparatively calm markets during times of financial turmoil will generally reduce the amount of overall risk taken.

But, of course, there also is a downside to multi-asset portfolios, which is an increased complexity that has to be dealt with. Individual asset classes usually have unique characteristics, such that highly standardized models most likely become inadequate. For example, stock and bond portfolios already have vastly different characteristics. While stocks (at least in theory) could be held arbitrarily long, buy-and-hold strategies for bond portfolios are not directly feasible due to pre-defined maturities. And even until maturity, returns of bond portfolios are non-stationary due to the inherent deterministic trend that is introduced by compounding obligations. This is a fundamental difference to stock returns, which are commonly modeled as being stationary. The advent of more exotic asset classes might require dealing with even further intricacies. For example, commodity prices might be subject to seasonal trends. Bearing all of these complexities in mind, it should come as no surprise that models of asset return moments usually come with high levels of estimation errors.

The imprecision in the modeling of asset return moments also is the weak spot of Markowitz mean-variance optimization. Its main deficiency is that it builds on known asset return moments, although inputs have to be estimated in real life. Straightforward application hence tends to maximize errors: assets that appear to be better due to estimation errors will be overweight in portfolios, which leads to suboptimal portfolio choices. In other words, Markowitz portfolio selection is too powerful for the quality of its inputs (Scherer 2002). The quality of any quantitative asset allocation strategy basically depends on two components: first, inputs should be modeled as precisely as possible. Obviously, if one could perfectly predict the market, there is no need for excessively complicated strategies to exploit this capability. Simply investing in the asset with highest return would do the job. Whenever market predictions are highly imprecise, however, an additional component becomes important: being able to also recognize the level of estimation errors around optimization inputs. While increased precision of inputs obviously leads to better portfolio choices, increasing the resilience to faulty inputs also does. One such approach we will examine in this thesis.

The asset allocation strategy proposed will deal with faulty inputs by reducing the sensitivity to inputs as compared to traditional mean-variance portfolio selection. The strategy is built on a diversification-aware approach that allows mean-variance portfolio selection under constraints on diversification. A minimum required level of diversification serves as single parameter to control sensitivity to inputs. Setting the required diversification level low enough, the solution will coincide with the efficient frontier portfolios of traditional Markowitz optimization. On the other end, for the limit case of maximally high required diversification, portfolios will coincide with the equal weights portfolio, making estimated asset return moments meaningless.

In an empirical application we show that this diversification-aware approach produces promising allocation decisions when used as part of a dynamic risk management strategy. Asset return moments used as inputs to the optimization are obtained from an unreflected application of exponentially weighted sample moments to assets of multiple asset classes. This estimator is predestined to generate poor portfolio optimization inputs, as it is backward-looking only and treats highly heterogeneous assets equally. Deliberately using an estimator with large imprecision will allow to examine the robustness of the allocation strategy with regards to faulty inputs.

The diversification-aware approach will be tested together with a dynamic risk management component that tries to keep portfolio risks stable despite the time-varying nature of asset moments. As response to increasing portfolio volatility, the risk management component will trigger portfolio re-balancing into assets with lower volatility. This way, portfolio downside risks like maximum drawdowns shall be reduced. Frequent portfolio re-balancing stabilizes portfolio risks and hence prevents persistent exceedance of target risks during lasting crisis periods. This is a natural way to make portfolio returns better matching to the assumptions of mean-variance utility.

In addition, the backtested dynamic strategy in the empirical application is set up with particular focus on real life feasibility, and hence reflects some aspects that go beyond the usual academic requirements. In order to address a broad range of clients, the strategy is built to be scalable with regards to a spectrum of risk aversion levels. Furthermore, within each risk category clients should be managed individually, in order to allow customization with regards to client-specific tax situations in a subsequent layer. The dynamic strategy hence has an additional optimization step based on relative tracking errors, designed to keep both trading costs and dispersion of client performances within bounds.

The remainder of the thesis is structured as follows. First, Chapter (1) proceeds with an overview of existing literature, followed by an introduction to quantitative portfolio management that lays out some basic formulas. Chapter (2) then deals with optimal portfolio choices in a single-period setting without re-balancing. It introduces Markowitz mean-variance portfolio selection, while also pointing out some of the major deficiencies of the approach. As a potential response to estimation errors, diversification-aware portfolio selection is developed in Section (2.3). The single-period setting is extended to multi-period asset allocation strategies that adapt portfolio choices to changing market conditions in Chapter (3). In the presence of transaction costs, dynamic risk management needs to carefully balance trading requirements and transaction costs. Accordingly, Section (3.3) introduces target weight filters in order to embed dynamic risk management into a cost-efficient and practicable framework. All theoretical concepts subsequently are backtested in an empirical application with multi-asset portfolios in Chapter (4). In order to maximize diversification benefits, multi-asset portfolios consist of market indices from several regions of the world. This implicitly also introduces currency risk into the application, and Section (4.4) entails a thorough treatment of advantages and disadvantages of currency hedging. Chapter (5) summarizes all results and points to potential further research related to this thesis.

1.1 Literature overview

Quite generally, the goal of this thesis is to come up with a quantitative approach to identify optimal investments. This, of course, is not a new challenge, and people participate in financial markets searching for optimal investment opportunities since centuries. Hence, this thesis is only another component to a huge branch of already existing literature, and it heavily builds on many of the well-thought findings of other authors. Optimal investment strategies, however, can only be derived with regards to certain assumptions on how financial markets work. In some sense, the understanding of financial markets and the investment strategies that seem worthwhile based on this understanding are heavily intertwined. Hence, we will both give an overview over some of the most influential theories in finance and asset pricing, as well as over the most popular investment approaches that build on top of these theories.

One of the first theories on asset pricing was developed by (Williams 1938), who stated that the price of a stock should equal the sum of the discounted future cash-flows. With this view, Williams was one of the founders of fundamental analysis, which tries to determine the intrinsic value of any individual company. The only relevant things for the evaluation of fair stock prices, according to this theory, are company specific cash-flows and discount rates, and hence any stock could be evaluated in isolation. Inspired by this approach, Markowitz was searching for a way to take into account the uncertainty of future cash-flows more explicitly (Markowitz 1990). He recognized that focusing on expected values of future cash-flows only would leave investors with extremely concentrated portfolios. The optimal portfolio choice would be to invest into the asset with maximum expected return only. Based on this perception, he set out to find the optimal portfolio choice for an investor quite differently. In (Markowitz 1952), asset returns are modeled as random variables, and investors are risk averse with mean-variance utility functions. Hence, investors are now not only interested in expected returns, but their utility is determined by a tradeoff between expected return and asset volatility. In a first step, this tradeoff can be solved in terms of Pareto-optimality: if higher expected returns are favored, while portfolio volatility is undesired, then any portfolio is Pareto-dominated when there exists another portfolio that either has higher expected returns for the same level of risk, or lower volatility for the same level of expected return. This way, the set of all feasible portfolios can be reduced to a set of Pareto-optimal ones, the so-called efficient frontier. Later on, in (Markowitz 1959), the concept of Pareto-optimality was refined and embedded into the framework of Von Neumann-Morgenstern utility theory, in order to explicitly resolve the tradeoff between mean and variance.

Extending the universe of assets with an additional risk-free asset, the efficient frontier becomes a linear function of the risk-free rate and the portfolio with maximum Sharpe ratio, and the optimum allocation between all risky assets becomes independent of actual investor's risk preferences (Tobin 1958). Risk preferences then are only reflected in the choice of an optimal allocation of wealth between the risk-free asset and the portfolio of risky assets.

While Markowitz portfolio selection examines the optimal behavior of financial investors for any set of given asset return moments, a follow-up branch of literature deals with the consequences for financial markets that result from this investor behavior. If all investors really followed the asset allocation principles described by mean-variance portfolio selection, under what pre-conditions would a financial market equilibrium emerge, and what would it look like? These questions led to the development of the capital asset pricing model (CAPM), which was independently formulated by (Sharpe 1964), (Treynor 1961), (Lintner 1965) and (Mossin 1966). Under the assumption that all investors have the same expectations on asset moments, combined with the separation theorem of Tobin, all investors will hold the same portfolio of risky assets. If additionally markets are required to be cleared, then the CAPM shows that the portfolio held by all investors has to equal the market portfolio: a portfolio that comprises all securities with individual asset weights given by company market capitalization. Due to this result the CAPM is frequently described as the origin of passive investment strategies, where the goal is to simply track a given market as closely as possible. If the market is represented by a market capitalization weighted index, passively tracking this index equals the equilibrium outcome suggested by CAPM. Furthermore, CAPM also has important implications for asset pricing theory. While mean-variance investors are only willing to increase overall portfolio volatility when expected portfolio returns increase, too, the relation between expected asset returns and asset volatility suggested by CAPM is more intricate. Here, it is not the full asset volatility that investors get compensated for, but only the systematic and non-diversifiable part of asset volatility. Idiosyncratic, asset-specific volatility can be fully eliminated in large portfolios, and hence the market portfolio consists of systematic risk only. As the market portfolio comprises all assets, it represents the only source of systematic portfolio risk. Hence, expected asset returns should not increase proportionally to overall asset risk, but to asset exposure with regards to the market portfolio.

A similar principle also underlies the multiple risk factors model introduced in (S. A. Ross 1976). Instead of a single market factor that drives expected returns, now the exposure with regards to a set of factors determines expected asset returns. If investors require compensation for holding certain types of risks, then risk premia associated with individual assets have to be priced consistently across different assets. In other words, two assets with equal factor exposures need to be priced equally, or otherwise arbitrageurs would be able to build portfolios with positive expected returns, although they have no systematic risk. The theory of pricing with multiple risk factors hence is called arbitrage pricing theory (APT). It is derived with less restrictive assumptions than CAPM, and CAPM may be interpreted as a special case of APT with only a single risk factor. Although APT defines a theoretical framework of how assets should be priced consistently, it does not give any indication on what the factors should be. It builds a theoretical framework for pricing, but relevant factors only have been empirically identified later on.

Among the first to fill this gap and identify relevant pricing factors were (N.-F. Chen, Roll, and Ross 1986). They define a set of economic state variables that possibly could change either discount factors or expected future cash-flows, and examine their influence empirically by estimating models with multiple factors. The most relevant economic factors found to have an influence on asset pricing are industrial production, yield curve changes and measures of unanticipated inflation or changes of expected inflation. Economic growth and inflation are also found to influence asset returns in (Ang 2014), together with market volatility and other macro-economic risks. A more elaborate treatment of expected returns and macro-economic factors also can be found in (Ilmanen and Asness 2011).

In addition to macro-economic factors, expected returns also have been found to be driven by individual company characteristics like company size (Banz 1981) and value (Basu 1977). Thereby value is a measure of cost, and the idea behind it is to describe whether a company's stock price is high compared to its fundamental company value. This perspective on company fundamentals has a long history and goes back to (Graham and Dodd 1934). Both size and value, however, only did obtain their full popularity with (Fama and French 1993), who introduced them as systematic risk factors embedded into a factor model in accordance with the principles of APT. Over time, the three risk factors used in (Fama and French 1993) have been gradually extended with further factors. Most prominently, the momentum strategy introduced in (Jegadeesh and Titman 1993) was added as additional factor in the four-factor model of (Carhart 1997). Momentum describes the tendency of securities that have performed well relative to their peers in the recent six or twelve months to further outperform in the near future. In other words, past winners have historically outperformed past losers. This phenomenon has been found across countries as well as across asset classes. Further important risk factors mentioned in the literature are quality (Bae 1999), low volatility (Ang et al. 2006) and carry (Meese and Rogoff 1983). Besides the empirical evidence, for many of the factors there is no general agreement on the theoretical justifications in the literature yet, and arguments can cover both rational and behavioral aspects of investors. While risk premia have been shown to persist for style factors over long periods of time, factor premia also can be shown to vary over time (Cochrane 2011).

Equipped with a basic understanding of asset pricing models, let's now get back to the original question at hand: how can we find optimal portfolio allocations? Let's take a more critical perspective on Markowitz's mean-variance portfolio selection, which still is the gold standard in quantitative portfolio management. The approach is built on two major assumptions. First, all asset return moments are fully known. And second, true investor preferences can be fully described in terms of mean-variance properties only. In other words, either the multivariate asset return distribution can be fully described by its first two moments only, or investor utility functions are such that only the first two moments matter. Although it is highly unlikely that mean-variance properties are sufficient to fully describe investor preferences in reality, it still can be shown that mean-variance utilities form a reasonably well approximation in many situations (Levy and Markowitz 1979). With more intricate return distributions, however, two assets of equal mean returns and variances might still have substantially different risk profiles. For example, (Ang 2014) compares the performances of the S&P500 index and a short volatility strategy over time. Although both end up with similar realized returns and variances, they substantially differ in terms of drawdowns. This difference, however, is largely driven by the time-varying nature of asset return moments, which also translates into time-varying portfolio moments. And, as in this case, this can lead to substantially different strategy risk-return profiles in certain sub-periods. In a setting of regular portfolio rebalancing, which explicitly keeps time-varying moments into account, this effect implicitly gets dealt with and hence should get dramatically reduced. Hence, when the time-varying nature of asset return moments is taken into account appropriately, mean-variance utility should be a sufficiently good approximation to true investor's utility in most cases.

The more important problem of Markowitz optimization is the assumption of fully known asset moments. In reality, asset moments are not directly accessible, but they have to be estimated from data. Asset moments are a fundamental input to the portfolio selection process, and hence portfolio choices directly depend on the goodness of estimators, quite in line with the saying “garbage in, garbage out”. For the estimation one can choose from a wealth of different asset moment estimators that are known in the literature. One approach is to use non-parametric estimators, like for example sample moments, or exponentially weighted sample moments (“RiskMetrics Technical Document” 1996) that try to deal with time-varying asset moments more explicitly. Another approach would be to use parametric estimators, either in the form of unconditional return distributions, or in the form of time series models like GARCH that explicitly allow for time-varying asset moments. More common, however, are approaches that aim to reduce estimation errors by incorporating additional structure into the estimation process. One way to do this is by making use of the conclusions drawn from asset pricing theory, imposing a factor structure for individual asset returns. This way style factor exposures can be used to estimate expected asset returns. A different way to use asset pricing theory is given by (Black and Litterman 1990), who try to back out expected asset returns from observable market capitalizations. However, one does not have to rely on the results of asset pricing theory in order to impose some structure on the estimation process. A different approach is to reduce estimation errors by shrinkage of moment estimators (Ledoit and Wolf 2004).

All these approaches try to deal with the sensitivity of Markowitz mean-variance optimization to faulty inputs by hopefully improving the quality of the inputs. A different way to overcome this problem, however, is to modify the portfolio selection process itself, such that the sensitivity to inputs gets reduced. In other words, the idea is to make the selection process more robust with regards to the high uncertainty of asset return moment inputs. A first such way to increase robustness of the portfolio selection procedure was given by Markowitz himself, who carved out the solution to a portfolio selection problem subject to various kinds of constraints (Markowitz 1956). This way, for example, upper bounds on single and joint asset weights can be used to prohibit excessively concentrated portfolio solutions. Another way of reducing the scope of active asset weight decisions is given in (Roll and Anderson 1992). The idea is to basically passively track a benchmark, only allowing slight asset weight deviations in order to generate excess returns. Deviations from benchmark weights are limited by tracking error constraints. Both approaches are able to limit the influence of possibly faulty inputs by prohibiting too aggressive portfolio weight choices. However, the magnitude of the interferences caused by the additional restrictions in no way depends on the actual amount of uncertainty that surround inputs. A different approach is to take this uncertainty explicitly into account, and make the interference into the plain Markowitz portfolio selection relative to the confidence that exists for the inputs. One way of doing this is by re-sampling of portfolio weights (R. O. Michaud and Michaud 2008). Asset moments are bootstrapped in order to determine the level of uncertainty involved, and portfolios are composed such that they are appropriate for the full range of possible inputs. Built on a quite similar idea, (Ghaoui, Oks, and Oustry 2003) construct portfolios that are optimal with regards to ranges of possible inputs. Using intervals to define ranges for individual asset moments, portfolios are constructed in such a way that they are suitable for the full range of possible input constellations. Dealing with uncertainty of estimators also can be done by Bayesian approaches. For a nice overview over portfolio approaches based on Bayesian principles see

(Avramov and Zhou 2010).

Even in light of the amount of sophisticated refinements for the original mean-variance portfolio selection, there is no evidence that naive investment approaches like the equally weighted 1/N portfolio achieve inferior out-of-sample results (DeMiguel, Garlappi, and Uppal 2007). Fixed weight strategies remain to play a significant role in the investment management world. The strength of such naive approaches is that they can quite naturally benefit from diversification effects, while regular rebalancing might add additional value by introducing a tendency to buy low and sell high on average (Ang 2014). A more sophisticated approach than simple fixed weight strategies is to determine individual asset weights based on risk contributions. For example, risk parity assigns an equal portion of risk contribution to each portfolio asset, instead of simply focusing on pure asset weights only (Maillard, Roncalli, and Teïletche 2010). This way the portfolio is guaranteed to be truly diversified in terms of risk, with no single asset dominating the overall risk of the portfolio. The portfolio chosen will be between the global minimum variance portfolio and the equal weights portfolio.

Many of the existing approaches pointed out in this literature overview also will be reflected in the dynamic asset allocation strategies developed in this thesis. At the very core, strategies will build on traditional mean-variance portfolio optimization developed in (Markowitz 1952). Similar to, for example, portfolio re-sampling (R. O. Michaud and Michaud 2008), robust portfolio optimization (Ghaoui, Oks, and Oustry 2003) or some Bayesian approaches to portfolio selection, the primary focus of the strategy will lie on increased robustness against estimation errors of inputs. This way, portfolio choices shall become meaningful even in the presence of asset moment estimators that do not do a particularly good job. Increased robustness will be achieved by tilting portfolio weights into the direction of well diversified portfolios, with the equal weights portfolio as limit case. The idea is to exploit widely documented benefits of diversification and use good out-of-sample performances of naive portfolio weights as anchor point (DeMiguel, Garlappi, and Uppal 2007). In addition, single-period strategies will be embedded into multi-period strategies that try to bypass elevated market risks by actively re-balancing portfolio weights in order to hold portfolio volatility fix. Thereby a particularly cost-efficient implementation based on tracking errors (Roll and Anderson 1992) will provide enough scope of action for dynamic risk management despite the presence of transaction costs. Further references to existing literature will also pop up during a more thorough analysis of backtest results for a particular selected set of strategies. For example, overall portfolio backtest results will also be decomposed to reflect individual assets' contribution to performances and risks (T. Roncalli 2013).

1.2 Introduction to quantitative portfolio management

In this chapter we will outline some basic notation and formulas that are helpful to formally describe concepts of quantitative portfolio management. We start with a mathematical description of portfolios. A portfolio shall consist of d different assets with S_i shares of asset i. Individual asset prices at time t are denoted by $(P_{t,i})_{i=1}^{d}$. Then the portfolio price is given as sum of all individual asset positions:

$$P_{t,P} = \sum_{i=1}^{d} S_i P_{t,i}$$

Relative changes of portfolio values are measured by discrete portfolio returns:

$$r_{t,P} = \frac{P_{t,P} - P_{t-1,P}}{P_{t-1,P}}$$

It is quite common to reformulate portfolio returns with regards to portfolio weights, which are defined as:

$$w_{t,i} := \frac{S_i P_{t-1,i}}{\sum_{j=1}^{d} S_j P_{t-1,j}}$$

Weight $w_{t,i}$ equals the proportion of wealth invested in asset i for the time interval from $t-1$ to t, and one can show that $\sum_{i=1}^{d} w_{t,i} = 1$. The portfolio return now can be expressed as:

$$\begin{aligned}
r_{t,P} &= \frac{P_{t,P} - P_{t-1,P}}{P_{t-1,P}} \\
&= \frac{\sum_{i=1}^{d} S_i P_{t,i} - \sum_{i=1}^{d} S_i P_{t-1,i}}{\sum_{i=1}^{d} S_i P_{t-1,i}} \\
&= \frac{\sum_{i=1}^{d} S_i (P_{t,i} - P_{t-1,i})}{\sum_{i=1}^{d} S_i P_{t-1,i}} \\
&= \frac{\sum_{i=1}^{d} S_i P_{t-1,i} \frac{(P_{t,i} - P_{t-1,i})}{P_{t-1,i}}}{\sum_{i=1}^{d} S_i P_{t-1,i}} \\
&= \frac{\sum_{i=1}^{d} S_i P_{t-1,i} r_{t,i}}{\sum_{i=1}^{d} S_i P_{t-1,i}} \\
&= \sum_{i=1}^{d} \left[\frac{S_i P_{t-1,i}}{\sum_{j=1}^{d} S_j P_{t-1,j}} \right] r_{t,i} \\
&= \sum_{i=1}^{d} w_{t,i} r_{t,i}
\end{aligned}$$

Furthermore, portfolio moments also can be written in terms of given portfolio weights $(w_i)_{i=1}^{d}$:

$$\mathbb{E}[r_{t,P}] = \mathbb{E}\left[\sum_{i=1}^{d} w_i r_{t,i}\right] = \sum_{i=1}^{d} w_i \mathbb{E}[r_{t,i}] = \sum_{i=1}^{d} w_i \mu_i = w'\mu \tag{1.1}$$

$$
\begin{aligned}
\mathbb{V}(r_{t,P}) &= \mathbb{V}\left(\sum_{i=1}^{d} w_i r_{t,i}\right) = \mathrm{Cov}\left(\sum_{i=1}^{d} w_i r_{t,i}, \sum_{i=1}^{d} w_i r_{t,i}\right) = \sum_{i=1}^{d}\sum_{j=1}^{d} \mathrm{Cov}(w_i r_{t,i}, w_j r_{t,j}) \\
&= \sum_{i=1}^{d} w_i^2 \mathbb{V}(r_{t,i}) + \sum_{j \neq i} w_i w_j \mathrm{Cov}(r_{t,i}, r_{t,j}) \\
&= \sum_{i=1}^{d} w_i^2 \mathbb{V}(r_{t,i}) + 2\sum_{i<j} w_i w_j \mathrm{Cov}(r_{t,i}, r_{t,j}) = w' \Sigma_R w
\end{aligned} \tag{1.2}
$$

1.2.1 Multivariate versus univariate portfolio return modeling

Deriving portfolio moments from asset moments amounts to a multivariate modeling approach. Quite generally, in a multivariate modeling approach each individual asset that is part of the portfolio needs to be modeled, together with the dependence structure that prevails between individual portfolio components. Only then the portfolio return distribution is derived from portfolio weights and multivariate asset return distribution. Focusing on portfolio return expectation and variance only, concrete asset distributions are not required, and the problem at least can be reduced to the first two asset moments, as was seen in Equation (1.2) and (1.1). For more intricate properties of portfolio returns, however, one generally needs full knowledge of the portfolio return distribution, and hence needs to explicitly model the full multivariate distribution of assets. Figure (1.1) shall illustrate this multivariate modeling approach: r_t^i denotes the asset return of asset i at time t, and the lower quantile of the portfolio return distribution is taken as an example for a return distribution property of interest. Using historic observations of $(r_t^i)_{i=1}^d$, the multivariate distribution of asset returns is estimated. Only then the portfolio return distribution is derived for given portfolio weights $(w_{t,i})_{i=1}^d$.

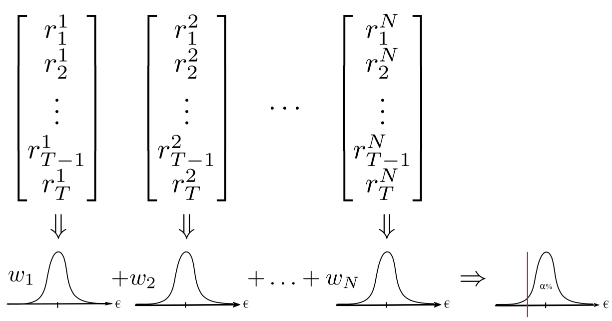

Figure 1.1: Multivariate portfolio return model

A fundamentally different approach immediately reduces the portfolio return modeling problem to a univariate model. Holding portfolio weights $(w_i)_i$ fix, historic asset returns $\left((r_t^1)_t, \ldots, (r_t^d)_t\right)$ can be used to infer historic portfolio returns $(r_t^P)_t$. The resulting univariate portfolio return series then can be used as historic observations from the currently held portfolio. Figure (1.2) illustrates this univariate modeling approach.

$$
w_1 \begin{bmatrix} r_1^1 \\ r_2^1 \\ \vdots \\ r_{T-1}^1 \\ r_T^1 \end{bmatrix} + w_2 \begin{bmatrix} r_1^2 \\ r_2^2 \\ \vdots \\ r_{T-1}^2 \\ r_T^2 \end{bmatrix} + \cdots + w_N \begin{bmatrix} r_1^N \\ r_2^N \\ \vdots \\ r_{T-1}^N \\ r_T^N \end{bmatrix} \Rightarrow \begin{bmatrix} r_1^P \\ r_2^P \\ \vdots \\ r_{T-1}^P \\ r_T^P \end{bmatrix}
$$

$$\Downarrow$$

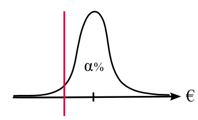

Figure 1.2: Univariate portfolio return model

Even though the current portfolio return distribution can be derived with given portfolio weights at each point in time, this return distribution will generally hold for only one period without re-balancing due to market moves. Holding a given portfolio unchanged without any further re-balancing will only fix asset volumes $(S_i)_{i=1}^d$, but not asset weights $(w_{t,i})_{i=1}^d$, as they change over time with asset price changes. Individual asset weights increase whenever the associated asset return exceeds the overall portfolio return:

$$
\begin{aligned}
w_{t+1,i} &= \frac{S_{t,i} P_{t+1,i}}{P_{t+1,P}} \\
&= \frac{S_{t,i} P_{t,i}(1 + r_{t,i})}{P_{t,P}(1 + r_{t,P})} \\
&= w_{t,i} \frac{(1 + r_{t,i})}{(1 + r_{t,P})}
\end{aligned}
$$

Hence, this poses problems when simulating portfolio returns more than a single period ahead. Even when asset moments remain constant, portfolio moments will change as portfolio weights change due to market moves. For example, let's assume a very simplistic setting where asset moments are constant over time. Let's now simulate multiple steps ahead with a multivariate modeling approach and given portfolio weights $w_t = (w_{t,i})_{i=1}^d$. Therefore, we draw asset returns $r_t = (r_{t,i})_{i=1}^d$ and get the associated portfolio return from $r_{t,P} = w_t' r_t$. Due to differences in asset returns, this will generally alter portfolio weights, such that next period portfolio returns then will be obtained by $r_{t+1,P} = w_{t+1}' r_{t+1}$, with $w_{t+1} \neq w_t$. Strictly speaking, even when asset distributions are constant over time, portfolio return distributions will change due to changes in portfolio weights.

These changes in portfolio return properties are hard to reflect in a univariate modeling approach. Again, let's start with given initial portfolio weights $w_t = (w_{t,i})_{i=1}^d$. Then with given weights and historic asset returns we can derive historic portfolio returns, and hence model the current portfolio return distribution. From this model, we can simulate one period ahead. Now, however, we would need to adjust portfolio weights in accordance with asset market moves. As individual asset returns were not explicitly simulated, however, new portfolio weights are generally not accessible. A common simplification to this problem is to just hold portfolio return properties fix, which basically would require that portfolio weights are constant, too. This could only be the case when portfolio positions are re-balanced each period, however.

1.2.2 Discrete versus logarithmic portfolio return modeling

Many financial applications build on models of logarithmic returns. For portfolio management, however, there are some drawbacks of using logarithmic returns. First, while discrete portfolio returns are a linear function of individual asset returns, this is not the case for logarithmic returns:

$$\begin{aligned} r_P^{log} &= \ln\left(w_1 \exp(r_1^{log}) + \ldots + w_d \exp(r_d^{log})\right) \\ &\neq w_1 r_1^{log} + \ldots + w_d r_d^{log} \end{aligned}$$

As a direct result, portfolio moments can not be calculated with the nice linear and quadratic functions of Equations (1.1) and (1.2), and hence are much harder to derive.

But even when logarithmic portfolio moments could be obtained, there is a second downside to using logarithmic returns. When making portfolio decisions, investors generally prefer higher returns and less risk. Higher returns are clear, as they ultimately want to increase wealth as much as possible. And less risk (uncertainty around a given payoff) is preferred because investors usually are risk-averse. Both risk and return, however, need to be measured in terms of discrete returns, as this is what investors actually will get.

Let's now assume that we have a model of logarithmic asset returns, from which we also can derive logarithmic portfolio moments for each portfolio. Then we only need to find those portfolios that do have the best risk-return properties. As we will see in Chapter (2), one way to find optimal portfolios with discrete portfolio properties is to keep a certain level of portfolio expectation fix and minimize portfolio variance. This way, one will always get an efficient portfolio. Applying the same approach on logarithmic portfolio moments will generally not work, as both expectations and variances of logarithmic returns do affect expected returns in discrete terms. In other words, holding logarithmic portfolio expectations fix, does not mean that discrete portfolio expectations are fix as well.

Let's illustrate this with an example. We assume that logarithmic portfolio returns do follow a normal distribution. Then discrete gross returns will follow a log-normal distribution:

Definition 1.2.1: ***[Log-normal distribution**, (Johnson, Kotz, and Balakrishnan 1995)]*

Let Y be a normally distributed random variable with mean μ and standard deviation σ. Then variable $X = \exp(Y)$ follows a log-normal distribution:

$$X = \exp(Y), \quad \Rightarrow X \sim \ln \mathcal{N}(\mu, \sigma^2)$$

Hence, normally distributed logarithmic returns $r^{\log}$ will generate log-normally distributed discrete gross returns:

$$(1+r) = \exp(r^{\log}), \quad \Rightarrow 1+r \sim \ln\mathcal{N}(\mu^{\log}, (\sigma^{\log})^2)$$

Furthermore, with given moments of the underlying normal distribution, moments of the log-normally distributed variable can also be derived.

Lemma 1.2.2: [Log-normal moments, (Johnson, Kotz, and Balakrishnan 1995)]

Let X be a log-normally distributed random variable, $X \sim \ln\mathcal{N}(\mu, \sigma^2)$. Then the arithmetic moments of X are given by

$$\begin{aligned} \mathbb{E}[X] &= \exp\left(\mu + \frac{\sigma^2}{2}\right) \\ \mathbb{V}(X) &= \left(\exp\left(\sigma^2\right) - 1\right)\exp\left(2\mu + \sigma^2\right) \end{aligned}$$

Using Lemma (1.2.2), moments of discrete net returns can be derived as:

$$\mu(r) = \mu(1+r) - 1 = \exp\left(\mu(r^{log}) + \frac{\sigma^2(r^{log})}{2}\right) - 1 \tag{1.3}$$

$$\sigma^2(r) = \sigma^2(1+r) = \left(\exp\left(\sigma^2(r^{log})\right) - 1\right)\exp\left(2\mu(r^{log}) + \sigma^2(r^{log})\right) \tag{1.4}$$

Hence, discrete portfolio expected returns do depend on both moments of logarithmic returns. In particular, increasing volatility of logarithmic returns will also increase expectations of discrete returns. Holding logarithmic expectations fix while minimizing logarithmic volatility, both expected returns and volatility will be decreased in discrete terms.

To conclude, the relevant quantities for an investor are given in terms of discrete returns, because this is what he really gets. Also building the model on discrete returns allows a straightforward identification of efficient portfolios, as risk and return have clear effects on investors' utility. Holding everything else fix, increased portfolio expectation is desirable, while increased volatility reduces utility. This straightforward interpretation will get lost when using logarithmic returns, as the effect of increased volatility becomes ambiguous. It simultaneously increases discrete volatility and expectation.

1.2.3 Scaling of moments

A common problem in many financial applications is a mismatch between different time horizons. Due to data availability, models are usually built on data of comparatively high frequency like daily, weekly or monthly data. Many quantities of interest, however, tend to be formulated on much lower frequencies. For example, investors might be interested in portfolio risk and return on an annual basis or on even longer time horizons when their goal is to build up wealth for their retirement. Naturally, their attitude towards risk will be formulated with regards to similar time horizons, independent of the time scale that is used for modeling. For example, investors might be required to state some level of percentage loss that they could afford on an annual basis. In contrast, it obviously wouldn't make sense to let an investor define this level in terms of five minute returns, where investors can not related to in any way. Hence, also risk targets are usually formulated in terms of comparatively long time horizons. In order to eliminate mismatches between time horizons of model and risk target, risk and return measures need to be made comparable. This is usually done by scaling higher-frequency measures to longer horizons. For example, one possible application could be to translate a daily Value-at-Risk into annual terms.

First and foremost, let's be very clear about what we are trying to achieve here. Scaling asset return moments like mean returns, volatility or correlations is far from being an obvious or easy thing. In reality, the exact type of scaling can be influenced by almost any time series property that one could think of. Whether it is marginal distributions that could be skewed, moments that could be subject to time-variation like volatility clusters or the existence of serial dependence like long memory. All such characteristics will affect the scaling behavior. Of course, for real data we do not know the true data generating process anyways. But even if we did, it is highly questionable that we ever would be able to infer analytic formulas that really capture the true scaling behavior.

These intricacies might easily be avoided by assuming independence over time while working with logarithmic returns. Multi-period returns then can be written as sum of single period returns, and moments scale easily according to

$$\begin{aligned} r_{0:T} &= r_1 + \ldots + r_T \\ \mathbb{E}[r_{0:T}] &= T\mathbb{E}[r_t] \\ \sigma(r_{0:T}) &= \sqrt{T}\sigma(r_t) \end{aligned}$$

In contrast, while aggregation over assets is simple and linear for discrete returns, aggregation over different time periods becomes a non-linear operation. Multi-period returns are not a sum of individual single period returns, but individual gross returns have to be multiplied:

$$(1 + r_{t:T}) = (1 + r_t) \cdot \ldots \cdot (1 + r_T)$$

Hence, even with independence over time, volatility does not scale with square-root-of-time for discrete returns, and expected returns do not scale linearly:

$$\begin{aligned} \mathbb{E}[r_{0:T}] &\neq T\mathbb{E}[r_t] \\ \sigma(r_{0:T}) &\neq \sqrt{T}\sigma(r_t) \end{aligned}$$

So we need to find another way (which actually means: based on different assumptions) to at least somewhat realistically scale return moments over different time horizons analytically.

That way daily or weekly risk measures shall be made more easily interpretable and better comparable to meaningful long-term targets.

The way we will do this is as follows. First, we translate discrete return moments to logarithmic return moments. We will do this by assuming that discrete returns are log-normally distributed. Then, we will scale logarithmic moments under the assumption of independence over time via the usual analytic formulas. And last, logarithmic moments are translated back to discrete moments. While the transformation from logarithmic to discrete moments was already determined in Lemma (1.2.2), we still need to lay out the opposite direction:

Lemma 1.2.3: [Discrete to logarithmic return moments]

Let $\mu(r)$ denote the mean value of net return r, and $\sigma(r)$ denote its standard deviation. Furthermore, let $r^{\log}$ denote the associated logarithmic return. Then expectation and variance of $r^{\log}$ can be calculated under the assumption of normally distributed logarithmic returns by

$$\mu(r^{\log}) = \log\left(\frac{(\mu(r)+1)^2}{\sqrt{(\mu(r)+1)^2+\sigma^2(r)}}\right)$$
$$\sigma^2(r^{\log}) = \log\left(1+\frac{\sigma^2(r)}{\mu^2(r)}\right)$$

Proof. See Appendix (6.0.1).

□

Summing up, we now have a way to analytically map daily asset return moments to annualized values. In order to maximize the number of observations available for asset moment estimation we use a daily frequency for modeling. Using some restrictive assumptions for return distributions and dependence over time we can analytically translate daily values into more meaningful annualized numbers. In case that the restrictive assumptions are not fulfilled in reality, calculated annualized values at least should be a reasonable approximation.

Chapter 2

Single-period portfolio selection

Quite generally, the portfolio selection problem can be separated into two components. First, for each possible portfolio weights the associated portfolio return distribution needs to be derived. This basically defines the outcome for each chosen portfolio weights. And second, among all possible portfolio outcomes, one has to choose a most favorable portfolio based on risk-return considerations. Therefore, one needs to define some kind of classification that allows the evaluation of portfolios with regards to investors' preferences. One possible way of defining such a classification would be by formulating the problem in terms of expected utility. Let w denote some given portfolio weights, U some given utility function, $\omega \in \Omega$ all possible states of the environment, and $P_t(w, \omega)$ the portfolio value of the portfolio with weights w in state ω. Then the expected utility maximization problem can be formulated as

$$w^{\star} = \arg\max_{w} \mathbb{E}[U(P_t(w, \omega))] = \arg\max_{w} \mathbb{U}(w)$$

In general, utility maximization requires knowledge of of the full distribution of portfolio values. However, the information required can be dramatically reduced for the case of either normally distributed returns or quadratic utility functions (Levy and Markowitz 1979). The only relevant values then are expectation and standard deviation of portfolio returns: μ_P and σ_P. More precisely: "the investor does (or should) consider expected return as desirable thing *and* variance of return an undesirable thing" (Markowitz 1952). This assumption for investors' preferences is the fundament of modern portfolio theory, and it is also an assumption that we will use in our dynamic asset management approach.

2.1 Markowitz portfolio selection

From the μ-σ utility function we can immediately follow that some portfolios are suboptimal and hence should never be chosen, independent of the actual risk aversion of the investor. The set of such inferior portfolios can be determined through the concept of Pareto-domination. A portfolio is Pareto-dominated if there is another portfolio with equal risk but higher return, or equal return but lower risk. Hence, we can reduce the set of all portfolios by removing all Pareto-dominated portfolios such that only an efficient set remains. The following derivation of this efficient set will largely build on (T. Roncalli 2013).

There are two different ways how the efficient set of portfolios can be characterized. First, portfolios can be directly identified as being Pareto-dominating. Therefore, a portfolio needs to be solution to the following two optimization problems. For some given level of portfolio expected return, it must minimize portfolio volatility:

$$\begin{aligned} w^\star = &\arg \min_w \quad \sigma_P \\ &\text{subject to} \quad \mu_P = \mu^\star \\ &\qquad\qquad\quad \sum_{i=1}^{n} w_i = 1 \quad \text{(weight constraint)} \end{aligned}$$

And for some given level of volatility it must maximize the expected portfolio return:

$$\begin{aligned} w^\star = &\arg \max_w \quad \mu_P \\ &\text{subject to} \quad \sigma_P = \sigma^\star \\ &\qquad\qquad\quad \sum_{i=1}^{n} w_i = 1 \quad \text{(weight constraint)} \end{aligned}$$

Any portfolio that fulfills these both criteria of optimality can not be Pareto-dominated and hence is part of the efficient set. A different way of characterizing the set of efficient portfolios is through the perspective of a μ-σ preference function. The portfolio selection problem then can be formulated as follows:

$$\begin{aligned} w^\star(\phi) = &\arg \max_w \quad \mu_P - \frac{\phi}{2}\sigma_P^2 \\ &\text{subject to} \quad \sum_{i=1}^{n} w_i = 1 \quad \text{(weight constraint)} \end{aligned}$$

It can be shown that any portfolio that maximizes this tradeoff between expected return and volatility for some value of ϕ will be part of the efficient frontier. The other direction also holds: for any portfolio of the efficient frontier there will exist a value ϕ such that the portfolio is the solution to this optimization problem. The border solutions will be obtained for extreme values of ϕ. On the one hand, for $\phi = 0$ it will maximize portfolio expectation μ_P independent of portfolio risk, thereby representing the case of risk-neutral investors. For $\phi = \infty$, on the other hand, $w^\star$ will be the global minimum variance portfolio.

Lemma 2.1.1: [Global minimum variance portfolio, (Kempf and Memmel 2005)]

The global minimum variance portfolio is the solution to the following optimization problem:

$$
\begin{aligned}
w_{\mathsf{GMV}} = & \arg\min_{w} \quad \sigma_P^2 \\
& \text{subject to} \quad \sum_{i=1}^{n} w_i = 1 \quad \text{(weight constraint)}
\end{aligned}
$$

With given covariance matrix Σ_R, its weights can be determined by:

$$
w_{\mathsf{GMV}} = \frac{\Sigma_R^{-1}\mathbf{1}_d}{\mathbf{1}_d'\Sigma_R^{-1}\mathbf{1}_d}
$$

Now that the set of portfolios has been reduced to the set of efficient portfolios, we still need to determine a single portfolio out of the full efficient frontier. One way of settling for a specific portfolio naturally derives from the optimization problem in terms of μ-σ preferences. ϕ can be interpreted as coefficient of risk-aversion, and it determines the tradeoff that an investor has between risk and return. Once ϕ has been uniquely determined, the optimization problem will have a single solution that implicitly meets the investor's risk-return preferences. This way, the optimal portfolio and hence the portfolio expectation and volatility will vary depending on the market environment. When return is rather costly in terms of risk, we might settle for a portfolio with lower return and lower volatility, while we might want to increase expected returns and risk in good times.

Although this approach of selecting a portfolio very naturally complies with μ-σ preferences, and also automatically adapts the risk taken with regards to the specific market environment, we still will choose a different way of determining the optimal portfolio. Our mechanism of choosing the optimal portfolio will use a risk target that the portfolio should match. In other words, among all portfolios of the efficient frontier we will pick the portfolio that matches the given target volatility the closest. Advantage of this approach is that the actually choice of the portfolio will not rely on the portfolio expectation. In a setting of perfectly known asset moments it would not make sense to not also make use of this information. As we will see later on, however, expected returns are very hard to estimate in reality, and by construction we will overestimate them for portfolios on the efficient frontier. Hence, it might be a more robust approach to leave the expected portfolio return out of this final portfolio decision.

In a world of known asset moments, however, such a volatility targeting approach should not be optimal, of course, as one will always settle for the same risk, no matter how beneficial it is. In other words, even if the efficient frontier is almost flat, and additional risk is basically not compensated at all, we will still pick a portfolio with large volatility. With estimated moments, however, individual portfolio decisions will become more robust when expected returns are not involved at this step also.

2.1.1 Computational efficiency

In practice, Markowitz portfolio selection is almost never implemented without additional constraints on the individual asset positions. The most prominent such constraint is to forbid short-selling: all weights $(w_i)_i$ then need to be larger than zero. As we will enforce no short-selling throughout the rest of the thesis, Definition (2.1.2) shall formalize this, such that we will not explicitly add these basic weight constraints to all following optimization problems anymore.

Definition 2.1.2: ***[Basic weight constraints]***

There are two basic weight constraints that shall apply to all of the portfolio selection optimization problems that follow:

$$\sum_{i=1}^{n} w_i = 1 \quad \textit{(weight constraint)}$$

$$w_i \geq 0 \quad \textit{(short-selling constraint)}$$

Many financial institutions are also subject to further regulatory requirements in order to prevent excessive risk taking. For example, pension funds might not be allowed to exceed a certain fraction of risky assets, such that the overall weight on equities needs to be capped from above. When additional constraints are incorporated into the optimization framework, however, the optimization problem needs to be solved numerically. However, some portfolio optimization problems with linear weight constraints are part of a subset of optimization problems with particularly nice properties: quadratic programming problems.

Definition 2.1.3: *[**Quadratic programming**, (Cornuejols and Tütüncü 2006)]*

Let w denote a d-dimensional vector with lower and upper bounds lb and ub. Let further A, B, a and b define linear equality and inequality constraints. Then the following optimization problem with quadratic objective function given by Q and c is called a quadratic programming problem:

$$\begin{aligned} w^{\star} = \arg\min_{w} \quad & \frac{1}{2} w' Q w + c' w \\ \textit{subject to} \quad & Aw \leq a \\ & Bw = b \\ & lb \leq w \leq ub \end{aligned}$$

Quadratic programming comes with two properties that are of major importance. First, there exist a couple of fast and efficient algorithms to find an optimal solution. And second, quadratic programming is part of a larger class of optimization problems called convex optimization problems. As such, any optimum automatically is known to be a global solution, ruling out the risk of accidentally picking up on a suboptimal local optimum.

Writing portfolio expectation and variance in terms of matrix multiplications like in Equation (1.1) and (1.2), Markowitz minimization of portfolio variance can be shown to be a quadratic programming problem.

Theorem 2.1.4: [Markowitz and quadratic programming, (T. Roncalli 2013)]

Minimization of portfolio volatility with given expected portfolio return is a quadratic programming problem.

Proof. Replacing μ_P and σ_P in the original formulation of the portfolio volatility minimization problem

$$\begin{aligned} w^{\star}(\mu^{\star}) = \arg\min_{w} \quad & \sigma_P \\ \text{subject to} \quad & \mu_P = \mu^{\star} \end{aligned}$$

immediately makes the optimization a quadratic programming problem:

$$\begin{aligned} w^{\star}(\mu^{\star}) = \arg\min_{w} \quad & w'\Sigma_R w \\ \text{subject to} \quad & w'\mu = \mu^{\star} \end{aligned}$$

□

2.1.2 Markowitz deficiencies

A lot of criticism exists on the original plain Markowitz portfolio selection approach, probably encouraged by its usually poor performance in empirical applications. One such criticism is that μ-σ optimization will only lead to utility maximizing portfolio choices in the context of either elliptically distributed returns or quadratic utility. For other, potentially more realistic settings, the optimal portfolio can not be found on the grounds of μ and σ alone, but either needs to be derived from the full distribution of portfolio returns, or from more sophisticated measures of risk like Value-at-Risk and Expected Shortfall. However, (Levy and Markowitz 1979) find that although μ-σ optimization might not always bring forth the true expected utility maximizing portfolio, it seems to be a very good approximation in most cases. Compared to other possible sources of errors, like measurement of risk-aversion and estimation of asset moments, the imprecision introduced by μ-σ optimization actually might be an issue of subordinate importance.

Another frequently mentioned criticism is that fully unconstrained Markowitz portfolios tend to be highly leveraged, while they are heavily concentrated in the case of no short-selling constraints. Although neither leverage nor concentration for itself should negatively affect expected utility, there still might be arguments to disfavor both in practice. For example, overly concentrated portfolios usually raise questions of investors regarding clustering of risks.

The most important criticism, however, is that Markowitz portfolios might be optimal when asset moments are fully known, but this not necessarily leads to good portfolio choices when asset moments are unknown and have to be estimated. As the full consequences of this point are not easy to see, we will now present the argument with more details. Thereby we will mostly rely on simulated examples, in order to make theoretical results more illustrative. As true underlying distribution for simulated data we will rely on the same set of assets that we will also use in the empirical application in Chapter (4). There, we will also introduce the data in further details. For now we only use sample means and sample covariance matrix to set up a realistic use case to illustrate the issue of estimation errors. For the distribution of asset returns we will use a simple multivariate normal distribution. The presentation of the theory on Markowitz with estimation errors will follow arguments made in (R. O. Michaud and Michaud 2008) and (Scherer 2002).

When asset moments are not known they need to be estimated from data. This will introduce estimation errors, which can be seen by repeatedly estimating asset moments from simulated data samples. Let μ and Σ_R denote the true vector of expected returns and the covariance matrix respectively, and let $F \sim \mathcal{N}(\mu, \Sigma_R)$ denote the multivariate distribution of asset returns. Let further $K = 1000$ be the number of repetitions, and T be the number of observations per repetition. For each repetition $i = 1, \ldots, K$ we simulate T observations from F and estimate asset sample moments $(\mu^{\{i\}}, \Sigma_R^{\{i\}})$. Figure (2.1) shows true asset moments as black dots, together with estimated asset moments for each simulated sample. Each color represents an asset, and the sample size T equals 40 observations in the left part of the plot, and 200 observations in the right part. Two well-known results can be directly seen from the figure. First, the less volatility, the smaller the estimation error. Hence, moments of low-risk assets on the left can be estimated more precisely. And second, the larger the sample size, the better the estimator.

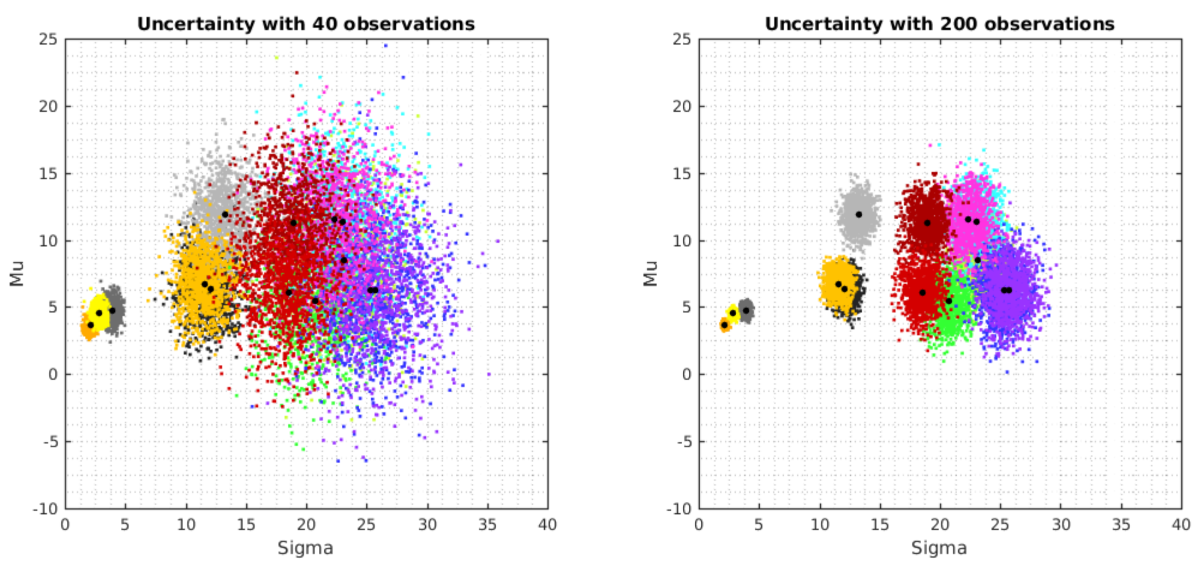

Figure 2.1: Estimation errors of asset moments with normally distributed returns

As portfolio moments are a function of portfolio weights and asset moments, the estimation error of asset moments directly translates into an estimation error of portfolio moments. Let $w^{\{1\}}$ and $w^{\{2\}}$ denote fix portfolio weights of two different portfolios, labeled as *FixPf1* and *FixPf2*. The left part of Figure (2.2) shows true portfolio moments together with true asset moments. In contrast, estimated portfolio moments of both portfolios are shown on the right, with one dot for each simulated data sample where portfolio moments have been estimated.

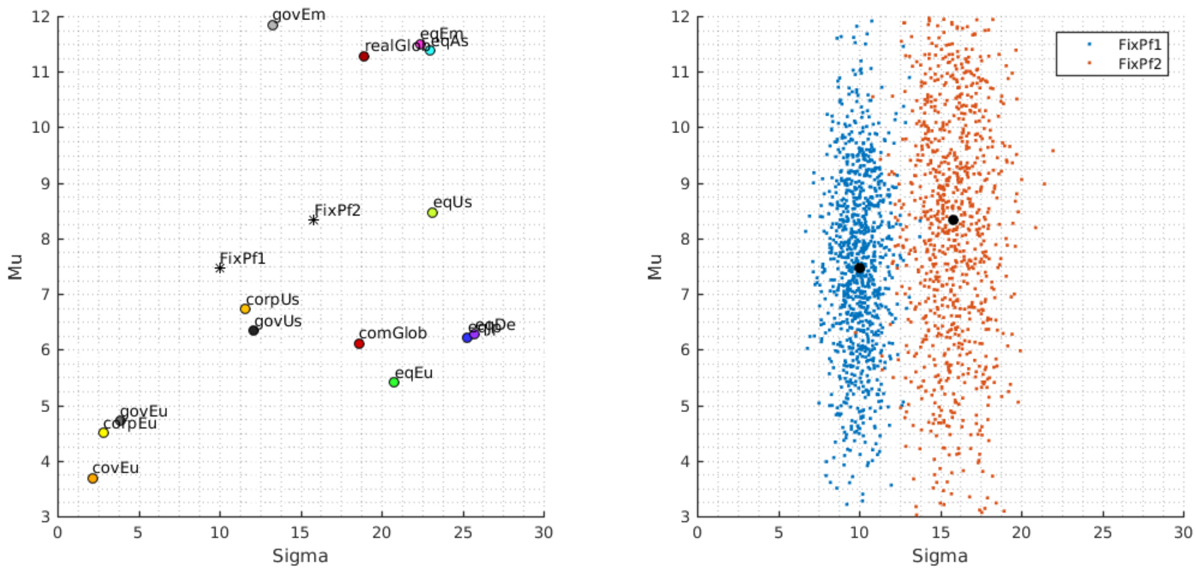

Figure 2.2: Estimation error of portfolio moments

The set of estimated portfolio moments gives an indication about the distribution of the portfolio moment estimator for both portfolios. This distribution could be put to use, for example, do define a measure for statistically significant portfolio differences. As can be seen, both portfolios are hard to distinguish in terms of expected returns, while they almost do not overlap in terms of portfolio volatility. Such a measure of portfolio differences could be used to determine whether it is worthwhile to trade from an initial portfolio to some portfolio target.

In reality, of course, we do not have the ability to simply simulate portfolio moments based on samples from true asset moments, as they are obviously unknown. However, estimation errors of portfolio moments can also be determined quite well through re-sampling techniques like bootstrapping. To show this, we simulate a single data sample from F, which we take as single observed historic data sample. From this data sample we once estimate sample expected portfolio returns and volatilities, shown as green dots in Figure (2.3). For comparison, the graphics also shows true portfolio moments in black. Furthermore, we re-sample from this original data sample in order to bootstrap portfolio moments.

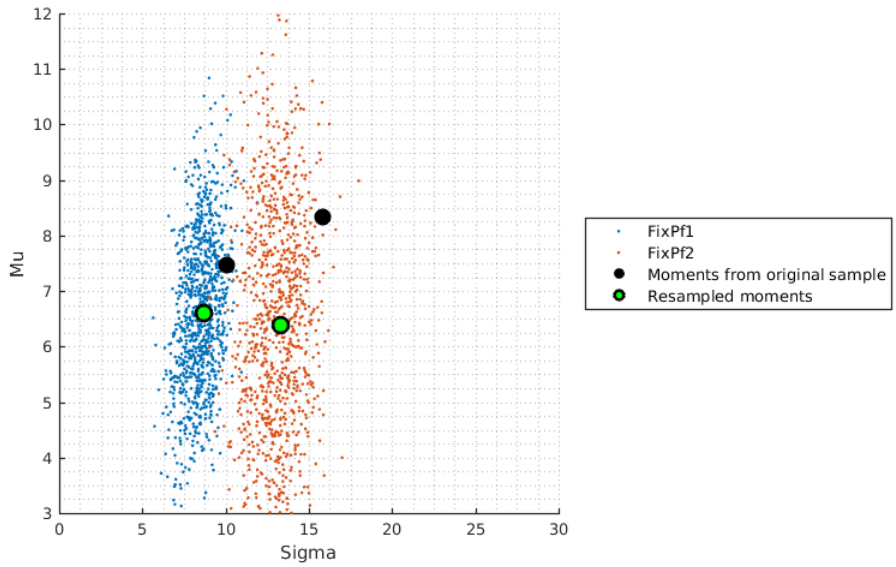

Figure 2.3: Bootstrapped portfolio moments

As can be seen, due to the particular sample that was initially simulated sample moments and true portfolio moments deviate. This deviation also carries forward into all re-sampled data samples, as they are all drawn from this original data sample. This way the centers of both re-sampled portfolio moment clouds do differ from the true portfolio values, but the dispersion of both clouds is almost comparable to what we did get through simulation in Figure (2.2). Hence, bootstrapping is a good way of quantifying the uncertainty around estimated portfolio moments. However, keep in mind that we were estimating portfolio moments for portfolios with fixed weights so far. In more realistic settings, portfolio weights will usually depend on what we estimate for asset moments, as we want to focus on assets with promising risk-return profiles. This is what we will change in the next step.

So from now on portfolio weights shall not be given upfront anymore, but they are chosen based on estimated asset moments. This way portfolio weights will depend on the actual data sample, and hence become a random vector themselves. Results for unconstrained Markowitz mean-variance optimization with expected return target of $\mu_P = 7$ are shown Figure (2.4). The leftmost chart shows box plots for individual asset weights, while the chart in the middle compares true optimal portfolio weights with averages of simulated optimal portfolio weights. And on the right we see the joint distribution of simulated asset weights for two different assets: European corporate bonds and European covered bonds. As can be seen, there exists large variation for individual optimal asset weights. For example, European corporate bond weights range from almost 400% to approximately -150%, with true optimal asset weight of roughly 90%. From the chart in the middle we can see that asset weights are not fully unbiased. In particular, European covered bonds on average get positive weight, although the true optimal weight is negative. This also could be caused by numerical instabilities, as solving for optimal portfolio weights requires to invert the covariance matrix, which sometimes introduces numerical problems. From the joint distribution in the right chart we can infer two things. First, the joint distribution itself looks rather well-behaved, with almost elliptical shape. And second, European corporate and covered bonds seem to be substitutes. The negative correlation between both assets indicates that together they should generate some level of exposure to European bond markets, and either one of both assets could be chosen to achieve this.

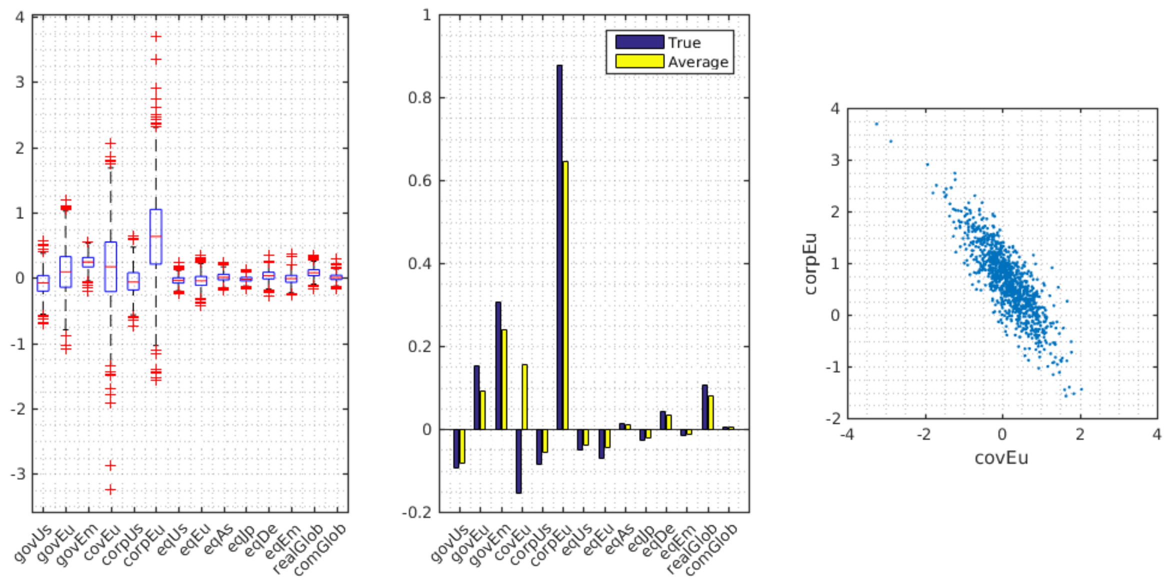

Figure 2.4: Estimation error of optimal portfolio weights

When using Markowitz with short-selling constraints, however, graphics change significantly, see Figure (2.5). As asset weights are capped from below, many assets are not part of the true optimal portfolio anymore. However, repeated simulation will create some “optionality” effect for these assets: due to randomness they could become part of the optimal portfolio in some samples, and this effect can not be offset in other samples, as weights are not allowed to become negative. This way, these assets will get slightly positive weights on average. This optionality also alters the joint distribution of asset weights, as can be seen in the right chart. The previously almost elliptical joint distribution now has become much more degenerated.

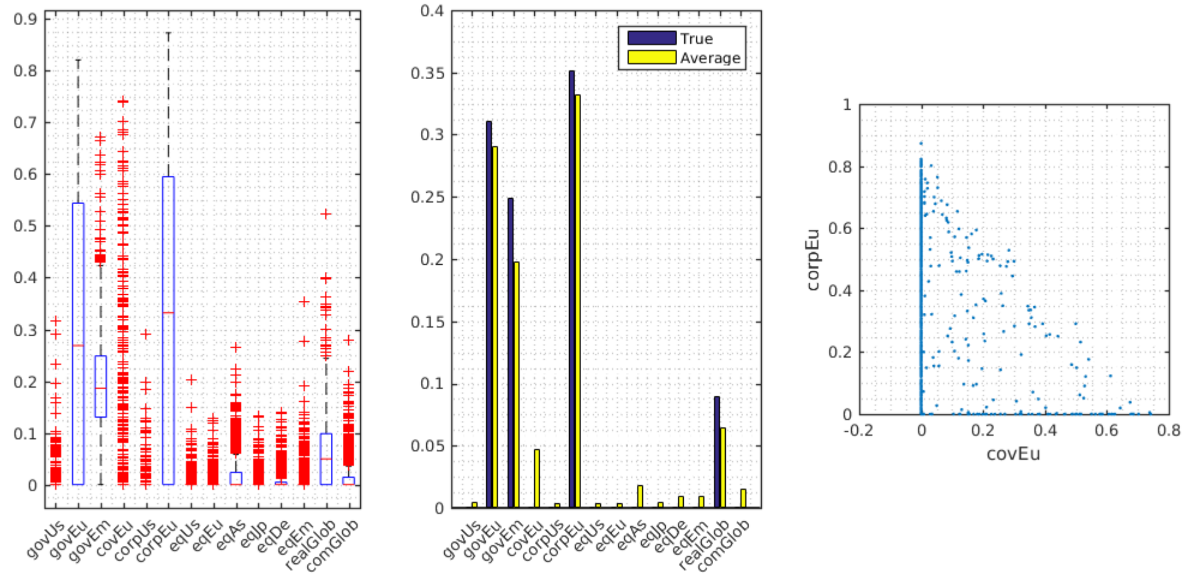

Figure 2.5: Estimation error of optimal portfolio weights with short-selling constraints

Portfolio weights themselves are usually not a direct input to utility functions, but utility depends on risk-return profiles of the chosen portfolio. Quite large differences in asset weights could still result in similar portfolio properties, and hence estimation errors of asset weights do not necessarily translate into proportional losses in terms of expected utility. Hence, in the next step we will analyze the distribution of final portfolio moments. For the case of unconstrained Markowitz this can been seen in Figure (2.6), where final portfolio moments are shown for different simulated data samples.

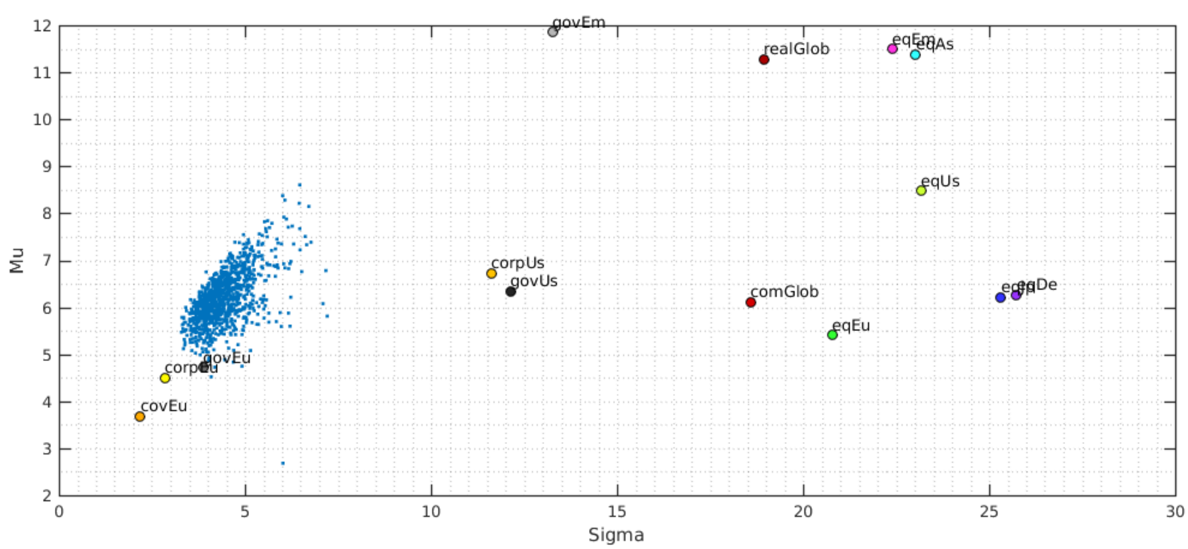

Figure 2.6: Portfolio moments of portfolios chosen from estimated asset moments

This randomness in portfolio moments also translates into randomness of expected utility. Consequently, we do not get the true optimal expected utility $\mathbb{U}(w^{\star})$ when asset moments are unknown, but expected utility depends on the actual weights $\hat{w}^\star$ that are chosen. Hence, any portfolio selection strategy should not be evaluated by the expected utility that is obtained when asset moments are known, but by taking into account that portfolio choices generally will be suboptimal when asset moments need to be estimated. The overall final utility that we have to expect can be obtained by integrating over all possible data samples:

$$\mathbb{E}[\mathbb{U}(\hat{w}^\star)]$$

This is the measure of utility that should be used to compare the appropriateness of different portfolio selection approaches, and it will be called out-of-sample utility in the following. Using out-of-sample utility, the evaluation of different portfolio selection approaches will depend on all components that were used during the determination of final portfolio weights. In particular, the following two components are crucial:

- the estimator used to derive asset moments
- the strategy that maps estimated asset moments to portfolio weights: $(\hat{\mu}, \hat{\Sigma}_R) \mapsto \hat{w}^\star$

So even if a given portfolio selection strategy maximizes expected utility when asset moments are known, it still might not maximize out-of-sample utility when asset moments can only be estimated with high imprecision. For example, forbidding short-selling will clearly decrease expected utility when asset moments are known. However, it could still provide larger out-of-sample utility when estimation errors are large. Figure (2.7) shows the distribution of portfolio moments for the case of Markowitz with no short-selling constraints.

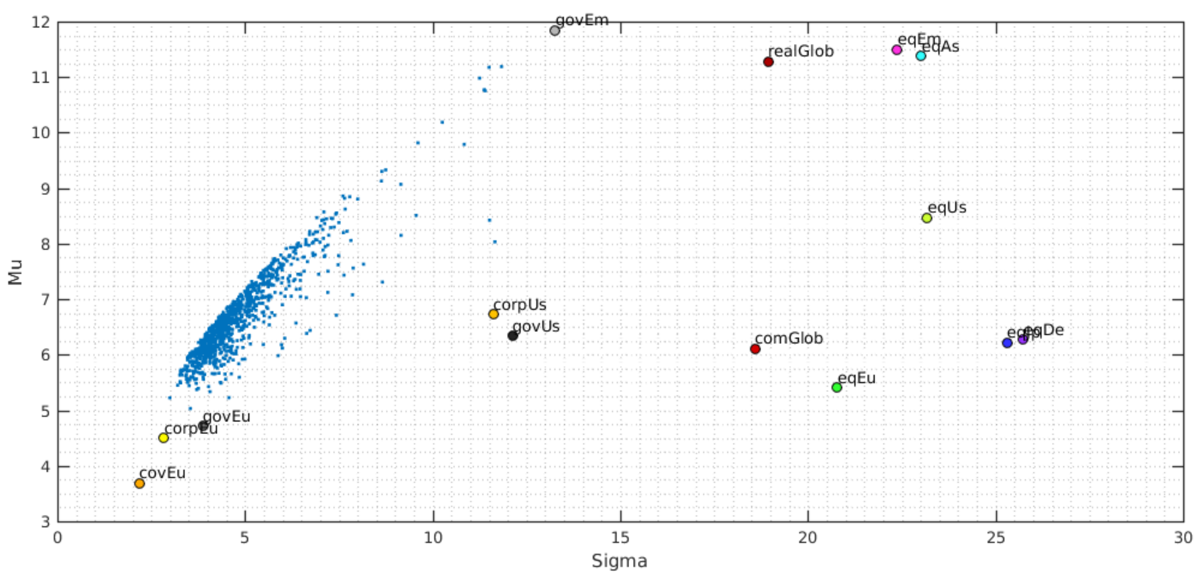

Figure 2.7: Out-of-sample portfolio moments with short-selling constraints

Comparing portfolio outcomes for unconstrained and no short-selling Markowitz in Figures (2.6) and (2.7) respectively, one can guess that unconstrained optimization really could become an inferior selection approach in terms of out-of-sample utility in some cases. However, neither of both approaches will be generally superior to the other one, but the best approach varies with regards to the level of estimation error. The more precisely asset moments can be estimated, the more we can benefit from leveraging by exploiting small differences in asset risk-return profiles. We will not go into further details in the present comparison of both approaches, but are content to pick up the following conclusion: in-sample superior portfolio selection approaches may become inferior in terms of out-of-sample utility due to estimation error. Overconfidence in estimates of asset moments will underestimate estimation errors, and hence might incline investors to try to exploit differences in assets' risk-return profiles too aggressively. This would ultimately lead to lower out-of-sample utility.

In reality we do not have the convenience to simply derive out-of-sample utilities by simulation, as we do not know the true asset moments to simulate from. Again, bootstrapping might be tried as fallback. However, when portfolio weights are not fixed upfront, but are chosen conditional on estimated asset moments, estimation of out-of-sample utility is subject to overestimation. In order to illustrate this point, let's just assume that expected utility is derived from μ-σ preferences. Then there exists some function H^U such that expected utility can be formulated as

$$\mathbb{U}(w) = H^U(\mu_P(w), \sigma_P^2(w)) = H^U(w'\mu, w'\Sigma_R w)$$

A prominent example for such a function is given by

$$H^U\left(\mu_P(w), \sigma_P^2(w)\right) = \mu_P(w) - \frac{\phi}{2}\sigma_P^2(w)$$

Thereby ϕ can be interpreted as risk-aversion coefficient. Let's further assume that we follow some given portfolio selection approach, which maps all possible estimated asset moments to some uniquely defined portfolio decision: $(\hat{\mu}, \hat{\Sigma}_R) \mapsto \hat{w^\star}$. Then for the true expected utility of one such chosen portfolio we get

$$\mathbb{U}(\hat{w^\star}) = H^U\left(\mu_P(\hat{w^\star}), \sigma_P^2(\hat{w^\star})\right) = H^U\left(\hat{w^\star}'\mu, \hat{w^\star}'\Sigma_R\hat{w^\star}\right)$$

As we do not know the true values of μ and Σ_R, we might try to derive an estimate of expected utility by plugging in asset moment estimates.

$$\hat{\mathbb{U}}(\hat{w^\star}) = H^U\left(\hat{w^\star}'\hat{\mu}, \hat{w^\star}'\hat{\Sigma}_R\hat{w^\star}\right)$$

This is problematic, however, as chosen optimal portfolio weights are generally not independent of estimated asset moments.

$$\hat{w^\star} \not\perp\!\!\!\perp \left(\hat{\mu}, \hat{\Sigma}_R\right)$$

In fact, we usually want to explicitly pick assets with good risk-return profiles, such that assets with overestimated μ values have a tendency to be picked more than assets that have lower estimated μ values by chance. This way, estimates of portfolio moments become biased, and so does the estimate of expected utility.

This biased estimation of portfolio moments can be seen quite nicely when looking at simulated portfolio selections. So let's now compare true and estimated portfolio moments for portfolios that are selected by short-selling constrained Markowitz. As target return we take two different targets: $\mu_P = 4.5$ and $\mu_P = 9$. For both portfolio selection strategies we simulate asset returns from a given joint asset return distribution, estimate asset moments for the simulated data and select the portfolio that we get by portfolio optimization. For this portfolio we then evaluate the moments twice: once with true and once with estimated asset moments. This procedure is repeated 1000 times, both for data samples of size 40 as well as for samples with 200 observations. Results are shown in Figure (2.8), where true out-of-sample portfolio moments are represented by red dots and in-sample portfolio moments by blue ones.

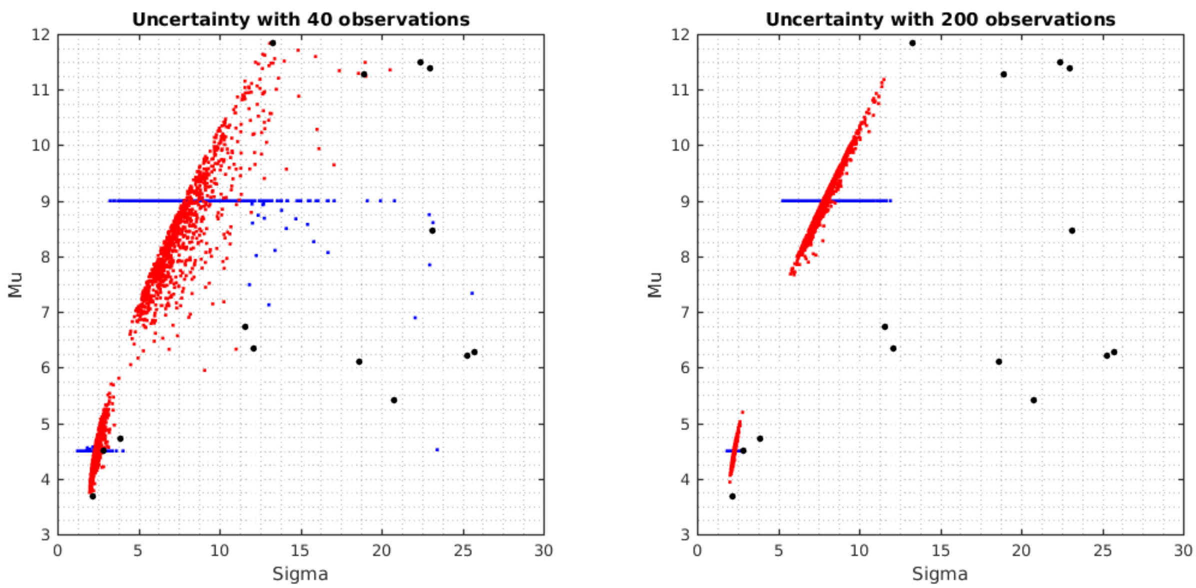

Figure 2.8: True versus estimated portfolio moments

As can be seen from the graphics, estimated portfolio moments do not just follow well-behaved and almost elliptical distributions like in Figure (2.2), but they basically all lie on a single line except for very few outliers. Reason for that is the interaction between the asset moments that are estimated and the portfolio that gets chosen. In other words, we just pick the portfolio weights such that the portfolio will fulfill the expect return target under estimated asset moments. The only way that the estimated portfolio expectation can deviate from its return target is when estimated moments are such that the target return is not feasible. In other words, there exists not a single portfolio with given constraints that would achieve the required expected return under the estimated asset moments. In these cases we just pick a portfolio that is as close to the expected return target as possible. In contrast to the blue dots, which show what we think we get as portfolio characteristics, red dots show the risk-return profiles that we would actually get. The higher the return target and the smaller the sample size, the more these actual risk-return portfolio profiles are spread apart. Lower return targets do coincide with less risky portfolios, which diminishes estimation errors. So do large sample sizes.

Although one already might guess from Figure (2.8) that there is some kind of systematic bias in estimated portfolio moments, there is an even better way of showing this. Instead of comparing true and estimated portfolio moments for single portfolios only, we will now do this for multiple portfolios simultaneously. Therefore we set up a range of 20 target portfolio returns, such that we effectively will compare true and estimated moments of the efficient frontier. First of all, however, we derive the true efficient frontier that we get when asset moments are known. Again, we will use Markowitz with short-selling constraints, and results are shown in Figure (2.9). In the left part of the figure we can see the portfolio moments for the 20 portfolios on the efficient frontier, while associated portfolio weights are shown on the right. The leftmost point of the efficient frontier coincides with the global minimum variance portfolio, which has slightly lower portfolio variance than the minimum variance asset due to diversification effects. In contrast, the rightmost portfolio simply represents the asset with maximum expected return, which in this case is government emerging markets bonds. As already pointed out before, efficient portfolios are highly concentrated, with no efficient portfolio comprising more than four assets with individual asset weight larger than 5%. The whole frontier is basically dominated by five assets: European covered bonds, European corporate bonds, European government bonds, global real estate and government emerging markets bonds.

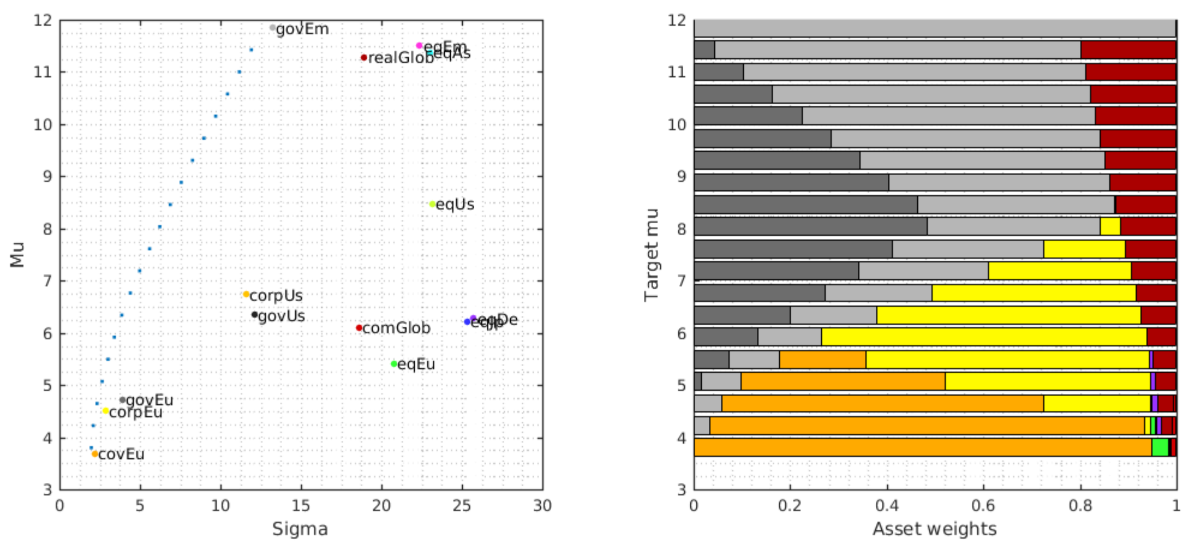

Figure 2.9: Efficient frontier portfolio moments and portfolio weights

We now estimate efficient frontiers for the same expected return targets as in Figure (2.9). Due to randomness, we get slightly different efficient frontiers for each of the 1000 repetitions. Again, we can translate the obtained portfolio weights into portfolio moments in two different ways. First, by using true asset moments, and second, by using the same estimated asset moments that were used to derive the efficient portfolios in the first place. Similar to before, Figure (2.10) shows true efficient frontier moments in red, while estimated moments are plotted in blue. The slightly thicker black line in the middle shows the true efficient frontier evaluated with true asset moments. By construction, all red lines have to lie below the true efficient frontier. Both red and black lines are evaluated with the same true asset moments, but only the true efficient frontier was also derived by using them in the optimization. Hence, no portfolio can be better than the true efficient frontier in terms of true asset moments. Comparing blue and red lines one can see that blue frontiers are systematically better on average, with either higher expected returns, lower portfolio variances or even both. In other words, in most cases we think we have portfolios with better risk-return profiles than what we actually get. Evaluating portfolio moments with the same estimated asset moments that were used to derive them generally leads to overestimation.

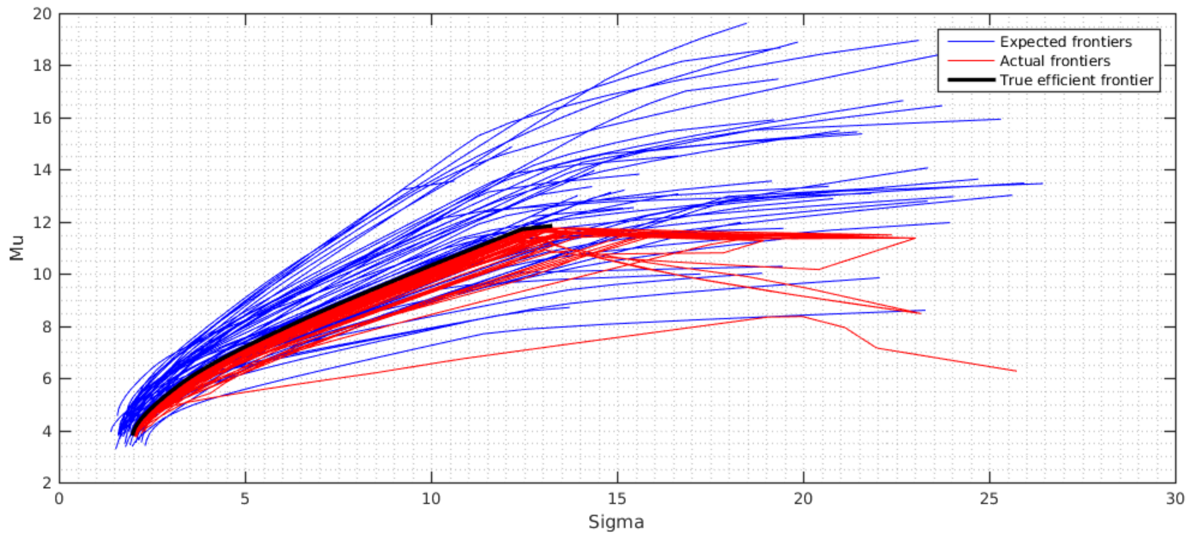

Figure 2.10: True versus estimated efficient frontiers

Now let's embed these insights into the context of our original problem. In real life, we only have one sample of historic data and possibly multiple different portfolio selection approaches at our hands. From all these portfolio selection strategies, we want to use the one that maximizes out-of-sample utility. But true out-of-sample utility we can only get from simulation, and estimating out-of-sample utility by using estimated portfolio moments just has been shown to be a highly unreliable and biased approach. It would always favor more flexible approaches that exploit differences in asset moments more aggressively, even though less flexible approaches could be maximizing true out-of-sample utility when uncertainty around asset moments is high. As a consequence, one might try to correct the bias introduced in the estimation of portfolio moments in order to make portfolio selection strategies of different levels of flexibility comparable. In some sense this is similar to the issue of overfitting, where one usually corrects for the number of parameters involved to provide a level playing field for different models. One possible way of correcting such a bias of portfolio moments could build on bootstrapping. However, we will not further pursue this approach due to the following reason: estimation risk is only one type of uncertainty that we have to face in reality, and it probably isn't even the most serious one.

Estimation risk generally denotes a type of uncertainty that can be eliminated through increasing sample sizes. For example, let's say that we want to estimate expectation and variance from some random variable. Given that we use consistent estimators, then the estimated values will converge to the true values as the number of observations increases to infinity. Hence, the uncertainty around estimated expectation and variance will vanish as soon as we have enough historic observations. And, even better, with estimation error we even can quantify the amount of uncertainty, as asymptotic distributions can be derived for many estimators.

A completely different type of uncertainty, however, is called model risk. It represents the risk that we estimate some quantity by using a fundamentally wrong structure, such that the estimate will not converge to the true value even for the case of infinite sample data. For example, estimating quantiles of a normally distributed random variable by fitting a Student's t-distribution will never be correct, no matter how much data we get. Of course, this is only an extreme example, and in reality model risk will most likely enter through much more subtle ways.

One such way how model risk could enter into the equation is due to time-varying asset moments. Especially when portfolio decisions have to be made with regards to short-term investment horizons, where it is not unconditional asset moments that matter, but conditional ones. In that case, recent historic observations might not be fully representative for the conditional moments in the near future. Let's illustrate this with an example, where we assume for simplicity that asset moments can be given by one of two possible regimes. Furthermore, let's assume that there was a regime change just five days ago, and we use sample moments of the last 100 days as estimator of current asset moments. Then by construction of this estimator, 95% of the observations that we use will come from the wrong regime and hence do not represent the true current state of assets. This would add a huge estimation error on top of what we would have if we just had constant asset moments and 100 observations to estimate them.

Summing up, in the presence of model risk out-of-sample utility of different portfolio optimization approaches is extremely hard to estimate reliably. Even with simulated data one would need to make assumptions on the actual type of misspecification that a certain estimator is exposed to. Especially for time-varying expected asset returns this might just end up in somewhat arbitrary manual interferences, possibly leading to distorted conclusions. Hence, we will refrain from further analyzing out-of-sample utilities in a single-period setting, but we will evaluate success of portfolio selection strategies in backtests. This way one can naturally test against realistic settings with uncertainty due to both estimation risk and model risk. Still, the analysis of single-period portfolio choices was far from pointless, as it did clearly deliver the following insight: not in-sample utility is the yardstick on which portfolio selection approaches should be evaluated, but out-of-sample utility. Simply optimizing with regards to estimated asset moments as if they were the true ones only appears optimal through the lens of the same wrongly estimated asset moments. Through this perspective unconstrained portfolio decisions are tremendously overestimated. In reality, approaches that acknowledge the fallibility of estimated moments more explicitly through additional robustness constraints will outperform. Being humble about the ability of asset moment predictions pays off more than blindly following estimated asset moments with overconfidence.

2.2 Fully uninformed market view

As we have seen in Section (2.1.2), being humble about the knowledge of asset moments might pay off in terms of out-of-sample utility. Let's for the moment exhaust this idea to the fullest and assume that we do not have any information about assets at all: no prior knowledge and no data. In reality, of course, there is plenty of data available for financial markets which generally could be used to draw conclusions about the behavior of individual assets. In particular, it is highly unlikely that all assets do have perfectly equal moments, and hence it might be worthwhile to try to use data in order to detect assets with particularly nice properties. Still, what would be the optimal investment strategy in case that we do not know anything about asset moments at all?

Even in that case we still might know that assets in reality are not equal. But as we do not have any data or knowledge about them, we can not distinguish between individual assets in any way. So basically each asset will appear the same, with equal expected value, equal volatility and the same correlation between all assets. We refer to this situation as "fully uninformed", and from this perspective the market can be modeled as follows:

Definition 2.2.1: ***[Fully uninformed market view]***

An investment situation shall be called fully uninformed market view when all assets are completely indistinguishable:

- *all assets have the same unknown expected return* $\bar{\mu}$
- *all assets have the same unknown variance* $\bar{\sigma}^2$
- *all pairwise correlations have the same value* $\bar{\sigma}_{ij}$

The fully uninformed covariance matrix hence is given by:

$$\bar{\Sigma}_R = \begin{bmatrix} \bar{\sigma}^2 & \bar{\sigma}_{ij} & \dots & \bar{\sigma}_{ij} \\ \bar{\sigma}_{ij} & \bar{\sigma}^2 & \dots & \bar{\sigma}_{ij} \\ \vdots & \vdots & \ddots & \vdots \\ \bar{\sigma}_{ij} & \bar{\sigma}_{ij} & \dots & \bar{\sigma}^2 \end{bmatrix}$$

As all assets have the same expected value $\bar{\mu}$, the expected portfolio return does not at all contribute to the portfolio decision: all portfolios have the same expectation $\bar{\mu}$ anyways. The portfolio variance, however, does depend on the actual portfolio weights that are chosen (at least when assets are not perfectly correlated). For example, investment into a single asset only will give a portfolio variance of $\bar{\sigma}^2$, while the variance of an equally weighted portfolio is

$$
\begin{aligned}
\sigma_{\mathrm{EW}}^2 &= w_{\mathrm{EW}}' \Sigma_R w_{\mathrm{EW}} \\
&= \sum_{i=1}^{d} w_{\mathrm{EW},i}^2 \sigma_i^2 + \sum_{i=1}^{d} \sum_{j=1, j\neq i}^{d} w_{\mathrm{EW},i} w_{\mathrm{EW},j} \mathrm{Cov}(R_i, R_j) \\
&= \frac{1}{d^2} \sum_{i=1}^{d} \sigma_i^2 + \frac{1}{d^2} \sum_{i=1}^{d} \sum_{j=1, j\neq i}^{d} \mathrm{Cov}(R_i, R_j) \\
&= \frac{1}{d^2} d \left(\frac{1}{d} \sum_{i=1}^{d} \sigma_i^2 \right) + \frac{1}{d^2} d(d-1) \left(\frac{1}{d(d-1)} \sum_{i=1}^{d} \sum_{j=1, j\neq i}^{d} \mathrm{Cov}(R_i, R_j) \right) \\
&= \frac{1}{d^2} d \bar{\sigma}^2 + \frac{1}{d^2} d(d-1) \bar{\sigma}_{ij} \\
&= \frac{1}{d} \bar{\sigma}^2 + \frac{(d-1)}{d} \bar{\sigma}_{ij} \qquad (2.1)
\end{aligned}
$$

Based on this formula we can derive two optimal portfolio allocation rules.

Theorem 2.2.2: [Optimality under fully uninformed market views]

Given that an investor has a fully uninformed market view, the following two investment strategies can be shown to optimize the portfolio selection:

1. for any given subset of assets the equal weights portfolio is a variance minimizing strategy
2. expanding the equal weights portfolio to as many assets as possible further reduces portfolio variance

In other words, when we do not have any knowledge about individual assets, we should invest in as many assets as possible, with wealth equally distributed between assets.

The second part of Theorem (2.2.2) can be seen by looking at the limit of Equation (2.1) in d. Holding $\bar{\sigma}^2$ and $\bar{\sigma}_{ij}$ fix, the portfolio variance is decreasing in d and converges to the average covariance $\bar{\sigma}_{ij}$:

$$
\lim_{d \to \infty} \sigma_{\mathrm{EW}}^2(d) = \bar{\sigma}_{ij}
$$

This limit can be interpreted as systematic risk that is common to all assets and can not be diversified away. The proof of Theorem (2.2.2), part 1, is a little bit more elaborate, and we need the inverse of the fully uninformed covariance matrix first.

Lemma 2.2.3: [Inverse of fully uninformed covariance]

Let Σ_R be a fully uninformed covariance matrix with variances equal to $\bar{\sigma}$ and covariances equal to $\bar{\sigma}_{ij}$. Then the inverse matrix of Σ_R is given by

$$\Sigma_R^{-1} = \begin{bmatrix} \bar{a}^2 & \bar{a}_{ij} & \dots & \bar{a}_{ij} \\ \bar{a}_{ij} & \bar{a}^2 & \dots & \bar{a}_{ij} \\ \vdots & \vdots & \ddots & \vdots \\ \bar{a}_{ij} & \bar{a}_{ij} & \dots & \bar{a}^2 \end{bmatrix}$$

with elements

$$\begin{aligned} \bar{a}^2 &= \frac{1-(d-1)\bar{\sigma}_{ij}\left(-\frac{\bar{\sigma}_{ij}}{\bar{\sigma}^4+(d-2)\bar{\sigma}_{ij}\bar{\sigma}^2-\bar{\sigma}_{ij}^2(d-1)}\right)}{\bar{\sigma}^2} \\ \bar{a}_{ij} &= -\frac{\bar{\sigma}_{ij}}{\bar{\sigma}^4+(d-2)\bar{\sigma}_{ij}\bar{\sigma}^2-\bar{\sigma}_{ij}^2(d-1)} \end{aligned}$$

Proof. [**Lemma (2.2.3)**]

To be the inverse of Σ_R, matrix Σ_R^{-1} must fulfill the following general equation:

$$\Sigma_R^{-1}\Sigma_R = \mathbf{I}_d$$

Due to the special nature of both matrices, all matrix elements of the product can be derived from one of two different equations:

$$\begin{aligned} \bar{\sigma}^2\bar{a}^2 + (d-1)\bar{\sigma}_{ij}\bar{a}_{ij} &= 1 \\ \bar{\sigma}^2\bar{a}_{ij} + \bar{\sigma}_{ij}\bar{a}^2 + (d-2)\bar{\sigma}_{ij}\bar{a}_{ij} &= 0 \end{aligned}$$

Solving the first equation for $\bar{a}^2$ gives:

$$\begin{aligned} \bar{\sigma}^2\bar{a}^2 + (d-1)\bar{\sigma}_{ij}\bar{a}_{ij} &= 1 \\ \bar{\sigma}^2\bar{a}^2 &= 1-(d-1)\bar{\sigma}_{ij}\bar{a}_{ij} \\ \bar{a}^2 &= \frac{1-(d-1)\bar{\sigma}_{ij}\bar{a}_{ij}}{\bar{\sigma}^2} \quad (\star) \end{aligned}$$

Plugging this into the second equation we can now solve for $\bar{a}_{ij}$:

$$
\begin{aligned}
\bar{\sigma}^2\bar{a}_{ij} + \bar{\sigma}_{ij}\bar{a}^2 + (d-2)\bar{\sigma}_{ij}\bar{a}_{ij} &= 0 \\
\bar{\sigma}^2\bar{a}_{ij} + \bar{\sigma}_{ij}\left(\frac{1-(d-1)\bar{\sigma}_{ij}\bar{a}_{ij}}{\bar{\sigma}^2}\right) + (d-2)\bar{\sigma}_{ij}\bar{a}_{ij} &= 0 \\
\bar{a}_{ij}\left(\bar{\sigma}^2 + (d-2)\bar{\sigma}_{ij}\right) + \frac{\bar{\sigma}_{ij}}{\bar{\sigma}^2} - \frac{\bar{\sigma}_{ij}^2(d-1)\bar{a}_{ij}}{\bar{\sigma}^2} &= 0 \\
\bar{a}_{ij}\left(\bar{\sigma}^2 + (d-2)\bar{\sigma}_{ij} - \frac{\bar{\sigma}_{ij}^2(d-1)}{\bar{\sigma}^2}\right) &= -\frac{\bar{\sigma}_{ij}}{\bar{\sigma}^2} \\
\bar{a}_{ij}\left(\frac{\bar{\sigma}^4 + (d-2)\bar{\sigma}_{ij}\bar{\sigma}^2 - \bar{\sigma}_{ij}^2(d-1)}{\bar{\sigma}^2}\right) &= -\frac{\bar{\sigma}_{ij}}{\bar{\sigma}^2} \\
\bar{a}_{ij} &= -\frac{\bar{\sigma}_{ij}}{\bar{\sigma}^4 + (d-2)\bar{\sigma}_{ij}\bar{\sigma}^2 - \bar{\sigma}_{ij}^2(d-1)}
\end{aligned}
$$

This, together with equation ($\star$) gives the final expression for $\bar{a}^2$.

□

Now we are ready to proof the variance minimizing property of the equal weights portfolio.

Proof. [**Theorem (2.2.2), part 1**]

From (2.1.1) we know that the global minimum variance portfolio is given by

$$
w_{\text{GMV}} = \frac{\Sigma_R^{-1}\mathbf{1}_d}{\mathbf{1}_d'\Sigma_R^{-1}\mathbf{1}_d}
$$

and its portfolio variance can be computed by

$$
\sigma_{\text{GMV}}^2 = \frac{1}{\mathbf{1}_d'\Sigma_R^{-1}\mathbf{1}_d}
$$

We will now prove that the variance of the equal weights portfolio is equal to the variance of the global minimum variance portfolio under fully uninformed market views. The variance of the equal weights portfolio is already known from Equation (2.1):

$$
\begin{aligned}
\sigma_{\text{EW}}^2 &= w_{\text{EW}}'\Sigma_R w_{\text{EW}} \\
&= \frac{1}{d}\bar{\sigma}^2 + \frac{d-1}{d}\bar{\sigma}_{ij} \\
&= \frac{\bar{\sigma}^2 + (d-1)\bar{\sigma}_{ij}}{d}
\end{aligned}
$$

Similarly, using the representation of the inverse covariance matrix for a fully uninformed market view given in Lemma (2.2.3), we get:

$$w_{\mathrm{EW}}' \Sigma_R^{-1} w_{\mathrm{EW}} = \frac{\bar{a}^2 + (d-1)\bar{a}_{ij}}{d}$$

Now the variance of the global minimum variance portfolio in a fully uninformed market view can be expressed in terms of $\bar{a}^2$ and $\bar{a}_{ij}$:

$$\begin{aligned}
\sigma_{\mathrm{GMV}}^2 &= \frac{1}{\mathbf{1}_d' \Sigma_R^{-1} \mathbf{1}_d} \\
&= \frac{1}{(d w_{\mathrm{EW}})' \Sigma_R^{-1} (d w_{\mathrm{EW}})} \\
&= \frac{1}{d^2 \frac{\bar{a}^2 + (d-1)\bar{a}_{ij}}{d}} \\
&= \frac{1}{d\left(\bar{a}^2 + (d-1)\bar{a}_{ij}\right)} \\
&= \frac{\frac{1}{(\bar{a}^2 + (d-1)\bar{a}_{ij})}}{d}
\end{aligned}$$

The remaining step to be done is to re-express the variance of the global minimum variance portfolio in terms of $\bar{\sigma}^2$ and $\bar{\sigma}_{ij}$, in order to verify equation

$$\bar{\sigma}^2 + (d-1)\bar{\sigma}_{ij} = \frac{1}{(\bar{a}^2 + (d-1)\bar{a}_{ij})}$$

This is shown in Appendix (6.0.2), so that altogether we get:

$$\begin{aligned}
\sigma_{\mathrm{GMV}}^2 &= \frac{\frac{1}{(\bar{a}^2 + (d-1)\bar{a}_{ij})}}{d} \\
&= \frac{\bar{\sigma}^2 + (d-1)\bar{\sigma}_{ij}}{d} \\
&= \sigma_{\mathrm{EW}}^2
\end{aligned}$$

□

To sum up, we now have several optimal asset allocation strategies, depending on the amount of information that we have for asset moments. If we really have no information at all, we will have to treat all assets as equal, and the optimal portfolio then is the equal weights portfolio with as many assets included as possible. Diversification pays off, and the portfolio variance will be below the individual asset variances. Next, if we have information about the covariance matrix only, without any knowledge about expected asset returns, then the optimal portfolio allocation is the global minimum variance portfolio. This way, lower risk assets will be preferred, and the portfolio will usually consist of several assets in order to also benefit from diversification effects. Theory suggests, however, that risk and return are related, so that low-risk assets generally also should be assets with lower expected returns. Given that this relation holds, the global minimum variance portfolio might not be the best way to bring knowledge of the covariance matrix to use anymore. Depending on the degree of risk aversion, higher risk assets might be preferable. This applies even more in a situation where expected asset returns are fully known. In such a case, Markowitz's portfolio theory will guide us to the optimal asset allocation, resulting in portfolios with optimal risk-return tradeoffs.

In reality, however, the situation that we face should generally be somewhere in-between. Clearly, we do not have perfect information about true asset moments, but they only can be estimated with uncertainty. Still, with data on historic price trajectories of assets we should at least get some idea about return distributions such that we could improve on naive equal-weights diversification. This is exactly what we will try to achieve in the following chapter: we want to design a portfolio selection approach that can explicitly be adjusted to reflect the actual level of uncertainty of asset moments. In the limiting case of perfect information about asset moments this selection approach should converge to Markowitz portfolio selection, while it should simply pick an equal weights portfolio for the other extreme case of full ignorance about asset moments. In other words, we want to both enjoy the benefits of diversification and still make informed bets on assets with supposedly good risk-return profiles.

2.3 Diversification-aware asset allocation

As we have seen in Section (1.1), there already exist a couple of strategies to derive optimal single-period asset allocations. Still, we will now develop an alternative approach. Thereby the main idea is that we want to find a strategy that produces well diversified portfolios. A well-known property of plain Markowitz is that it will generally result in highly concentrated portfolios, with only few assets that will be invested into. Even when uncertainty around asset moments is explicitly modeled and taken into account, diversification usually does not gain increasing importance in a plain Markowitz approach. When estimators are assumed to be unbiased, estimation uncertainty will basically only bloat up asset variances, but expected asset returns stay unaffected. Hence, portfolio decisions will not change dramatically. Also, higher diversification does not directly impact estimation uncertainty of portfolio moments. Standard errors are proportional to portfolio variance, but beyond that diversification itself does not play any role. Whether there are two or twenty assets in a portfolio does not change the precision of estimates, given that portfolio variances are the same.

In an environment of time-varying asset moments with huge uncertainty of the estimates, however, more diversified portfolios tend to perform better. For example, constrained Markowitz approaches enforce higher diversification by setting upper bounds on certain assets or groups of assets. Also, empirical performances for constrained Markowitz approaches generally tend to outperform plain Markowitz. In line with these findings, also the equal weights portfolio is commonly accepted as hard to beat portfolio strategy (DeMiguel, Garlappi, and Uppal 2007), which is just another indicator that portfolio diversification seems to pay off. In line with all these arguments, we now want to find a single-period asset allocation strategy that somehow favors diversified portfolios. Still, one point remains to be determined: how exactly is diversification defined and measured? We will thereby use the following definition.

Definition 2.3.1: ***[Portfolio diversification]***

Let P denote a portfolio with weights w, $\sum_{i=1}^{d} w_i = 1$. Then diversification of the portfolio is derived from the sum of squared deviations to the equal weights portfolio w_{EW}, and it is defined as:

$$\mathcal{D}^{||2||}(w) = 1 - \sqrt{\left(\sum_{i=1}^{d}\left(w_i - \frac{1}{d}\right)^2\right)}$$

The measure attains its maximum when all weights are equal, where its value is equal to 1:

$$\mathcal{D}^{||2||}(w_{EW}) = 1$$

Remark The measure of diversification can also be formulated in terms of the euclidean norm:

$$\mathcal{D}^{||2||}(w) = 1 - \sqrt{\left(\sum_{i=1}^{d}\left(w_i - \frac{1}{d}\right)^2\right)} = 1 - \|w - w_{\mathrm{EW}}\|_2$$

This way it could easily be generalized to use other norms as well:

$$\mathcal{D}^{||p||}(w) = 1 - \left(\sum_{i=1}^{d}\left|w_i - \frac{1}{d}\right|^p\right)^{\frac{1}{p}} = 1 - \|w - w_{\mathrm{EW}}\|_p$$

Depending on the context, we sometimes also might use different notations for the measure of portfolio diversification:

$$\mathcal{D}^{||2||}(w),\ \mathcal{D}(w),\ \mathcal{D}_w \text{ or } \mathcal{D}_P$$

Of course, one could easily think of alternative ways of measuring diversification or - if regarded from the opposite side - measuring portfolio concentration:

- number of assets in a portfolio
- number of assets with asset weight above some significant threshold
- Gini coefficient of portfolio weights

However, the definition of diversification that was chosen here has several properties that could make it a particularly good choice. First of all, the relation to the equal weights portfolio allows some justification to the view that diversification itself might be a desirable property. In the limit, maximum diversification is achieved with equal weights, which has proven to be a particularly well performing strategy in empirical applications. Also, as pointed out in Theorem (2.2.2), it is the best portfolio choice in a situation of fully uninformed market views. Hence, depending on how much trust we have into the estimation of asset moments, we could require higher levels of diversification, up to the case where we will end up with an equal weights portfolio. Furthermore, the measure also comes with desirable properties from a mathematical point of view. Derived from the euclidean norm, it can be shown to be a convex function, allowing it to preserve convexity in certain kinds of optimization problems. This way not only a global solution to the optimization problem can be guaranteed, but also fast and computationally efficient algorithms can be used.

Let's now define different approaches of how portfolio diversification could be taken into account in single-period asset allocation problems. The basic framework in such a single-period optimization is characterized by a couple of assumptions. First, we do have an estimate for expected return and volatility for all individual assets, as well as for the covariance matrix that describes their linear dependence. Second, transaction costs are assumed to be irrelevant for now, such that potential initial weights from former periods can be fully disregarded. Third, there are no other objectives besides maximization of expected out-of-sample utility and diversification. In particular, we do not have to track any benchmark or the like. And fourth, the single-period is the only period that influences expected utility. It is not part of a multi-period investment target, but can be considered separately. Historic performances, for example, do not need to be taken into account in any way, like it would be the case for example for portfolio insurance strategies.

In addition to this basic framework, all portfolios shall again fulfill the basic weight constraints of Definition (2.1.2). Hence, short-selling is forbidden.

Definition 2.3.2: ***[Minimum level of diversification]***

Let $\phi > 0$ be a penalty factor to express risk aversion, and $c_{MIN} \in [0, 1]$. Then the following constrained optimization problem will find an optimal portfolio for a given level of risk aversion ϕ within the set of portfolios that exceed a certain level of diversification c_{MIN}:

$$\begin{aligned} w_{[\mathcal{D}]}(\phi, c_{MIN}) = \arg\max_{w} & \quad \mu_P - \phi\sigma_P^2 \\ \textit{subject to} & \quad \mathcal{D}_w \geq c_{MIN} \end{aligned}$$

Hence, with varying levels of risk-aversion ϕ we will get a diversification-aware efficient frontier: all (μ_P, σ_P) optimal portfolios that also fulfill some minimum level of diversification. Sub-index $[\mathcal{D}]$ denotes the constraint, which in this case is a constraint on portfolio diversification $\mathcal{D}$.

Remark Optimization problem (2.3.2) can be reformulated into a convex optimization problem (Cornuejols and Tütüncü 2006). Both objective function and constraints therefore need to be written in terms of convex functions:

$$\begin{aligned} w_{[\mathcal{D}]}(\phi, c_{\text{MIN}}) = \arg\max_{w} & \quad w'\mu - \phi w'\Sigma_R w \\ \text{subject to} & \quad \|w - w_{\text{EW}}\|_2 \leq 1 - c_{\text{MIN}} \end{aligned}$$

With the $w_{[\mathcal{D}]}$ optimization one can easily fix a desired level of diversification. Given this constraint, we get a set of μ-σ optimal portfolios similar to the traditional Markowitz efficient frontier. While the constraint fully determines the diversification-aware frontier, we need parameter ϕ to finally determine a unique portfolio out of this set. ϕ basically determines the optimal tradeoff between risk and return. Or, in other words: it determines how far we walk along the frontier to the right. ϕ generally should be determined by the investors risk-aversion. Holding it fix, the risk that we ultimately get depends on how risk is compensated in the current financial environment. In times of high risk premia, when expected returns are comparatively large for the risks that are prevailing in the financial system, investors will end up in comparatively risky portfolios. In contrast, the same value of ϕ will lead to portfolios with comparatively lower volatility in times of low risk premia. While this automatic adaption mechanism may very well be advantageous, it is not really obvious how a correct value for ϕ should be chosen. Mapping investors' risk aversion to adequate values of ϕ is far from being obvious.

In many situations it would be desirable to select individual portfolios of the frontier through a parameter that is given on a more meaningful scale. For example, using $w_{[\mathcal{D}]}$ optimization to determine a ramp of σ targets, individual points on the frontier ideally could be referenced by portfolio volatility directly. This problem, however, can be solved numerically. As portfolio volatility σ_P is decreasing in ϕ, the diversification-aware frontier can be evaluated on a grid of values $(\phi_i)_i$, with further fine-tuning of the values through bisection methods. This way we get a numeric solution to the following optimization problem.

Definition 2.3.3: ***[Optimality with σ and diversification targets]***

Let $\sigma^\star$ denote a target portfolio volatility, and c_{MIN} denote a desired minimum level of diversification. Then we define the optimal portfolio with given volatility and diversification targets by the following optimization problem:

$$\begin{aligned} w_{[(\mathcal{D},\sigma)]}(c_{MIN}, \sigma^\star) = \arg\max_{w} \quad & \mu_P \\ \textit{subject to} \quad & \mathcal{D}_w \geq c_{MIN} \\ & \sigma_P = \sigma^\star \end{aligned}$$

Both optimization problems so far are built in a way that directly allows the specification of a minimum level of diversification. But how exactly should this level be chosen?

The optimal level of diversification might be changing with regards to the financial environment. When individual assets show exceptionally good risk-return profiles, it might be worthwhile to forgo some portfolio diversification for the sake of more aggressively exploiting existing investment opportunities. The following different approach to diversification-aware optimization tries to treat this tradeoff between expected returns and diversification more explicitly.

Definition 2.3.4: ***[diversification-aware σ-targeting]***

Let $\phi > 0$ be a scaling factor, and $\sigma^\star$ be a given level of volatility. Then the following optimization *problem will find a portfolio with optimal tradeoff between expected return and diversification among all portfolios that fulfill a given $\sigma^\star$ target.*

$$\begin{aligned} w_{[\sigma]}(\phi, \sigma^\star) = \arg \max_{w} \quad & \mu_P + \phi \mathcal{D}_w \\ \textit{subject to} \quad & \sigma_P = \sigma^\star \end{aligned}$$

The value of ϕ hence can be interpreted as tuning parameter to determine how much portfolios with higher diversification levels will be favored. Higher levels of ϕ will lead to higher levels of diversification, but at the cost of lower expected portfolio returns. Again, sub-index $[\sigma]$ represents the portfolio constraint, which in this case is a constraint on portfolio volatility.

Remark Optimization problem (2.3.4) can also be reformulated as convex optimization problem (Cornuejols and Tütüncü 2006). Writing it in terms of convex functions and constraints, we get:

$$\begin{aligned} w_{[\sigma]}(\phi, \sigma^\star) = \arg \max_{w} \quad & w'\mu + \phi \left(1 - \|w - w_{\mathrm{EW}}\|_2\right) \\ \text{subject to} \quad & \sqrt{w'\Sigma_R w} = \sigma^\star \end{aligned}$$

Now that we have a suitable set of diversification-aware portfolio selection strategies at our hand, we will combine the individual parts in order to formulate a robust strategy that can deal with any σ target and any required level of diversification. For example, depending on the choice of $\sigma^\star$, optimization problem $w_{[\sigma]}$ might become infeasible. If $\sigma^\star$ is chosen below the variance of the global minimum variance portfolio, for example, no portfolio exists that could fulfill the optimization constraints. A similar problem also could apply to pairs of risk and diversification levels $(\sigma^\star, c_{\mathrm{MIN}})$. Formulating a robust strategy requires some knowledge about extreme portfolios that are edge points of the range of feasible solutions.

Definition 2.3.5: ***[Edge case portfolios]***

For any given level of portfolio volatility $\sigma^\star$, we define the maximally diversified portfolio as the solution to the following optimization problem:

$$\begin{aligned} w_{[\max(\mathcal{D})|\sigma=\sigma^\star]} := \lim_{\phi\to\infty} w_{[\sigma]}(\phi, \sigma^\star) = \arg \max_{w} \quad & \mathcal{D}_w \\ \textit{subject to} \quad & \sigma_P = \sigma^\star \end{aligned}$$

Furthermore, for any given level of portfolio diversification c_{MIN}, we define the edge points of the associated diversification-aware portfolio frontier as solutions to the following optimization problems:

$$\begin{aligned} w_{[\min(\sigma)|\mathcal{D}=c_{MIN}]} := \lim_{\phi\to\infty} w_{[\mathcal{D}]}(\phi, c_{MIN}) = \arg \min_{w} \quad & \sigma_P^2 \\ \textit{subject to} \quad & \mathcal{D}_w \geq c_{MIN} \end{aligned}$$

$$w_{[\max(\sigma)]|\mathcal{D}=c_{MIN}]} := w_{[\mathcal{D}]}(0, c_{MIN}) = \arg\max_{w} \quad \mu_P \quad \text{subject to} \quad \mathcal{D}_w \geq c_{MIN}$$

Given these edge case portfolios, we now can define a robust diversification-aware strategy. The strategy will exploit different optimization problems for different ranges of σ.

Definition 2.3.6: ***[diversification-aware σ ramp]***

Let $(\sigma_i^\star)_i$ be a set of portfolio volatility targets, and c_{MIN} be a level of minimum diversification. Then for each target $\sigma_i^\star$ we select the following portfolio weights:

$$w_i^\star = \begin{cases} w_{GMV}, & \text{if} \quad \sigma_i^\star \leq \sigma_{GMV} \\ w_{[\max(\sigma)|\mathcal{D}=c_{MIN}]}, & \text{if} \quad \sigma^\star \geq \sigma\big(w_{[\max(\sigma)|\mathcal{D}=c_{MIN}]}\big) \\ w_{[(\mathcal{D},\sigma)]}(c_{MIN}, \sigma^\star), & \text{if} \quad \sigma\big(w_{[\min(\sigma)|\mathcal{D}=c_{MIN}]}\big) < \sigma^\star < \sigma\big(w_{[\max(\sigma)|\mathcal{D}=c_{MIN}]}\big) \\ w_{[\max(\mathcal{D})|\sigma=\sigma^\star]}, & \text{if} \quad \sigma_{GMV} < \sigma^\star \leq \sigma\big(w_{[\min(\sigma)|\mathcal{D}=c_{MIN}]}\big) \end{cases}$$

Although the definition looks rather involved at first sight, the idea is quite simple. At the core of the strategy we still just enforce the concept that we want to promote diversification. However, when the diversification target is formulated too ambitiously, such that no feasible solution exists, we need to define fallback portfolios. So whenever it is possible, we will take a portfolio on the diversification-aware frontier. If the required level of diversification can not be achieved for a given σ target, we will settle for the portfolio that at least maximizes diversification. In addition, we now also need to define fallback portfolios for the case when the σ target itself is chosen out of range. The lowest possible portfolio volatility that can be achieved is given by the global minimum variance portfolio. So if the σ target falls below this threshold in some given market environment, we will simply choose the global minimum variance portfolio instead. On the opposite side, when the σ target is chosen too high, we will use the right endpoint of the diversification-aware frontier as fallback. If we were further following the contour line with given level of diversification, we would end up with a Pareto-dominated portfolio: the same level of diversification and expected return also could be achieved with less risk. Alternatively, we could also allow the fallback strategy to settle for portfolios with less diversification, such that we effectively would pick portfolios along the borderline of the $(\mu, \sigma, \mathcal{D})$ Pareto-optimal area. However, this way we might end up in highly concentrated portfolios with high volatility, which basically amounts to putting large bets on a few unpredictable assets. This does not seem to be a good idea in order to maximize out-of-sample utility, so that the right endpoint of the diversification-aware frontier will mark the maximum volatility that we allow. Overall, Definition (2.3.6) just extends the diversification-aware frontier such that we also get a solution for σ targets outside of its range.

Let's now examine a couple of aspects of diversification-aware portfolio selection in more details. Therefore, we will use the same example data as in Section (2.1.2) for illustrative purposes. First of all, in contrast to traditional μ-σ portfolio selection we now evaluate portfolios according to three properties: expected return, variance and diversification. Hence, each portfolio needs to be represented by a three-dimensional point, such that the traditional μ-σ space needs to be extended with an additional dimension. In general, each feasible (μ, σ) combination could be achieved with multiple portfolios. As portfolio diversification now also is a desirable property, let's only look at those portfolios that maximize the level of diversification for each (μ, σ) combination. This is a necessary condition for a portfolio to be Pareto-optimal, because otherwise we could find a portfolio with equal risk-return profile but higher diversification. When looking at portfolios with maximum diversification only, each feasible (μ, σ) combination can be mapped to a diversification value. This way, we get a maximum diversification surface above the μ-σ domain. Due to the convexity property of diversification measure $\mathcal{D}$, contour lines of this surface will enclose convex sets. In other words, the set of (μ, σ) combinations where the maximum achievable diversification level exceeds a certain threshold is convex. From these contour lines, in turn, only the upper-left part entails Pareto-optimal portfolios. All other portfolios that lie on the same contour line will have equal diversification level, but sub-optimal risk-return properties.

Figure (2.11) illustrates how the traditional μ-σ perspective can be extended with diversification properties. Individual asset moments are shown as colored dots with text labels, and the μ-σ efficient frontier from Markowitz portfolio selection with short-selling constraints is shown as thick blue line. In addition, upper-left sections of maximum diversification contour lines are shown for selected levels of diversification as black lines. These sections of the contour lines correspond to the diversification-aware frontiers that are defined by the optimization problem of Definition (2.3.2). The higher the required level of diversification, the closer the frontiers will get to the equal-weights portfolio shown as blue circle. Furthermore, the red line shows the maximum level of diversification that is achievable for each σ target between the global minimum variance portfolio and the equal weights portfolio. Given these components, the diversification-aware strategy defined in Definition (2.3.6) can easily be described. It consists of all portfolios on the respective diversification-aware frontier, extended to the left by portfolios on the red maximum diversification line.

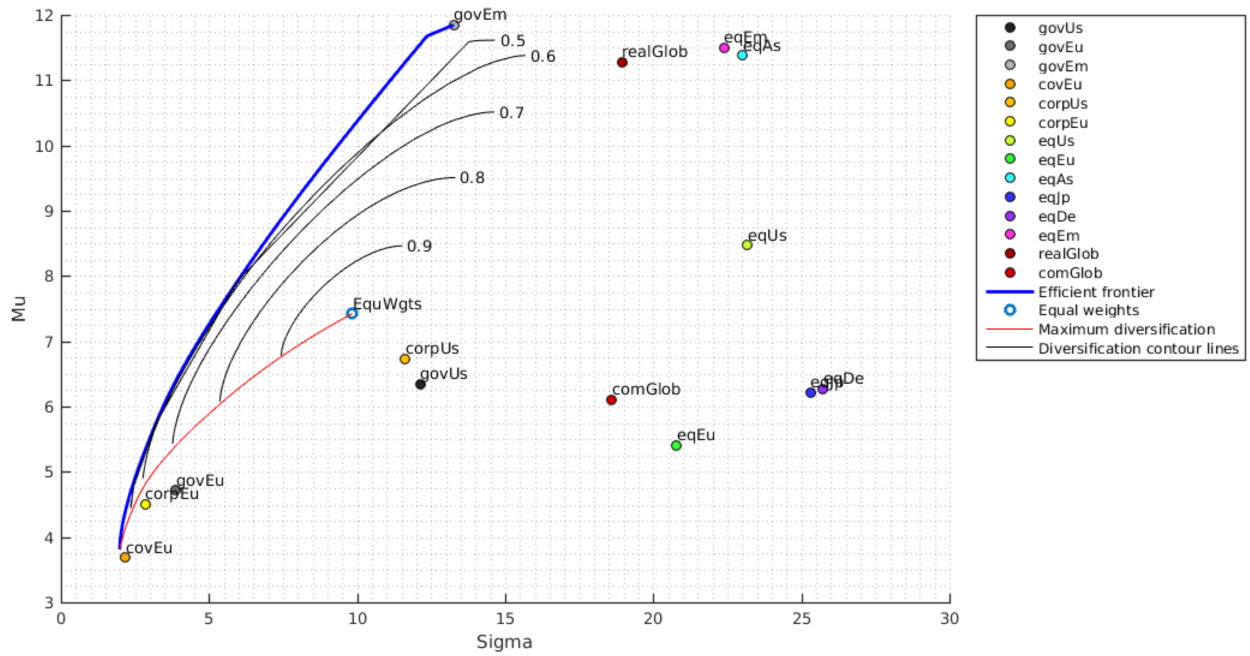

Figure 2.11: Efficient frontier together with diversification-aware frontiers

While in the traditional μ-σ setting the set of Pareto-optimal portfolios forms a single line, the efficient frontier, the set now forms a surface in three-dimensional space, see Figure (2.12). Thereby the orange line depicts the traditional μ-σ efficient frontier, however with additional illustration of the level of diversification of individual portfolios. As one can see, portfolios are generally below a diversification level of 50%, and they get even more concentrated at both endpoints of the frontier, where diversification levels get close to zero. This is just a quantitative measure of what we have found in Figure (2.9) already: efficient frontier portfolios consist of very few assets only. For the high-risk endpoint of the efficient frontier the portfolio even consists of a single asset only, while the global minimum variance portfolio at least has some fraction of wealth invested into assets other than the minimum variance asset. The red line again shows the maximum diversification line for each σ target, which also is a boundary of the Pareto-optimal set of portfolios. The Pareto-optimal surface encapsulated between efficient frontier and maximum diversification line is shown in gray, with some artifacts at the fringes due to rendering problems of the numerical software that was used. In reality, the gray surface would fully extend to the orange and red boundary lines. However, visualization requires the surface to be defined on an equidistant grid, causing numerical issues in steep areas of the surface.

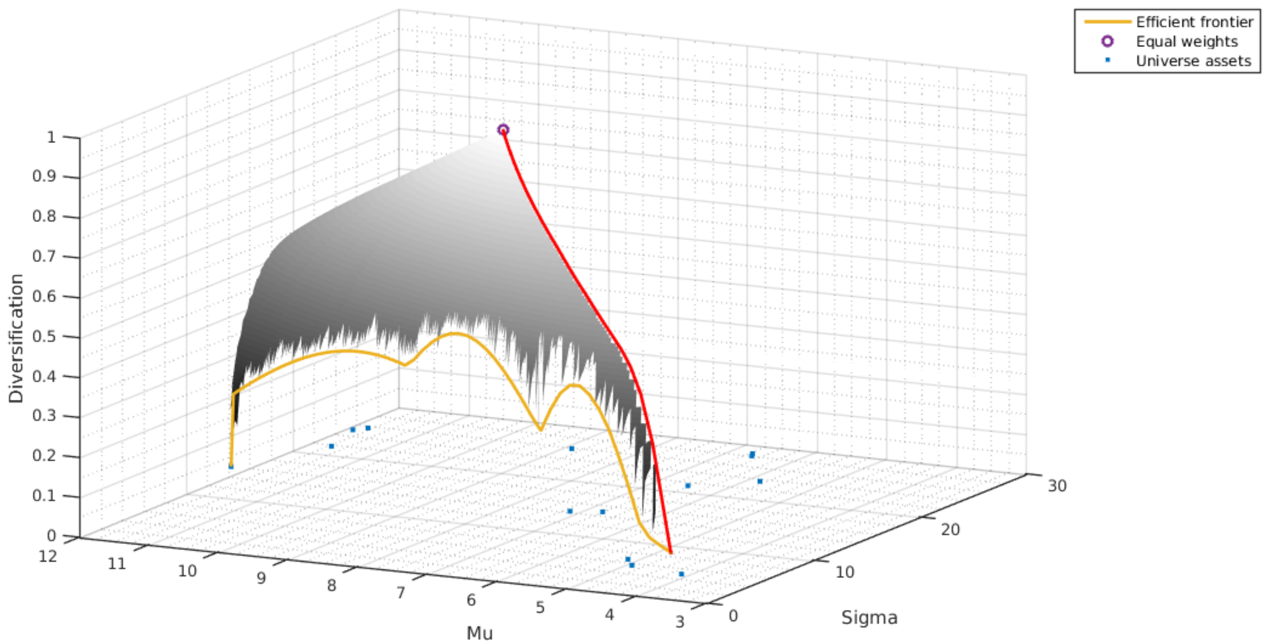

Figure 2.12: Surface of $(\mu, \sigma, \mathcal{D})$ Pareto-optimal portfolios

When the mapping from (μ, σ) to the diversification maximizing portfolio is unique, we can further extend this mapping into a mapping from (μ, σ) to other portfolio properties. For example, we can associate each (μ, σ) combination with the portfolio weight of a particular asset. This way we can see how much diversification we need to require such that assets with poor risk-return profile will become part of the portfolio choice. And more generally, we can see for which risk-return regions individual assets will be important. Figure (2.13) shows for each individual asset the asset weight in portfolios on the Pareto-optimal surface. In each subplot we can see all individual asset moments depicted by blue dots, with the moments of the currently chosen asset depicted with a red asterisk. Furthermore, the $(\mu, \sigma, \mathcal{D})$ Pareto-optimal region is colorized according to the fraction of weight that the asset has within the portfolio. Dark blue colors represent weights close to zero, and dark red colors weights equal to or above 20%. The full color scale can be seen in the subplot in the lower right corner.

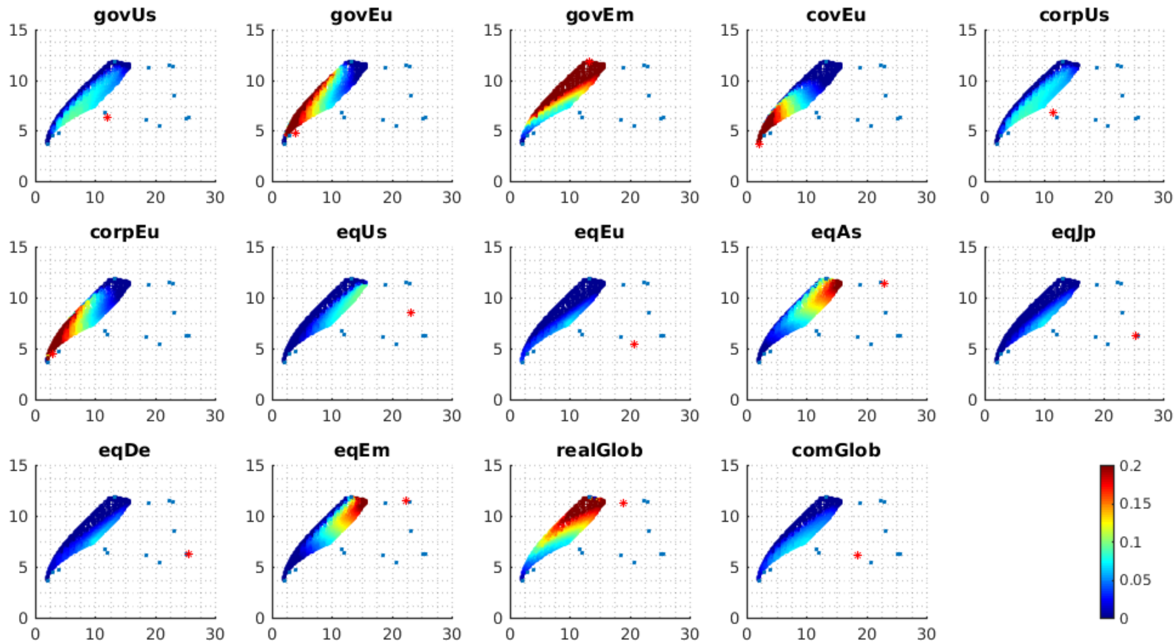

Figure 2.13: Individual asset weights for different portfolio risk-return regions

This is only one possible example that shows how a unique Pareto-optimal surface could be used to gain additional insights, but similar applications could be thought of. Instead of relating risk-return regions to asset weights, it could be more meaningful to directly relate different regions to risk factors. For example, one could use this approach to gain insights into portfolio exposure with regards to market beta, interest rate sensitivities or foreign exchange rate exposure. This way one can analyze how relaxation of portfolio diversification requirements could lead to more aggressive investments into certain market sectors that also open the door to increased exposure with regards to certain risks.

One of the original intentions of a diversification-aware approach was to construct portfolios that more strongly exploit the benefits that can be achieved from diversification. In other words, having a portfolio with multiple assets in itself is not yet a desirable thing, but only becomes useful when individual assets do not have perfect correlations between each other. Only then the overall portfolio volatility will be lower than the sum of the individual asset volatilities. In order to roughly evaluate whether this is actually achieved with the proposed diversification-aware approach, we will now define a measure of diversification benefits. The idea is very simple: for each portfolio we calculate the true portfolio volatility and then compare it to the portfolio volatility that one would get if all correlations were equal to one. The difference between both volatilities will be a measure for the benefit achieved by diversification. For the case of perfect linear correlations the portfolio volatility is equal to the weighted sum of the individual asset volatilities, so that we get for the measure of diversification benefit:

$$\sum_{i=1}^{d} \sigma_i w_i - \sqrt{w' \Sigma_R w}$$

This measure of diversification benefit is visualized in Figure (2.14) as a function of diversification and portfolio volatility. Individual lines represent different levels of diversification-aware frontiers, portfolio volatility is shown on the x axis, and diversification benefits on the y axis. As can be seen, higher levels of diversification coincide with higher levels of diversification benefits, such that diversification-aware portfolio selection really seems to benefit in terms of inherently reduced variances. This effect also seems to be more pronounced for higher-risk portfolios. Not unexpectedly, as there also is a larger amount of fluctuations in the first place that potentially could be offset.

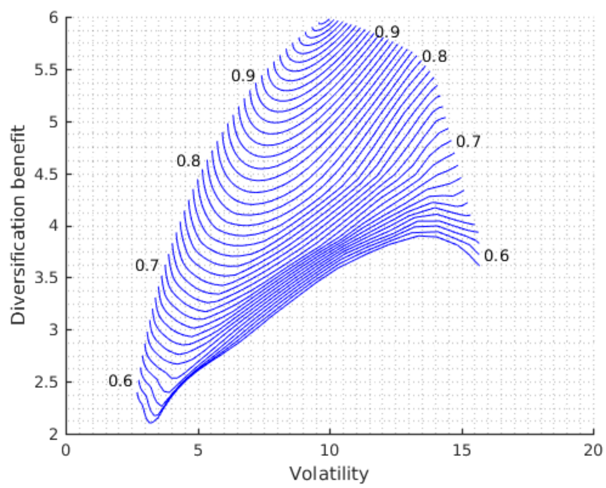

Figure 2.14: Diversification benefits as a function of diversification and portfolio volatility

As already mentioned above, the diversification-aware portfolio strategy defined here will ultimately achieve higher levels of diversification by bringing Markowitz portfolios closer to the equal weights portfolio. This same idea also could be achieved by some other approaches. The first such approach does simply shrink Markowitz portfolios directly to the equal weights portfolio. The final optimal portfolio then will be given as convex combination of a Markowitz efficient portfolio and the equal weights portfolio:

$$w^{\star} = \alpha w_{\mathrm{EW}} + (1-\alpha) w_{\mathrm{Markowitz}}, \quad \alpha \in [0, 1]$$

This way we somewhat loose control over the final portfolio volatility that we get, as we only can define a volatility target for the original Markowitz portfolio that is used. A simple workaround for this, however, could be to convex-combine a full efficient frontier with the equal weights portfolio, and then just pick a portfolio that matches the desired volatility target on the modified version of the efficient frontier. A more serious drawback to this approach is that it does not allow to fully circumvent individual assets that temporarily have very bad risk-return profiles. Convex combination with equal weights will automatically include all assets into the final portfolio. This is where our diversification-aware approach could be superior, as it still allows to leave out such assets and achieve higher diversification through the rest of the assets. Only when the required diversification level is excessively high we can not avoid unprofitable assets anymore.

A second approach that might lead to similar results is by shrinking estimated asset moments themselves to some default case where assets are indistinguishable. For indistinguishable assets the best portfolio strategy would be the equal weights portfolio, such that it again emerges as a natural limit of the approach. However, volatility targeting is hard to enforce, as the optimization itself will only work with modified and distorted moments. It is not obvious how this could be made compatible with a target for real and unbiased portfolio volatility.

Chapter 3

Dynamic risk management

In the previous chapter we have dealt with optimal portfolio choices from a single-period perspective: if we had to make a single investment decision, what would be the optimal one. In reality, however, we usually face a rather long investment horizon as financial investor, and it is unlikely that the financial landscape will remain unchanged over such long horizons. Furthermore, such changes of the financial environment are almost impossible to predict upfront. Hence, one might not want to stick to an initial investment decision infinitely, but want to adjust the original portfolio at later points in time. Portfolio selection hence quite naturally becomes a multi-period concept, and this is what we will address in the present chapter.

3.1 Objectives

First and foremost, the objective of effective asset management should be to achieve high risk-adjusted returns. Performance alone is not a sufficient evaluation criteria, but returns always need to be judged based on the risk that was taken. Our target hence is finding an asset management strategy that has a good risk-return profile. As trading costs will negatively affect final payoffs, we will also try to keep trading activity within some reasonable bounds in order to end up with good net returns. Net returns thereby refers to returns after subtraction of transaction costs.

In addition to these rather indisputable objectives, we also want to fufill a couple of “soft” goals that arise when asset management itself shall be offered as a product. First, investors generally differ with regards to investment horizon and risk-aversion, and hence any single “one-size-fits-all” asset management strategy is unlikely to appeal a majority of investors. Hence, the strategy should come with a tuning parameter that reflects the actual level of risk that is desired. This way, a large spectrum of risk categories can be covered, such that a custom-tailored solution can be provided to investors with different levels of risk-aversion. In other words: we do not search a single investment strategy, but a broad set of related strategies that cover a significant spectrum in the dimension of risk. We will also speak of a "ramp of risk" in this context, refering to a spectrum of strategies. And second, if an asset management strategy is exercised for multiple persons, then it also should lead to similar results in all cases. An easy way to achieve this would be to implement the strategy by means of an investment fund, where shares for a single joint portfolio are issued and sold to clients. That way clients with perfectly equal holding periods will be guaranteed to also get the same performance. However, such a solution will prohibit full client-tailored asset management strategies that try to make use of client specific characteristics in an optimal way. For example, active tax loss harvesting will require each client portfolio to be optimized individually, as payable taxes from asset price gains depend on the client's specific tax situation. How taxable gains translate into taxes, for example, depends on tax allowances and effective tax rates. If such additional wealth management strategies are to be put in place as well, a fund solution will not be possible. Hence, clients within each risk category have to be optimized individually, possibly opening the door for undesired performance dispersion within categories. A suitable asset management strategy should try to keep this within reasonable bounds.

Now that the broader context is laid out, let's further pin down some desired properties that individual category strategies should have from a risk-return perspective. Wide parts of the financial econometrics literature revolve around the issue of time-varying asset moments, with most research showing evidence of moments changing over time. One of the first appearances of time-varying volatility was (Mandelbrot 1963), who found that "large changes tend to be followed by large changes - of either sign - and small changes tend to be followed by small changes". Meanwhile such variation of levels of volatility is commonly accepted in the literature, and picked up as "volatility clusters" as so-called stylized fact of financial data. Somewhat less evident is the existence of time-varying expected returns. This mainly has to do with the fact that estimation of first moments of returns is tremendously more difficult, due to a low signal-to-noise ratio of asset returns: compared to the size of expected returns, variation around expected values is huge. Still, at least for long-term horizons the evidence on predictable time-varying expected returns seems to be convincing (Fama and French 1988).

So how do time-varying asset moments affect asset management? With changing asset moments, also optimal portfolio weights will change over time. Or, put differently: holding asset weights fixed will most likely result in sub-optimal portfolio choices over time. Hence, one ideally would like to react to changing market environments, holding different portfolios at different points in time. So even if a certain asset is not a really promising investment on average, there could be sub-periods where it is worthwhile to be held. This would require, however, that we are able to identify good and bad periods for individual assets, which is a rather ambitious and probably even unrealistic goal. Asset moments are generally difficult to estimate already in the long run, and the time-varying nature of moments either requires an ability to predict changes of risk-return characteristics, or to at least detect them in a timely manner. This in turn usually requires to put large weight on recent observations in the estimation process, thereby exposing moment estimates to large estimation risk. There are ways of making use of contemporary information less aggressively, however, and this is what we will try to do with our strategies. Time-variation of asset moments automatically passes on to portfolio moments, given that portfolio weights are held fixed. In particular, this would mean that portfolio returns will be exposed to volatility clusters as well, such that a static portfolio would be subject to different levels of risk over time. Risk is largely related to volatility, which can be estimated adequately and timely with econometric models like GARCH (Bollerslev 1986). This information then can be used to rebalance portfolios in such a way that risks are stabilized over time: if portfolio volatility increases above a certain level, wealth is shifted into assets with more appropriate risk characteristics. Given that investors' risk-aversion stays rather constant over time, portfolio risk could more closely align with investors' preferences that way. The more aggressive way of finding assets with temporarily nice risk-return profiles is usually referred to as active management, with primary focus on the exploitation of expected return predictions. Opposed to that, the approach of dynamically rebalancing portfolios in order to stabilize risks will be referred to as dynamic risk management from now on. In effect, the asset allocation strategies applied in the following chapters will be somewhere in between. Although stabilization of risk is the dominant feature, at the same time we also will try to slightly benefit from time-varying expected return estimates as well.

Steady portfolio adaptions to changing market environments will require trading of securities, thereby incurring trading costs. Depending on the amount of rebalancing these trading costs can quickly become non-negligible, affecting net investment returns. Hence, these costs need to be taken into account, making net returns the primary measure of investment success.

In addition to that, dynamic risk management also requires another fundamental change to the optimization problem. Without any rebalancing, the asset allocation problem can be formulated as single-period investment choice with holding period equal to the investment horizon. In other words: if we want to invest money for the next ten years, then the only thing we should be interest in are the ten-year portfolio return moments. With a strict investment maturity, the path of portfolio values itself is not relevant at all, but the only thing that matters is the final return at maturity. When managing money for multiple investors, this single strict maturity is not existent anymore, as the investment strategy needs to comply with all investment horizons of all investors simultaneously. This effectively will render the single-period portfolio optimization problem inadequate, and the trajectory of the portfolio performance now matters at each point in time. Both constant rebalancing and multiple investors will require to consider portfolio optimization as a multi-period challenge.

In principle, an optimal multi-period strategy would also require a multi-period aware optimization. In other words: a multi-period setting is more than just an iterated single-period setting. An easy example that can clarify this point is as follows. Let's assume that there is a given annual risk target. And, given this amount of risk that may be taken, we want to maximize expected returns. The first way to achieve this is by selecting portfolios in each sub-period such that each portfolio will load on risks proportionately to the overall risk. A portfolio that is held for one week hence should take risks that amount to $\frac{1}{52}$ of the annual target risk. The overall expected return of this strategy might be exceeded by a more flexible approach, however. Let's now imagine that we flexibly increase exposure in those periods, where risk is rewarded above average, while cutting back on risks whenever risk is rewarded comparatively less. Both strategies could be calibrated in such a way that they have the same risk on an annual basis. Still, the additional leverage that the more flexible strategy provides in good times should boost expected returns. Although this is only a quite sloppy argumentation, the example still intuitively points out the difference between an iterated single-period optimization and a full multi-period optimal strategy.

Summing up, the main targets that we want to achieve are the following:

- designing a set of strategies with increasing risk profiles
- stabilizing risk over time for individual portfolios
- conducting asset management on a client level without introducing portfolio dispersion within categories
- taking transaction costs into account

3.2 Multi-step asset management approach

In this chapter we will now try to find a solution to the problem of optimally managing assets over multiple time periods. But how exactly is an optimal strategy described, and is it realistic to find one? Quite generally, an optimal asset management strategy should be derived as the solution to some optimization problem, where the objective function of the optimization usually is given by some function of expected utility, optionally with additional constraints in place. But how should all of the objectives that are laid out in Section (3.1) be translated into a single optimization problem?

In order to see, how far we could get, let's first start and inspect a reduced variant of the full set of objectives and think about the maximization of expected net returns for some given target of risk. Net returns, again, refer to returns after subtraction of transaction costs. A major challenge of this optimization problem is that individual sub-period portfolio decisions are highly interrelated. Hence, the overall optimization problem can not be split into smaller sub-period optimizations. As already mentioned before, there might be better ways of matching an annual risk target than simply spreading the risk taken on sub-periods equally. Increasing exposure in periods where risk is rewarded above average while reducing risk in all other periods could improve the overall risk-return profile. Hence, in the true optimal allocation strategy individual sub-periods become interrelated: the sub-period risk target depends on what is chosen as risk target in other sub-periods. This problem, however, gets eliminated by one of the additional strategy objectives: keeping risks stable over time will break the existing interrelation that is caused by a long-term risk target.

There is, however, another aspect that will prevent the split-up of the overall optimization problem: transaction costs. In the presence of transaction costs, any rebalancing decision becomes a tradeoff between costs and benefits. While benefits are given by holding an improved portfolio, costs are determined by the transaction costs that are required to get from some initial portfolio to the desired portfolio. The problem with time-varying asset moments is, however, that while today's trading might bring you closer to some optimal portfolio this period, it might simultaneously bring you further away from an optimal portfolio some periods ahead. Either one then rebalances again, or one has to live with an inferior investment in the future. Bottom line is: today's portfolio decision will either have a long-lasting effect on the performance of the investment, or it will require additional trading costs in the future. This way individual investment sub-periods become interrelated again, and the multi-period optimization problem does not reduce to an iteration of single-period optimizations. The global optimal solution would need to be derived for all periods simultaneously. Therefore, a single global optimization problem would require an unreasonably large computational complexity. To keep the computational burden in bounds, we will split the global optimization into several smaller optimization steps. Doing this, of course, comes at the cost of not getting a globally optimal solution.

We will now split the optimization problem into two major components. First, we consider the problem of finding an optimal solution in a world without any transaction costs. And second, we will take the portfolio weights of the first step as target weights, and we will try to realize them as cost-efficiently as possible.

The part of finding optimal portfolio weights with transaction costs set to zero will reduce to an iterated single-period optimization in our case. By explicitly requiring the same risk target in each period, the globally optimal strategy will be the maximization of expected returns in each individual period. Hence, each period can be optimized for itself. We will call such a single-period optimization problem without consideration of any transaction costs a single-period asset allocation problem. One famous example for such a single-period asset allocation strategy is a straightforward application of Markowitz portfolio selection. An alternative would be the diversification-aware portfolio selection introduced in Section (2.3).

Let's now get to the part of cost-efficiently replicating a given series of single-period optimal assets. Thereby, we need to balance between two different objectives: first, matching the portfolio targets as closely as possible, and second, trading as less as possible in order to keep transaction costs limited. While trading costs are easily determined, measuring the gain of trading is far less obvious. For a single period, how much do we gain by trading to the optimal portfolio? There are a couple of ways how one could define such a measure, but for the moment we will defer this problem to a later chapter. So let's now assume that we already have an appropriate measure that we can use to quantify how much better the optimal portfolio is. Still, there is another crucial problem: we do not know yet how long this benefit will last. With changing market situations, a given portfolio might be the better choice now, but after some time the situation easily could be reversed. So it is not sufficient to determine the instant single-period benefit of a better portfolio, but we also need to take into account the persistence of the situation. For example, completely converting a stock portfolio into a bond portfolio is not worthwhile when we can already suspect that the bond portfolio will be better for only a very short period. We would just have to redo the trades again, and the costs of changing the portfolio twice might dwarf the temporary benefits. Summing up, evaluating the gains of trading can be decomposed into two different parts. Instant gains that measure the immediate gains that will arise in the next investment period. And persistence of gains, measuring how long-lasting the effect will be.

To conclude, instead of directly striving for a globally optimal dynamic asset management strategy, we will separate the problem into three distinct parts:

- finding optimal single-period portfolios without consideration of trading costs
- adjusting optimal portfolios in order to account for persistence
- substituting target portfolios with cost-efficient similar alternatives

The first part is related to the single-period perspective that was treated in Chapter (2). The other two parts, however, we will both treat in the next sections as part of a common framework that we will refer to as weight filters. Thereby the underlying idea is that we get a series of optimal target portfolios from single-period optimizations, which then will be iteratively adapted through weight filters. In a world without trading costs, we would fully realize single-period target portfolios in each period. With trading costs, however, we will use weight filters in order to modify original target portfolio weights such that they become cost-efficient but still sufficiently similar alternatives to the original targets. Or, as extreme case, target portfolios are modified so much that they are set equal to current portfolio weights, such that no trading will occur at all.

3.3 Target portfolio weight filters

First, we will now look at weight filters that deal with part two of the above mentioned challenges: the adjustment of target weights in order to reflect the endurance of possible investment opportunities. To correctly adapt for endurance, one ideally needs to have a good guess about the future evolution of the financial environment. Only then it is possible to reliably assess how long a certain investment opportunity might persist. However, asset moments are already hard to estimate contemporaneously, let alone to estimate how they will change in the future. Inferring endurance from the informational granularity of individual asset moments hence could be an overly ambitious target.

A different approach might be to rely on measures that emerge from historic portfolio positions. For example, historic fluctuations of individual asset weights might be an indicator of how persistent certain investment opportunities should be in the future. An asset that was an optimal investment opportunity every now and then in the past, but each time remained optimal for a short period only, might also not be a beneficial investment for longer time horizons now. Similar to that, some assets might only prove to be good investment choices under particular market circumstances. For example, when global stock markets experience simultaneous downturns due to worldwide financial distress, bonds might temporarily become a comparatively attractive investment. However, as soon as markets start to recover, equities will become the asset of choice again. Depending on how long such downturns have persisted in the past one could decide whether a temporary investment into bonds is worth the costs of trading an equity portfolio into a bond portfolio and back again.

Similarly to historic asset weight fluctuations one could also think of other measures that relate to historic asset positions. For example, historic contributions of individual assets might be a proxy for how reliable current investment opportunities really are. An asset that frequently did appear to be a good investment opportunity in the past, but in the end did generate almost no profits nevertheless, might indicate one of two problems. Either these investment opportunities exist, but they only persist for such a short time that we are not able to exploit them, or they simply do not exist at all and we are just not good at identifying truly good opportunities.

However, we will use an even cruder proxy to identify persistence of investment opportunities: we only invest in assets after they have appeared to be a good investment opportunity for a certain amount of time. In other words, just because an asset is part of the single-period portfolio target today, we do not want to invest into it directly. Only when the asset is part of the target weights for a certain time span we will trust it to be a good investment choice. In some way we hence use past persistence as a proxy for future persistence.

A rather simple approach to implement this logic is through moving averages of target weights. This way, assets that have been part of single-period target portfolios for longer periods tend to get larger shares of the portfolio than those that only made it into single-period targets very recently.

Definition 3.3.1: *[**Moving average target weight filter**]*

Let $(w_{t,A})_t$ denote the single-period portfolio target weights in period t. Then the moving average target weight filter of window size n is defined by the following weight transformation function h:

$$h^{MA}(n) := h(w_{0,A}, \ldots, w_{T,A}, n) = \frac{1}{n}\sum_{i=0}^{n-1} w_{T-i,A}$$

A very similar approach to correct target weights for persistence is through exponential weighting of portfolio weights. This way, more recent portfolio targets get higher weight in the calculation of the average, and hence they have more impact on final portfolio weights.

Definition 3.3.2: *[**EWMA target weight filter**, (Würtz et al. 2009)]*

Let $(w_{t,A})_t$ denote the single-period portfolio target weights in period t. Then the exponentially weighted moving average target weight filter with decay factor λ is defined by the following weight transformation function h:

$$\begin{aligned} h^{EWMA}(\lambda) := h(w_{0,A}, \ldots, w_{T,A}, \lambda) &= \sum_{i=0}^{T} \bar{\lambda}_i w_{T-i,A} \\ \bar{\lambda}_i &= \frac{\lambda^i}{\sum_{i=0}^{T} \lambda^i} \end{aligned}$$

Although these two persistence filters are rather simple and of heuristic nature, we will still leave it at that and focus more strongly on cost-reduction filters now.

Quite generally, the idea of cost-reduction filters is to choose portfolio weights that differ from originally proposed target weights in order to minimize transaction costs. Whenever portfolio weights of two different strategies do deviate, however, this also could open the door to differences in performances. Hence, we could get a performance that differs from the performance associated with the original targets. And if target weights have been chosen wisely, then this should rather lead to outcomes that are worse.

However, differences in weights only represent a necessary condition for differences in performance, but not a sufficient one. In other words, two portfolios could be quite different in their assets but still lead to very similar performances. One example is when two portfolios invest into different representative ETFs for the same regional market, like S&P500 and NASDAQ indices for U.S.. Then although their weights fully differ for this region, as they own different ETFs, they basically invest into the same underlying risk factors.

So the final target of all cost-efficiency filters is to reduce overall transaction costs, but without affecting performance outcomes too seriously. Still, we will now start with a very commonly used approach to cost reduction that does not at all take into account any possible consequences on portfolio and performance deviations: trading with fixed trading frequency only. This means that although target weights are produced on a daily basis, we only trade at a lower frequency, like for example on a weekly, monthly or quarterly basis. This way, naturally, trading costs will be reduced compared to a daily realization of target weights. Following a fixed trading frequency can also be written in terms of a weight filter framework.

Definition 3.3.3: ***[Fix trading frequency]***

Let w_0 denote current portfolio weights and w_A denote single-period portfolio target weights. Let further f denote some trading frequency. Then a fixed trading frequency approach can be written as weight filter with transformation function h:

$$h^{\mathit{FixFreq}}(f) := h(w_0, w_A, t, f) = \begin{cases} w_A, & \mathit{if} \quad (t \bmod f) = 0 \\ w_0, & \mathit{if} \quad (t \bmod f) \neq 0 \end{cases}$$

This fix trading frequency does not at all take into account whether currently held portfolio weights and targets actually differ or not. Hence, a natural improvement is to base trading decisions on information about differences of portfolio weights.

Definition 3.3.4: ***[Minimum turnover filter]***

Let w_0 denote current portfolio weights and w_A denote single-period portfolio target weights. Let further ϕ denote some threshold of maximally allowed weight differences. Then the minimum turnover filter is defined by the following transformation function h:

$$h^{\mathit{TOmin}}(\phi) := h(w_0, w_A, \phi) = \begin{cases} w_0, & \mathit{if} \quad \frac{1}{2}\sum_{i=1}^{d} |w_{0,i} - w_{A,i}| \leq \phi \\ w_A, & \mathit{if} \quad \frac{1}{2}\sum_{i=1}^{d} |w_{0,i} - w_{A,i}| > \phi \end{cases}$$

Using this filter we now only trade whenever the turnover required to trade to the target portfolio exceeds some minimally required threshold. Thus, differences in weights that are smaller are ignored, and no trading is triggered in these cases.

Although this cost-reduction filter still is quite simple in nature, it nevertheless already introduces a fundamental change to trading behavior, especially when asset management is done for multiple clients per strategy. Previously, trading did occur at fixed dates, where all clients are traded to the same target portfolio. In other words, any differences in individual client portfolio weights were also reset in regular intervals, such that clients should have maximally aligned portfolios after each trading event. Now, using a minimum turnover filter, trading events do not occur in equidistant intervals anymore, but proportionally to how fast target weights do change over time. This way target weights might change more quickly in turbulent market periods, which in turn would trigger more frequent trading, too. In addition, however, clients now might become traded at different points in time, as they do not necessarily need to exceed their respective minimally required turnover threshold simultaneously. Small differences in client portfolio weights will always emerge at some point, whether they are caused by different rounding due to different portfolio sizes, different client cash-flows related to withdrawals and incoming deposits, or different tax situations. In a setting of fixed trading dates such differences are offset regularly, and clients align after each trading event. When trading events are triggered individually, however, at different points in time, there is no regular anchoring anymore, such that client portfolios will generally differ at all times.

Setting the minimum turnover threshold ϕ low enough, final portfolio weights chosen can be guaranteed to remain sufficiently close to target portfolio weights. This way, portfolio performances will also not differ too much from the performances that would be achieved by target portfolio weights themselves. As pointed out before, however, deviating weights are indeed the source of performance deviations, but they do not necessarily have to cause them. One hence could try to optimize cost-reduction filters such that they allow weight deviations as long as portfolio properties and performances are not affected too much. This way, we would not use weight differences as a proxy for portfolio differences anymore, but would directly evaluate portfolio differences themselves.

In order to come up with such an approach, one first needs to define a way to measure portfolio similarity other than through portfolio weights. A first guess might be to define similarity based on portfolio moments, such that two portfolios are similar whenever their risk-return profiles are just matching closely enough. However, even when expected returns and volatilities perfectly match, portfolios still can be subject to significantly different risk factors, such that they experience their ups and downs in completely different financial environments. For example, imagine a U.S. and an emerging markets equity portfolio with equal absolute risk-return profiles. Still, their performances will generally differ substantially over time. A more appropriate measure is to directly evaluate differences between portfolio returns on sub-periods, like e.g. daily portfolio returns.

Definition 3.3.5: *[**Tracking error**, (Roll and Anderson 1992)]*

Let $(r_{t,0})_t$ and $(r_{t,A})_t$ denote portfolio return series of two different portfolios. Then the tracking error between both portfolios is defined as the standard deviation of the differences of their respective portfolio return series:

$$\mathit{TE} = \sigma(r_{t,0} - r_{t,A})$$

Tracking error is widely used in the passive portfolio management industry, where the target is to track a pre-defined market index as closely as possible. Thereby tracking error is used as a measure for how closely the actual portfolio follows the index. Tracking errors are minimal whenever portfolio variances match and also portfolios do fluctuate based on the same risk factors, i.e. ups and downs of the portfolios occur at the same time. Expected portfolio returns, however, do not directly influence tracking error, such that portfolios with small tracking errors still could significantly differ in terms of expected returns. When both portfolios are based on an equal universe with modest number of assets, however, expected portfolio returns should usually not differ widely when tracking errors are small.

Computation of tracking errors also can be reformulated in terms of the asset covariance matrix. This way, one does not need to construct historic return series in order to derive tracking errors between all portfolios, but only needs to estimate a single covariance matrix. With portfolio weights $w_{t,0}$ and $w_{t,A}$, the variance of return differences can be reformulated according to

$$\begin{aligned}\mathbb{V}(r_{t,0} - r_{t,A}) &= \mathbb{V}\left(w_{t,0}' r_t - w_{t,A}' r_t\right) \\ &= \mathbb{V}\left((w_{t,0} - w_{t,A})' r_t\right) \\ &= (w_{t,0} - w_{t,A})' \Sigma_R (w_{t,0} - w_{t,A})\end{aligned}$$

Hence, for tracking error we get

$$\sigma(r_{t,0} - r_{t,A}) = \sqrt{(w_{t,0} - w_{t,A})' \Sigma_R (w_{t,0} - w_{t,A})}$$

Measuring tracking errors this way will generally lead to different magnitudes of tracking errors for different risk categories. High risk categories do fluctuate more, and hence there also is more scope for deviations between portfolio returns. In contrast, low-risk categories usually have returns close to zero for short time horizons, so even when portfolios widely differ, portfolio returns can only deviate to a certain degree. In order to offset this effect, we will look at relative tracking errors instead: tracking errors that are scaled by the level of volatility of the target portfolio.

Definition 3.3.6: ***[Relative tracking error]***

Let $(r_{t,0})_t$ and $(r_{t,A})_t$ denote portfolio return series of two different portfolios. Then relative tracking error with regards to portfolio A is defined as tracking error divided by the volatility of the target portfolio:

$$\mathit{RTE}_A = \frac{\sigma(r_{t,0} - r_{t,A})}{\sigma(r_{t,A})}$$

Obviously, this measure of tracking error is not symmetric, as it depends on which of both portfolios is considered to be the target portfolio. As this is basically self-evident in our applications, we will nevertheless use relative tracking errors in the following parts, in order to make values more comparable across different risk categories. We are now ready to define cost-reduction filters based on portfolio similarity measured differently than on weight differences. First, let's generalize the minimum turnover filter of Definition (3.3.4).

Definition 3.3.7: ***[Minimum RTE filter]***

Let w_0 denote current portfolio weights and w_A denote single-period portfolio target weights. Let further ϕ denote some threshold of maximally allowed relative tracking error. Then the minimum RTE filter with regards to portfolio A is defined by the following transformation function h:

$$h^{\mathit{RTEmin}}(\phi) := h(w_0, w_A, \phi) = \begin{cases} w_0, & \text{if} \quad \mathit{RTE}_A \leq \phi \\ w_A, & \text{if} \quad \mathit{RTE}_A > \phi \end{cases}$$

Hence, this way trading gets triggered whenever portfolio behavior did drift apart too much, instead of when differences in portfolio weights are too large. This prevents portfolios from becoming too different from single-period target portfolios, which ensures that portfolio properties are close to optimal ones. Still, keeping portfolio properties close to targets is only one side of the equation, and we do not yet take into account in any way at what costs this is achieved. Ideally, we will only ensure portfolio similarity as long as it does not cause too much trading costs. Hence, we now will set up a trading cost aware way of tracking error minimization. This can be achieved quite similarly to diversification-aware portfolio selection by introducing an objective function that reflects some tradeoff between portfolio similarity and trading costs. A simple example is by using a linear tradeoff between both quantities, with a penalty factor for trading costs that regulates how strongly costs will have to be taken into account. We will assume that trading costs are variable costs and hence increase with the amount that is traded.

Definition 3.3.8: ***[Trading cost aware tracking error minimization]***

Let w_0 denote current portfolio weights and c_i denote variable trading costs for asset i. Then the following optimization problem will determine optimal portfolio weights under consideration of both portfolio similarity to target portfolio A and trading costs.

$$h(w_0, w_A, \Sigma_R, \phi, (c_i)_i) = \arg\min_{w} \quad \mathit{RTE}_A + \phi \sum_{i=1}^{d} c_i |w_i - w_{0,i}|$$

Primarily, the linear tradeoff function used in this trading cost aware tracking error minimization is motivated by the fact that it is simply difficult to relate portfolio similarity and trading costs in a meaningful way. Both measure completely different things, and it is not clear how an optimal tradeoff between them should look like. The linear function at least reflects the obvious and minimal requirement that smaller portfolio similarity and larger trading costs are both undesirable properties. Nevertheless, even in this simple setting it is hard to determine a suitable and meaningful penalty factor ϕ.

In order to avoid this problem, we will deal with cost-efficient tracking error control in a slightly different and more intuitive way. Therefore, the optimization problem will be structured such that transaction costs can be minimized for any given level of maximally allowed relative tracking error.

Definition 3.3.9: ***[Cost-aware capped relative tracking error]***

Let w_0 denote current portfolio weights and T_{MAX} a given level of maximally allowed tracking error with regards to portfolio A. Then the following optimization problem will minimize transaction costs among all portfolios that comply with the required level of portfolio similarity.

$$\begin{aligned} h^{RTEcap}(T_{MAX}) := h(w_0, w_A, \Sigma_R, T_{MAX}, (c_i)_i) = \arg\min_{w} \quad & \frac{1}{2}\sum_{i=1}^{d} c_i |w_i - w_{0,i}| \\ \textit{subject to} \quad & RTE_A \leq T_{MAX} \end{aligned}$$

Now one does not need to explicitly define a tradeoff function between costs and portfolio similarity anymore, but one needs to define a cap for relative tracking error instead. Once one has decided on the magnitude of portfolio tracking errors that one is willing to live with, the optimization will realize this requirement with minimal costs. Compared to an explicit tradeoff function it comes at a disadvantage though: the level of portfolio similarity will be the same no matter how costly it can be achieved. In reality it might be worthwhile to give more room to tracking errors when large trading costs would have to be paid.

A huge benefit of the capped tracking error approach, however, is that it can be reformulated as convex programming problem.

Lemma 3.3.10: [Reformulation as convex optimization problem]

The optimization problem defined in Definition (3.3.9) can be reformulated as convex optimization problem.

Proof. Let's first inspect the objective function. With $S' = [c_1, \ldots, c_d]$ it can be rewritten as vector multiplication:

$$\sum_{i=1}^{d} c_i |w_i - w_{0,i}| = S'|w - w_0|$$

By artificially bloating up the optimization problem to higher dimensions, the absolute value function can be eliminated and the objective function becomes linear (T. Roncalli 2013).

Furthermore, the constraint can be reformulated as quadratic constraint. With

$$(w - w_A)'\Sigma_R(w - w_A) = w'\Sigma_R w - 2w'\Sigma_R w_A + w_A'\Sigma_R w_A$$

we get:

$$
\begin{aligned}
& & \mathrm{RTE}_A &\leq T_{\mathrm{MAX}} \\
&\Leftrightarrow & \frac{\sqrt{(w-w_A)'\Sigma_R(w-w_A)}}{\sqrt{w_A'\Sigma_R w_A}} &\leq T_{\mathrm{MAX}} \\
&\Leftrightarrow & \sqrt{(w-w_A)'\Sigma_R(w-w_A)} &\leq T_{\mathrm{MAX}}\sqrt{w_A'\Sigma_R w_A} \\
&\Leftrightarrow & (w-w_A)'\Sigma_R(w-w_A) &\leq T_{\mathrm{MAX}}^2 w_A'\Sigma_R w_A \\
&\Leftrightarrow & w'\Sigma_R w - 2w'\Sigma_R w_A + w_A'\Sigma_R w_A &\leq T_{\mathrm{MAX}}^2 w_A'\Sigma_R w_A \\
&\Leftrightarrow & w'\Sigma_R w - 2w'\Sigma_R w_A &\leq T_{\mathrm{MAX}}^2 w_A'\Sigma_R w_A - w_A'\Sigma_R w_A
\end{aligned}
$$

□

Summing up, we now do have a couple of weight filters at our disposal. The only input that all of them share are current target portfolio weights, but other than that they generally use different information. Persistence filters are independent of current portfolio weights but work on historic target weights. Cost-reduction filters require current portfolio weights, target weights and additional information like trading dates for fixed trading frequencies or covariance matrices for relative tracking error calculations. In addition, filters also differ in complexity: some are simple transformations of inputs, while others involve the solution of specific optimization problems.

One very useful aspect of capturing all of these weight transformations in a common weight filter framework is that it is very easy to combine them that way. For example, we can apply some chosen persistence filter in a first step, and use additional cost-reduction filters to cost-efficiently realize target weights after they have been corrected for persistence. Furthermore, one can use different combinations of cost-reduction filters in order to calibrate a given trading strategy to almost any level of desired turnover and trading costs. While weight filters can be arbitrarily combined, these combinations are generally not commutative. As an example, let's assume that we want to combine a moving average filter with a minimum turnover filter. Applying the moving average filter first, we might end up without any trading whenever the minimum turnover threshold is not exceeded after the targets have been modified by averaging. If the sequence of filters is changed, however, and the minimum turnover threshold again is not exceeded, we will simply take the moving average over historic target weights and current portfolio weights. This average will generally differ from current portfolio weights, such that we would end up trading in this case. Weight filters also have to be applied recursively, one trading day after the other, and can not be applied to full series of target weights directly. The reason is that individual trading days are interconnected, as finally realized target weights today are the initial weights of the trading decision tomorrow. In other words, if I trade today and hence align the portfolio more closely with the desired target portfolio weights, then tomorrow I will already have a portfolio that is reasonably close to the optimum and hence most likely will not need to be traded again.

A particularly powerful combination of weight filters is the combination of minimum RTE filter and cost-aware capped RTE. The minimum RTE filter will trigger trading whenever current portfolio properties are not sufficiently similar to target portfolio properties anymore.

And, when trading is triggered, one will not just re-balance current portfolio weights to target weights completely, but cost-aware capped RTE optimization will align the portfolio with the target portfolio only up to some required degree of relative tracking error. Effectively, final portfolios hence will always remain within certain bounds of relative tracking error, and this will be realized in a cost-efficient manner. Also, relative tracking error is just one way of measuring portfolio similarity, and the same approach could also be applied to one of many other deviation measures (Rachev, Stoyanov, and Fabozzi 2008).

Overall, the general idea of these weight filters is to realize a series of desired portfolio properties over time in a slightly different but less costly way. Their purpose is not, however, to reduce noise in original target portfolio decisions. In other words: in a world without trading costs, filters generally should deteriorate portfolio decisions, and realizing target weights every day should be a superior trading strategy. If this is not the case, one should step back to single-period weight decisions and evaluate whether the market model used to derive these target weights is good enough, or whether one needs to deal with noisy estimates more elaborately already at this step. When single-period target weights are good enough, however, and trading costs exist, weight filters hopefully increase returns net of transaction costs.

Chapter 4

Empirical application

In previous chapters we have introduced a couple of components that can be used to build a dynamic risk management strategy. In particular, Chapter (2) has introduced classic Markowitz portfolio selection, as well as an alternative diversification-aware portfolio selection strategy. Both can be used to derive single period target weights without consideration of transaction costs. Furthermore, in Chapter (3) we developed weight filters that can be used to realize a series of desired single-period portfolio targets in a slightly modified and less costly way. In this section we will now put these individual components together and backtest several dynamic risk management strategies in an empirical application of multi-asset portfolios.

4.1 Backtest settings

For the empirical application we first need to define a universe of assets on which the dynamic risk management strategy should operate on. Note that this is a highly influential decision for performance results of backtests, and one could potentially introduce a significant selection bias at this step. For example, feeding the dynamic risk management strategy in backtests with historically nicely performing assets only, it is almost guaranteed that final strategy returns will be pleasant as well. In order to reduce the danger of compromising results by any kind of selection bias, we will not put the asset universe together ourselves, but try to replicate a universe chosen by others as closely as possible. For this, we will adopt the same universe that is used by Scalable Capital, the currently largest European digital wealth manager, as closely as possible. This universe comprises fourteen exchange-traded funds (ETFs), representing market indices from different asset classes and regions denominated in Euro currency.

ETFs are baskets of multiple underlying securities that jointly trade for a single price at an exchange. Each of these baskets is closely related to an underlying benchmark index that shall be tracked: ETFs try to match the evolution of a benchmark index as closely as possible. This way, a single ETF will provide exposure to a full market segment already. For example, an ETF that tracks the S&P500 will undergo almost the same market movements as the

S&P500 index itself. Holding the ETF hence is as if one would own shares of all companies of the S&P500 with fractions equal to the individual assets' weights in the index. A single ETF hence is already invested into multiple individual positions, thereby serving the objective of diversification.

When adopting the ETF universe of Scalable Capital, we also benefit from additional side-effects other than the selection of markets and asset classes themselves. Individual ETFs are also chosen with regards to characteristics that alleviate any real world implementation of dynamic trading strategies. In particular, they are chosen with regards to liquidity aspects, and favoring low total expense ratios, physical replication and pleasant tax properties from the perspective of European investors.

As ETFs are a rather new financial product, however, historic prices for most of the selected ETFs are only available going back a couple of years. This would highly limit the informative value of backtests, as we only could examine what the dynamic trading strategy would have done for a short period of time. Hence, we will not use real ETF price data for backtests, but use total return data for the benchmark indices that are tracked by the individual ETFs instead. Even then, four of the fourteen indices do not go back far enough, so that we will replace them with different indices of the same market that go back further into the past. This way, we manage to put together a universe with data going back to January 2nd, 1999. The following table contains a list of all indices that are used, with checkmarks in the second column indicating those indices where we did have to deviate from the original target universe of Scalable Capital. The third column shows short labels that we will use to reference individual assets in the following parts.

Market index	Modified	Label
Markit iBoxx Pfandbriefe Index TRI	✓	covEu
Bloomberg Barclays US Government 10 Year Term Index	✓	govUs
Bloomberg Barclays Euro Treasury Bond Index		govEu
J.P. Morgan EMBI Global Core Index		govEm
Markit iBoxx USD Liquid Investment Grade Index		corpUs
Bloomberg Barclays Euro Corporate Bond Index		corpEu
STOXX Europe 600		eqEu
DAX		eqDe
MSCI Japan Net Total Return Local Index	✓	eqJp
S&P 500 Total Return		eqUs
MSCI Pacific ex Japan Total Return Net		eqAs
MSCI Emerging Markets Total Return Net		eqEm
FTSE EPRA/NAREIT Developed Index Total Return EUR	✓	realGlob
Thomson Reuters/Jefferies CRB Total Return		comGlob

Most of the universe hence consists of assets that belong to either equities (**eq**) or bonds , and one real estate (**real**) and one commodity (**com**) index add to these more traditional asset classes. Within the class of equities, however, the universe comprises a widespread selection of global asset markets: country indices for U.S. (**Us**), Germany (**De**) and Japan (**Jp**), as well as regional indices for Europe (**Eu**), the Pacific region (**As**) and several emerging markets (**Em**). On the equity side we hence have a globally diversified mix, comprising mainly large and to a minor extent also mid cap stocks. Similarly, the asset class of bonds does also consist of a number of heterogeneous sub-groups. European and U.S. government bonds (**gov**) are part of it, which are the most popular bond markets. In addition, the universe also comprises corporate bond (**corp**) indices for both regions. On the low-risk end of bonds a covered bond (**cov**) index extends this selection, while emerging market government bonds extend the more traditional bond markets on the high-risk end. Altogether, universe assets cover wide ranges of risk: from low-risk covered bonds to European and U.S. government bonds, to medium risk assets like corporate bonds and emerging market government bonds, to high-risk equities and real estate and commodities.

There is another aspect to the selection of the universe, however, that also deserve attention: the universe does not include any risk-free asset. From mean-variance portfolio optimization we know that the addition of a risk-free asset will extend existing investment opportunities. Portfolios on the efficient frontier then get Pareto-dominated by portfolios that are a combination of the risk-free asset and the tangency portfolio only. From a methodological point of view the diversification-aware approach easily could be extended to the case with an existing risk-free asset. As there is no uncertainty regarding the properties of the risk-free asset, the single-period portfolio choice does not need to be made robust against estimation errors of the moments of the risk-free asset. Hence, diversification only needs to be required for all remaining assets. Optimal portfolio choices in the presence of a risk-free asset then should simply use a different tangency portfolio: the tangent to the diversification-aware frontier is the new determining factor of the tangency portfolio, and not the one to the efficient frontier. Despite the methodological changes due to the inclusion of a risk-free asset, however, the more crucial question is whether a risk-free asset even exists for our purposes. Credit risk of U.S. government debt usually is assumed to be negligible, such that U.S. government bonds are generally perceived riskless when held until maturity. Due to the stochastic nature of our dynamic trading strategy, however, individual trading events are not predictable, such that we do not know in advance when the next re-balancing event will occur. Hence, it is impossible to determine a bond that perfectly matches the maturity to the next trading event upfront. Another possible proxy that could be used is an interest bearing bank account. However, this only works for small portfolios, as deposit insurance only holds for a limited amount of cash, and any cash in exceedance of this limit is subject to credit risk. A true representative risk-free asset hence is not really easy to find in the first place. Let's for the moment, however, assume that a risk-free asset would exist. In the presence of short-selling constraints, the only portfolios that would change are portfolios left to the tangency portfolio. This range of affected portfolios is usually rather small, however, as the universe already comprises a couple of low-risk assets that will let the tangency portfolio emerge in low-risk ranges already. Hence, for all these reasons, we will not use a specially treated risk-free asset, but we are pleased with the amount of low-risk assets that are already part of the present selection.

Now that the universe is set up, we also need to define a couple of other settings that will influence backtest results. In particular, we will apply some standard simplifications that introduce deviations from what one would get when the trading strategy actually would be realized. We will start with a deviation that we already have mentioned before: using total return index data instead of ETFs themselves.

Even though ETFs usually have an explicitly stated tracking goal, they still will at least slightly deviate from their benchmark index in general. The main reason for this is that the index itself does not take into account any transaction costs, and even large changes in the weights of the underlying benchmark index occur without any transaction related drain of money. In contrast to such a synthetic index calculation, any ETF really needs to buy and sell securities in order to change its weights, which in turn causes transaction costs. To keep these transaction costs within bounds, an ETF usually tries to track its benchmark index without directly investing in all of its underlyings, a technique call stratified sampling. Instead of buying all individual index components, the ETF only invests into the major drivers of the index, thereby trying to capture the main properties of the full index. Through this strategy, the underlyings of an ETF and its respective benchmark index will generally deviate. Hence, ETF price and index value will not always coincide perfectly, a tracking error exists.

In addition to the tracking error, there is another deviation between ETF prices and benchmark index values. At the end of each day, the ETF provider calculates and publishes the net asset value (NAV) of the ETF, which is the total per share value of the basket of securities that the ETF is holding, plus cash and any other assets, less any liabilities. Sometimes this NAV might differ from the price of the ETF. In other words, the sum of the individual prices of all positions does not fully match the overall price of the ETF that one would get when selling the ETF at the exchange. Some of these deviations actually might be justified, as for example ETFs usually trade with much more liquidity than some of the individual positions that they are invested into. Hence, ETF prices sometimes reflect new pricing information earlier than individual illiquid assets. Whenever there is no justification, however, there exists an arbitrage opportunity that so-called authorized participants could use to their advantage. These authorized participants are basically allowed to exchange a share of an ETF into the underlying basket of assets and vice versa. This way they can benefit from any unjustified deviations between ETF prices and NAVs, and in their role as arbitrageurs they automatically will help to further align both prices again.

In addition to tracking errors and deviations between ETF prices and NAVs, there exist two further aspects where our backtest based on index prices will introduce deviations from reality. First, we will work on total return data, such that we implicitly omit any kinds of distributions that occur in real life. When distributions exist, there will be additional cash-flows that basically cause two things. First, distributions need to be re-invested again, or otherwise they will just lie idle as cash on a bank account. Any such re-investment of cash will also cause trading costs. And second, distributions generally are tax relevant, and as such they could cause tax payments when investors are not exempt from taxes. A second aspect of simplification related to total return index prices is that we will simply use closing prices for all assets in backtests. As individual indices represent different regions, however, these closing prices will be asynchronous: they do not happen at the same point in time, but there is a time shift of several hours between individual closing times. In other words, there is no time of the day where simultaneous trading of all assets at their respective closing prices is actually possible. At the time that European stock exchanges close, the New York Stock Exchange will still be open for some hours. Hence, S&P500 closing prices will generally include some hours worth of additional information more than European market closing prices. Even worse, this effect will usually not only affect trading prices, but also estimates of correlations and volatilities (Burns, Engle, and Mezrich 1998).

In addition to slightly simplified price data, we will also do not deal with transaction costs in full detail. Transaction costs usually consist of two components: direct and indirect transaction costs. Direct transaction costs are what will be charged by custodian or exchange for trading. These costs could be a flat trading fee, variable trading costs that increase with trading volume, or variable trading costs that increase per asset that is touched. In contrast, indirect trading costs are those that do not occur explicitly on any invoice or bill. For example, trading usually does not happen at mid prices, but at bid or ask prices. Hence, even without direct transaction costs, buying and immediately selling an asset again will slightly reduce wealth as one will lose half of the bid-ask spread per trade on average. Bid-ask spreads usually differ per ETF, and they are lost together with direct trading costs instantly per trading. In contrast to that, we will simply subtract both direct and indirect trading costs proportionally to turnover from final portfolio returns.

Another major difference to reality is that we will compute backtest returns based on portfolio weights instead of portfolio volumes. In reality, almost all ETFs are still tradable in integer volumes only, and fractional shares are not possible yet. Hence, we can not realize any desired fractional weight target in reality, but only approximate it with regards to whole-numbered asset shares. In reality this will slightly distort target weights, and generally increase idle cash positions as it is not possible to invest money into ETFs down to the last cent. The larger the portfolio values are in comparison to asset prices, the smaller the distortions of rounding will get.

Not only do we make simplifications with regards to market data and trading, but also with regards to investor characteristics. So we will assume that portfolios are unaffected by any external cash-flows, and investors do neither add further deposits nor withdraw money in-between. All such external requests would have to be reflected through additional trading, which in turn could lead to portfolio weight distortions and trading costs.

Last but not least, another major simplification is that we will artificially set investor tax rates to zero. In reality, many countries do settle taxes immediately at individual trade events by using sale revenues as tax liquidity. Hence, proceeds of asset sales could be substantially lower than market prices. For example, let's assume a tax rate of 25%, an asset that has been bought for 60 Euro and now is sold for 100 Euro. Then we will have to pay 25% taxes on a price gain of 40 Euro, such that we will only get 90 Euro instead of 100 Euro at sale. In a world with taxes, re-balancing is not a self-financing event anymore, but portfolio values can decrease due to tax liabilities.

4.2 Missing values

Price data of the chosen indices are from Bloomberg and are available from January 1st, 1999, until April 10th, 2017. Individual time series come with prices for business days that are not perfectly aligned, so that some assets have missing prices for some dates. This is not unexpected, as individual indices represent different regional markets and hence are affected by different regional holidays as well. Figure (4.1) shows the number of missing observations per asset, as well as the maximum number of consecutive missing values.

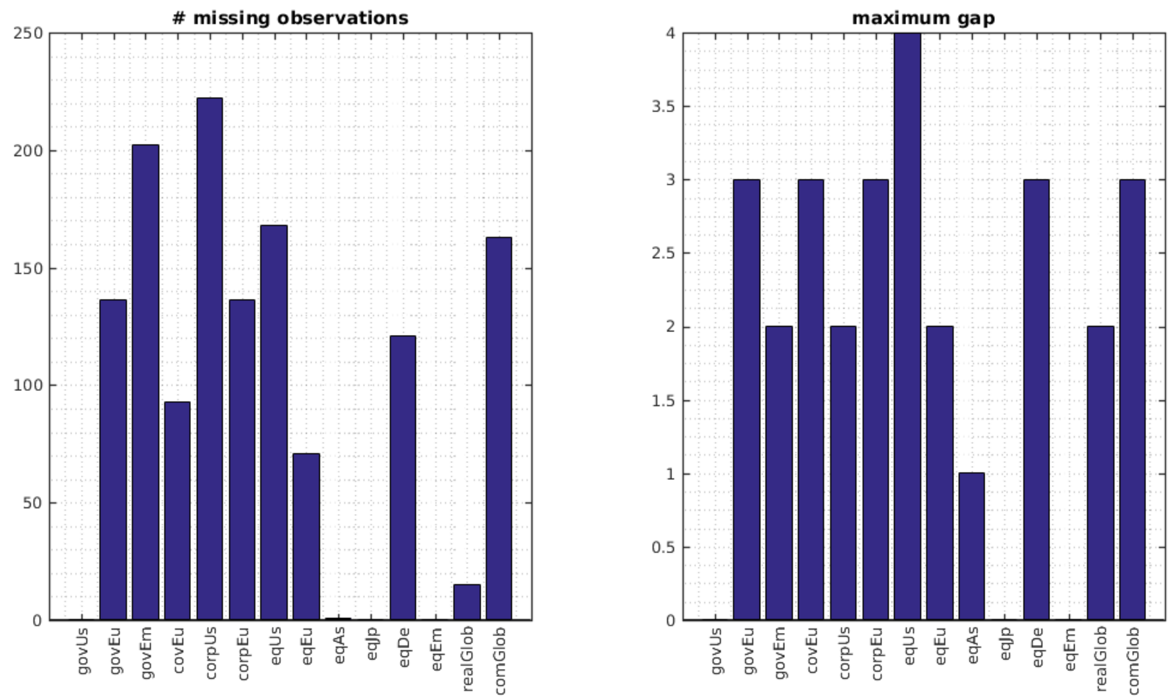

Figure 4.1: Number of missing observations and maximum gap size per asset

There are a couple of ways to deal with missing index prices, but we will choose a quite simple and prominent way to do so: Last Observation Carried Forward (LOCF). It naively imputes missing observations by simply using the last existing price instead. As with any other method of data imputation, one should at least put some thought into whether the imputation introduces any artificial patterns. As we will see, LOCF is far from optimal, but we will use it nevertheless as this is in line with our target to find an asset allocation strategy that is robust against errors in inputs.

Missing prices quite generally could be attributed to two different causes. First, there simply never was a price, as it is the case for days when stock exchanges are closed (e.g. holidays). And second, cases where there originally was a price, but due to incomplete data collection or some data error it has not been included into the database. Based on a simple reasoning with equity prices one can see that both causes for missing prices might have to be dealt with differently. Imagine that a given missing observation is due to closed stock exchanges on a public holiday. Then we generally can assume quite safely that the existence of this day should not have an influence on stock prices. Even though stock prices tend to increase on average, there is no convincing evidence that stock prices also increase over weekends or holidays (an effect that would be reflected in prices of morning auctions). This argument does not hold anymore when missing values are caused by data errors, where in reality there was a price movement that is just not reflected in the database. Assuming that stock prices on average tend to increase, then we ideally should assign a positive expected value to the unobserved price movement in that case.

LOCF, in contrast, simply fills missing observations with last existing prices in any case. This way, it artificially inflates the number of returns that are equal to zero. And, given that prices actually did exist, but only are not reflected in the observations due to data errors, it also distorts single day return distributions by mixing multiple-day returns into the series of single day returns. An example for this is given in the following table. Thereby Δ denotes the difference between consecutive price entries, and NA represents missing values.

Table 4.2: LOCF for multiple consecutive missing values

Prices	LOCF	LOCF Δ
100	100	
105	105	5
NA	105	0
NA	105	0
NA	105	0
145	145	40

4.3 Moment estimation

In this section we will introduce the estimator that is used for asset return moments. As already pointed out before, we are not interested in finding a particularly good estimator for the asset returns at hand. Ultimately, we rather want to check whether the diversification-aware approach is robust enough to be able to deal with sub-optimal moment estimates, too. The only requirements that we have for the estimator hence are more of a practical nature. First, it should not require any data additional to the historic returns of the respective assets. Second, it should not use any particular knowledge with regards to asset class peculiarities, but simply treat all assets the same. And third, it should not require any complicated numerical optimization during estimation. For all these reasons, we will settle for the EWMA estimator proposed in ("RiskMetrics Technical Document" 1996), an estimator that is quite popular among practitioners.

We will now give a short introduction to the EWMA estimators. Thereby we will closely follow ("RiskMetrics Technical Document" 1996) for all definitions and derivations.

Definition 4.3.1: ***[EWMA mean estimator]***

Let $0 < \lambda_\mu < 1$ denote a decay factor, and $r_{t,i}$ denote the return of asset i at time t. Then the exponentially weighted moving average model for the mean asset return $\mu_{t,i}$ is given by

$$\mu_{t,i} = \sum_{k=1}^{t} \frac{\lambda_\mu^{k-1}}{\sum_{k=1}^{t} \lambda_\mu^{k-1}} r_{t-k+1,i}$$

As $0 < \lambda_\mu < 1$, the most recent observations will carry the largest weight in the EWMA formula, while observations far in the past will only enter with tiny weights. Due to convergence of the geometric series, for the limit case $t \to \infty$ the formula can be further simplified:

$$\sum_{k=1}^{\infty} \lambda_\mu^{k-1} = \frac{1}{1-\lambda_\mu}$$

$$\Rightarrow \;\; \mu_{t,i} = (1-\lambda_\mu) \sum_{k=1}^{\infty} \lambda_\mu^{k-1} r_{t-k+1,i}$$

Based on this expression, the EWMA mean estimator can be reformulated into recursive form:

$$\mu_{t+1,i} = \lambda_\mu \mu_{t,i} + (1-\lambda_\mu) r_{t+1,i}$$

This way, λ_μ has a clear interpretation, as new observations will enter with weight $(1-\lambda_\mu)$ to the existing estimate of $\mu_{t,i}$, which itself gets weighted by the decay factor λ_μ. A decay factor closer to 1 will put less weight on the last observation, and hence represents a less responsive estimator. Still, by putting larger weight on recent observations, the estimator will react faster to data than what a sample mean would do. For variances and covariances we can define EWMA estimators in a similar way.

Definition 4.3.2: ***[EWMA variance/covariance estimator]***

Let $\sigma_{t,ij}$ denote the covariance between asset i and asset j. Let further $\bar{r}_{t,i}$ denote the sample mean of asset i with observations up to time t:

$$\bar{r}_{t,i} = \frac{1}{t} \sum_{k=1}^{t} r_{k,i}$$

Then the exponentially weighted moving average model for covariance $\sigma_{t,ij}$ is given by:

$$\sigma_{t,ij} = \sum_{k=1}^{t} \frac{\lambda_\sigma^{k-1}}{\sum_{k=1}^{t} \lambda_\sigma^{k-1}} (r_{t-k+1,i} - \bar{r}_{t,i})(r_{t-k+1,j} - \bar{r}_{t,j})$$

Again, with infinite history the estimator can be reformulated in recursive form:

$$\sigma_{t+1,ij} = \lambda_\sigma \sigma_{t,ij} + (1-\lambda_\sigma)(r_{t+1,i} - \bar{r}_i)(r_{t+1,j} - \bar{r}_j)$$

According to ("RiskMetrics Technical Document" 1996), there are two advantages over sample moments, which implicitly would use equally weighted observations. First, volatility will react faster to market shocks than with sample moments. And second, volatility declines exponentially, as the weight of observations that did cause a certain shock to volatility will get increasingly less over time.

Despite these potential advantages, however, the value added from EWMA estimators is far from being unanimously acknowledged in academic literature. For example, this is what Andrew Ang writes in (Ang 2014) about backward-looking estimators that are based on historic data only:

> "Using short data samples to produce estimates for mean-variance inputs is very dangerous. It leads to pro-cyclicality. When past returns have been high, current prices are high. But current prices are high because future returns tend to be low. While predictability in general is very weak, chapter 8 provides evidence that there is some. Thus, using a past data sample to estimate a mean produces a high estimate right when future returns are likely to be low. These problems are compounded when more recent data are weighted more heavily, which occurs in techniques like exponential smoothing."

Hence, we have good reasons to believe that EWMA moment estimates will generate large enough estimation errors in order to pose a challenge on the robustness of diversification-aware portfolio selection.

Furthermore, however, we still need to define actual values for decay factors λ_μ and λ_σ. While some optimal values for decay factors ideally would be estimated from data, we will simply fix some reasonable values in the first place. The values that we choose are:

- $\lambda_\mu = 0.99$
- $\lambda_\sigma = 0.95$

This way, moment estimates will quite quickly react to observations for variances and covariances, while EWMA mean estimates are more persistent. Hence, the responsiveness of estimators coincides with two wide-spread beliefs. First, second moments of asset returns are generally easier to estimate. They hence can be made more responsive and be estimated with less data than expected returns. And second, volatility tends to spike quickly as a reaction to market shocks. Hence, historic data tends to get outdated faster, which also requires higher responsiveness of estimators.

To make the responsiveness of different decay factors more meaningful, λ values also can be translated into half-life periods. Half-life periods measure the number of most recent observations required in order to get a cumulative weight larger than 50%. Assuming infinitely many observations, the half-life period can be analytically derived using geometric series. The joint weight of all observations that are at least K periods in the past is given by:

$$(1-\lambda)\sum_{k=K}^{\infty}\lambda^{k-1} = (1-\lambda)\lambda^{K-1}\sum_{k=1}^{\infty}\lambda^{k-1} = (1-\lambda)\lambda^{K-1}\frac{1}{1-\lambda} = \lambda^{K-1}$$

Given this formula, for any given joint probability of the most historic observations we can derive the associated threshold period K:

$$
\begin{aligned}
\lambda^{K-1} &\overset{!}{=} \alpha \\
(K-1)\log(\lambda) &= \log(\alpha) \\
K-1 &= \frac{\log(\alpha)}{\log(\lambda)}
\end{aligned}
$$

Solving for K, with $\alpha = 0.5$ and λ equal to $\lambda_\mu = 0.99$ and $\lambda_\sigma = 0.95$ respectively, we get:

$$
\begin{aligned}
K_\mu &= \frac{\log(0.5)}{\log(0.99)} + 1 = 69.97 \\
K_\sigma &= \frac{\log(0.5)}{\log(0.95)} + 1 = 14.51
\end{aligned}
$$

Hence, for the estimation of mean returns the most recent 70 days approximately get the same weight as all previous observations together. In contrast, for variances and covariances the half-life period is only 15 days.

4.4 Currency risk

Investing into global asset markets opens the door to currency risk: some asset prices can be denominated in foreign currency, thereby exposing investors to potentially adverse currency movements. Even when prices increase in terms of foreign currency, they still might decrease when viewed through the perspective of domestic currency.

We will assume that investors predominantly care about the value of their portfolio in terms of the currency of their country of living. Although not every investor might be related to a single currency only, there usually should be some currency where purchasing power matters most, and country of residence should be a reasonable determining factor. In other words, European investors are assumed to care about the portfolio value in EUR, while U.S. investors' interest should lie in prices given in USD. Hence, currency risk is an inherently asymmetric matter: what EUR is for some, is USD for others. To avoid asymmetries in notation we will generally denote currencies independent of actual countries or currency names, but relative to the perspective of investors. Any individual investor's currency of predominant interest will be referred to as **homeland** currency, while all other currencies become **foreign** currencies from his perspective. In addition, **local** currency denotes currencies with respect to countries: EUR is local currency of Europe, while USD is local currency of the U.S..

Another ambiguity with multiple currencies arises in the quotation of exchange rates: they can either be quoted in price notation or in quantity notation (Mankiw 2008). Price notation specifies how much units of homeland currency are required in order to buy a single unit of foreign currency. In contrast, quantity notation states the number of units of foreign currency associated with a single unit of homeland currency. In the following, we will generally state exchange rates in terms of quantity notation, such that an increase in foreign exchange rates will represent an appreciation of homeland currency.

Definition 4.4.1: *From the perspective of an individual investor, $\mathsf{FX}_t^{H/L}$ denotes the foreign exchange rate in quantity notation. When the homeland currency is unambiguous, we will also use the short-hand notation FX_t^{L}.*

Example A EUR/USD exchange rate of $\mathrm{FX}_t^{\mathrm{EUR/USD}} = 1.2$ means that a single EUR will buy 1.2 USD:

$$\begin{aligned} \mathrm{FX}_t^{\mathrm{EUR/USD}} &= 1.2 \qquad \Leftrightarrow \\ 1\,\mathrm{EUR} &= 1.2\,\mathrm{USD} \end{aligned}$$

With given exchange rates, foreign asset prices and asset returns can be converted into in homeland currency.

Theorem 4.4.2: [Foreign price and return conversion]

Let P_t^L denote a foreign asset's price in local currency. Then the corresponding price P_t^H denoted in homeland currency is given by

$$\begin{aligned} P_t^H &= P_t^L \frac{1}{\mathsf{FX}_t^{H/L}} \\ &= P_t^L \mathsf{FX}_t^{L/H} \end{aligned}$$

Furthermore, the return of a foreign asset can also be converted into homeland currency:

$$\begin{aligned} (1 + r_t^H) &= \frac{P_t^H}{P_{t-1}^H} = \frac{P_t^L}{P_{t-1}^L} \frac{\mathsf{FX}_t^{L/H}}{\mathsf{FX}_{t-1}^{L/H}} \\ &= \left(1 + r_t^{\mathsf{L}}\right)\left(1 + r_t^{\mathsf{FX(L/H)}}\right) \end{aligned} \tag{4.1}$$

Hence, the return in homeland currency can roughly be decomposed into two parts: an asset return in local currency, and a component that reflects changes in foreign exchange rates.

$$\begin{aligned} (1 + r_t^H) &= \left(1 + r_t^{\mathrm{L}}\right)\left(1 + r_t^{\mathrm{FX(L/H)}}\right) \\ &= 1 + r_t^{\mathrm{L}} + r_t^{\mathrm{FX(L/H)}} + r_t^{\mathrm{L}} r_t^{\mathrm{FX(L/H)}} \end{aligned}$$

$$\Rightarrow \quad r_t^H \approx r_t^{\mathrm{L}} + r_t^{\mathrm{FX(L/H)}}$$

Foreign exchange rates hence do have an influence on asset returns in homeland currency. For example, even when foreign assets achieve positive returns in local currencies, the picture might significantly differ after currency conversion. Over a period of multiple years, exchange rates can easily change in a magnitude similar to equities. Figure (4.2) shows pairwise exchange rates between four major currencies over the full period of our data sample: U.S. dollar (USD), Euro (EUR), Japanese Yen (JPY) and British pound (GBP). For better comparison, exchange rates are normalized such that they start with a value of one.

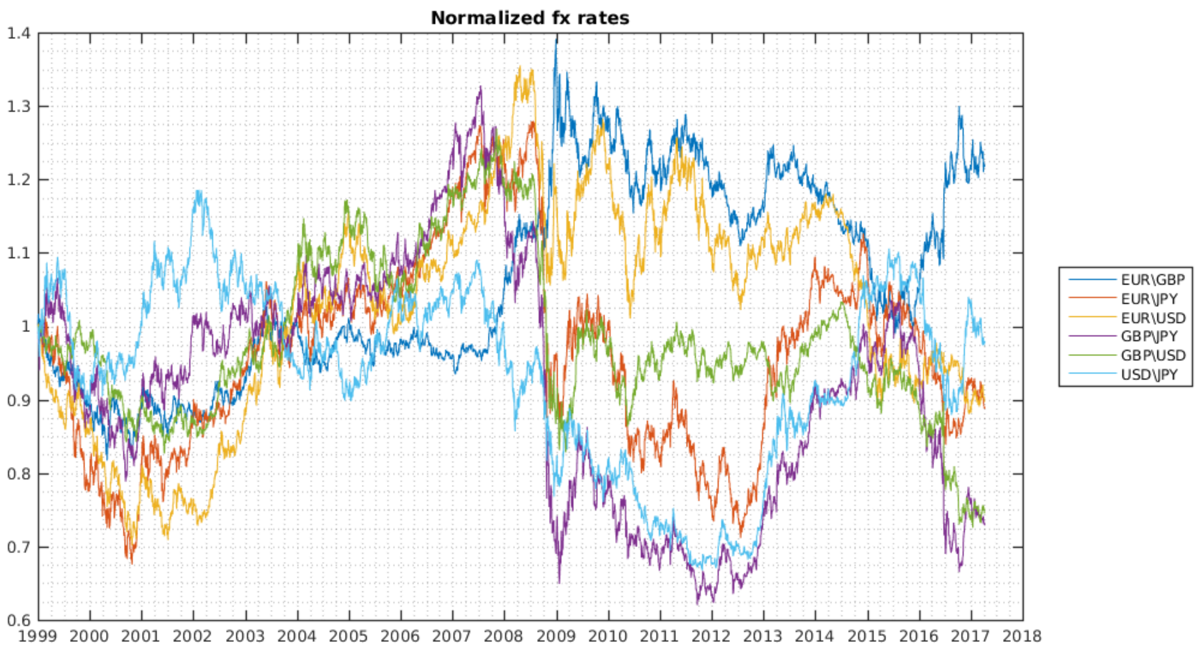

Figure 4.2: Normalized exchange rates

When looking at exchange rate changes over the full sample period, some currencies had rather moderate overall changes. Over the period of roughly 17 years, EUR did depreciate approximately 10% against both USD and JPY, and USD/JPY exchange rates almost did not change in value at all. Still, these exchange rates did fluctuate widely at least temporarily. From mid 2000 to mid 2008, EUR appreciated more than 70% against USD and JPY. Exchange rates quite generally have been subject to huge variability. For example, from mid 2007 to beginning of 2009 GBP did depreciate roughly 50% compared to JPY. UK asset prices would have had to rise 100% in value in local currency in order to offset such a drop in exchange rates from the perspective of Japanese investors.

This variability can be seen even better when looking at changes of exchange rates. Figure (4.3) shows logarithmic percentage changes for all currency pairs. As one can see, there is significant variability in daily exchange rate changes similar to the level of variability of equity returns. Also, all time series exhibit patterns of heteroskedasticity, with increased volatility during turbulent market periods.

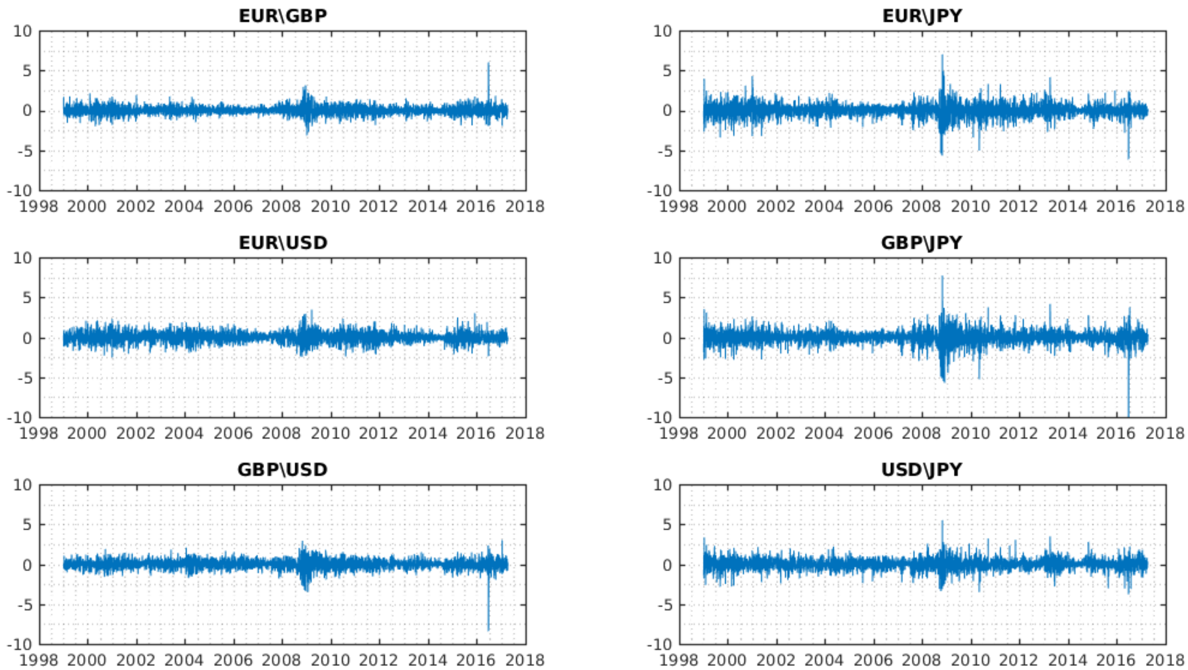

Figure 4.3: Logarithmic percentage changes of exchange rates

From Equation (4.1) we know that asset returns in homeland currency are both driven by local asset returns and foreign exchange rate changes. Hence, the stochastic properties of exchange rates will directly pass on to homeland currency asset returns. This might influence asset allocation decisions, as it might alter three components of the portfolio allocation decision:

- risk-return profiles of individual assets
- correlations between individual assets
- exposure to certain risk factors (e.g. “beta” to some benchmark)

While the first two points are quite obvious and important for any portfolio allocation problem, the last point usually only comes into effect under special requirements. In other words, it only becomes relevant when not only the risk-return profile of a portfolio matters, but also the exposure or co-movement with regards to some benchmark. For example, European funds invested in the U.S. stock market might find it attractive to have a certain degree of co-movement with the U.S. market, as it is one of the most popular yardsticks by which private investors will evaluate the asset manager’s capabilities. In other words: when global markets, and in particular the U.S. market, go up, clients also expect their portfolios to go up. If this is not the case, as it might happen due to adverse exchange rate movements, a fund manager might have to engage in cumbersome explanations. Clients sometimes forget that the U.S. market in USD is something that is not available to investors with different homeland currencies. News about rising U.S. stock market indices build up expectations, such that a possibly falling U.S. stock market position due to adverse currency movements hits a lack of understanding. This is something that fund managers might want to avoid.

The other two more obvious effects of currency movements, changing asset moments and correlations, we will now examine based on the present data sample. Therefore, we estimate asset sample moments and covariance matrix on weekly discrete asset return data for several currencies. First, estimation is done on asset returns in local currency, with results shown in Figure (4.4). Values are scaled to annual basis to make them more meaningful.

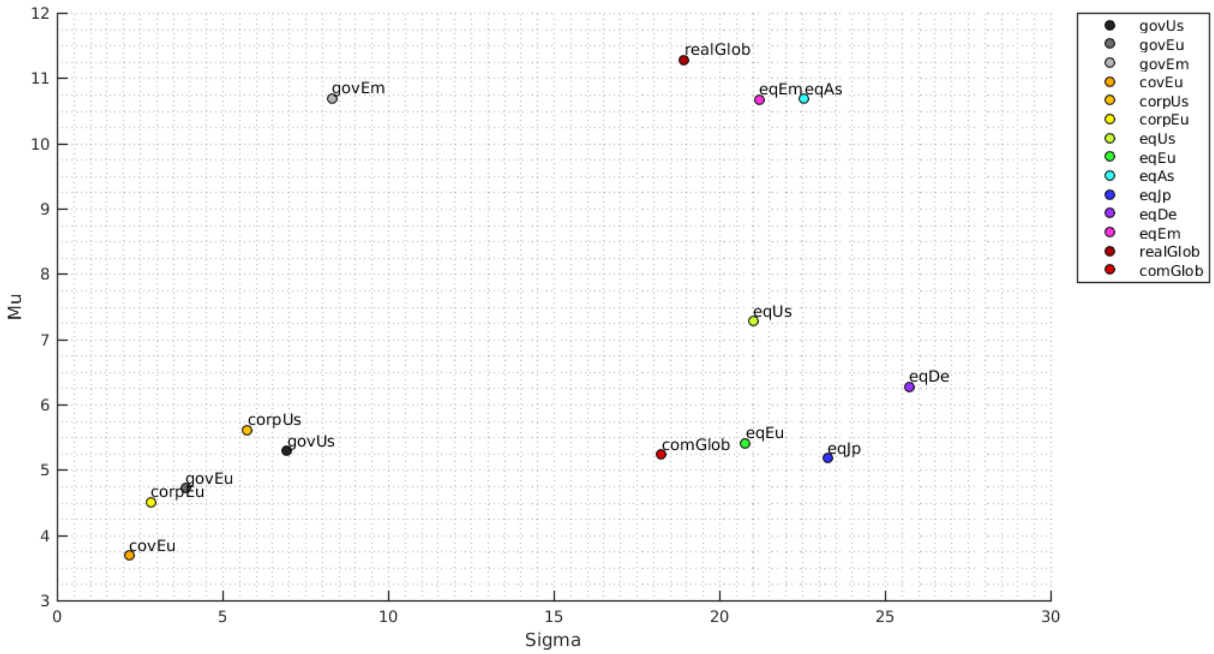

Figure 4.4: Annualized historic moments of discrete percentage returns in local currency

As can be seen, all assets basically can be split into two categories. On the left side, with lower values of volatility, are all bonds: covered bonds, government bonds and corporate bonds. And among all bonds, emerging market government bonds have the highest volatility, coinciding with usually higher underlying economical and political risks. Then, with a small gap, the more risky assets follow on the right side of the figure, with equity, real estate and commodity indices. So far, individual asset risks are as expected, and the universe assets cover a broad range of levels of risk.

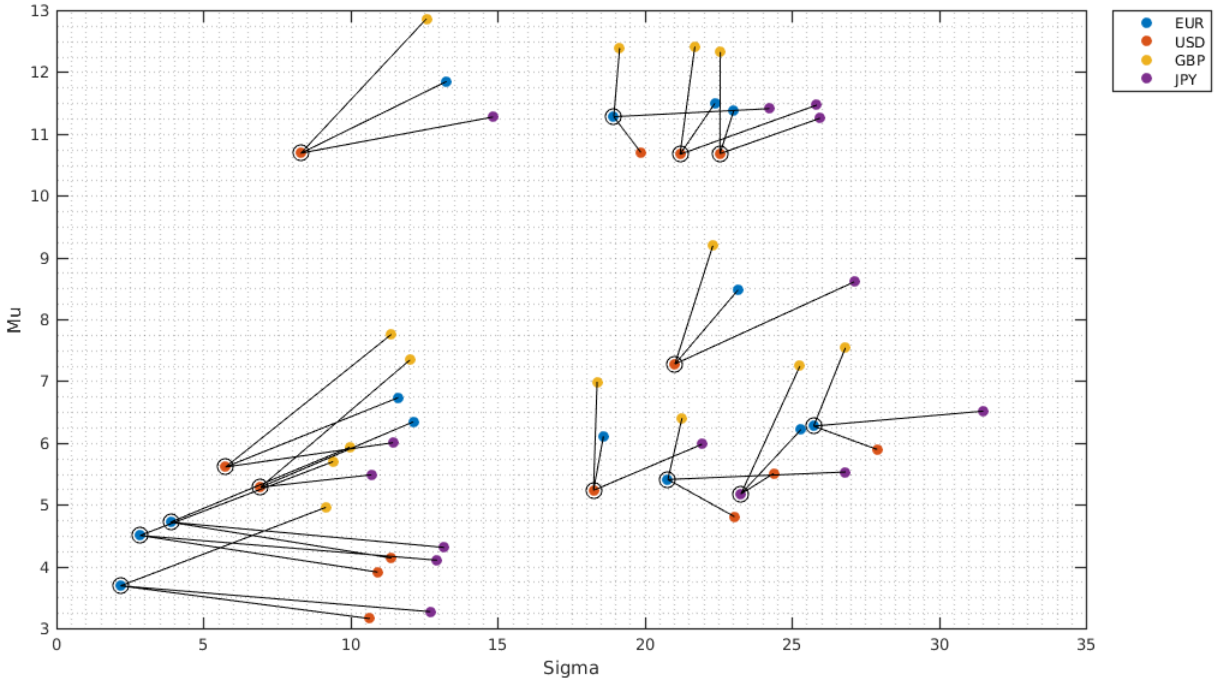

Figure 4.5: Asset return moments for different currency denominations

Next, let's examine how much estimated sample moments do change when asset returns are denominated in different currencies. For each asset, returns over time are inspected in four different currencies: EUR, USD, GBP and JPY. We again estimate asset moments and scale them to annual basis. Figure (4.5) shows estimated mean return and volatility for each asset and each of the four currencies. The original local currency moments are encapsulated by black circles, and black solid lines connect all estimated moments for a given asset. As can be seen, currency conversion away from local currency returns almost exclusively increases individual asset volatilities. This is not surprising, as exchange rate movements incorporate an additional source of risk into asset returns denominated in foreign currency. Furthermore, this effect is far more pronounced for low-risk assets (bonds), than it is for higher risk asset classes. Rather stable bond return series become quite volatile when exchange rate changes need to be taken into account. For example, EUR denominated covered, government and corporate bonds are no longer low-risk assets through the eyes of a U.S. investor, with annual volatility exceeding 10% in all three cases. In local currency, all of them had a volatility of below 5%.

In addition to changes in volatility, one also can see effects of currency conversions on expected returns. However, these changes are most likely of stochastic nature primarily, as obvious trends in exchange rates are usually not present in empirical data. It is doubtful that such patterns will continue to hold out-of-sample.

In contrast, effects on pairwise asset correlations should be meaningful and persistent. Two assets that are almost independent in local currency could easily become dependent after currency conversion, as they then share the exchange rate as a common risk factor. This can be seen in Figure (4.6), where local currency correlations are compared to correlations of asset returns in EUR. The left side of the figure shows asset correlations after currency conversion, while local currency correlations are given on the right. Assets are arranged in such a way that in local currency co-moving assets are arranged closer together. As can be seen, in a first layer the asset universe splits up into two broad blocks. While European and U.S. bonds can be found in the lower-right corner of the heatmap, all other assets are grouped in the upper-left part. Only emerging market bonds can not clearly be assigned to any of the two blocks, as they exhibit positive correlations with both. Within the two broad blocks, some smaller clusters can be found. Bonds can be sub-divided into European and U.S. bonds, and also European and German equities form a block within equities. Keeping the order of assets fixed, the left heatmap of Figure (4.6) shows how the structure of the correlation matrix is blurred when asset returns are converted to EUR. Bonds now get divided with regards to regions: European bonds form a joint block with only minor dependence to the other bonds. And emerging market bonds suddenly get closely linked to U.S. bonds, as all of them share EUR/USD exchange rates as common risk factor.

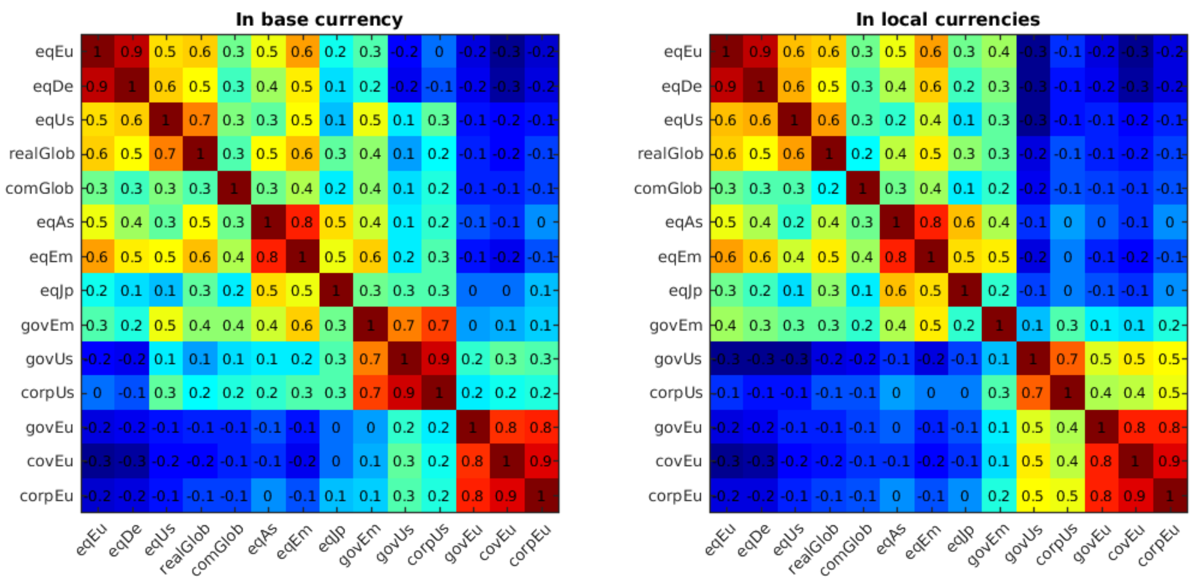

Figure 4.6: Estimated pairwise correlations for different currencies

Given these findings, exchange rate fluctuations are a non-negligible risk factor in international portfolio allocation decisions. However, the present data sample is just too short to draw reliable conclusions especially with regards to mean expected exchange rate changes. Exchange rates just have too much variability in order to precisely pin down expected values with such a small sample. We will now complete the picture by looking at findings in the literature.

(Dimson, Marsh, and Staunton 2002) have looked at financial market data of 16 countries over a 101-year period from 1900 to 2000. In particular, they also investigated historic exchange rates, for which they found that nominal exchange rates have changed up to 5.7% annually against the dollar. Over a period of 101 years this amounts to the depreciation of a given currency to roughly only 0.27% of its original worth. Certainly, the time period chosen might be considered an extreme case, with plenty of political and economical upheavals. But even when the sample is reduced to the period from 1950 to 2000, some nominal exchange rates shift substantially, with annualized changes ranging from -4.3% to 3.3%.

Let's for a moment assume a rather moderate geometric mean nominal exchange rate change of 1% annually for country A against country B. After a period of 50 years, this will still lead to an appreciation of 64% for currency of country A, while country B's currency depreciates 39% (due to $(1 - \frac{1}{1.64}) \approx 1 - 0.61 = 0.39$). Keeping all other things equal, this would mean that residents of country A can now buy 64% more goods for a given unit of currency in country B than 50 years before. Can this be?

Purchasing power measures the amount of goods and services that can be bought with a single unit of currency. In a global market, there are two ways to buy a certain good. First, the good can be purchased domestically, with homeland currency. And second, homeland currency can be converted into foreign currency, which in turn is used to buy the desired good abroad. The so-called law of one price states that in equilibrium both ways of buying the good should ultimately have the same price. Whether this purchasing power parity (PPP) holds, however, depends on the nominal exchange rate between countries. With given domestic and foreign prices P_t^H and P_t^L, there exists a unique exchange rate such that a unit of currency has the same purchasing power domestically and abroad. This rate is called the PPP exchange rate. Global exchange rates generally differ from PPP exchange rates, such that some currencies possess more purchasing power than others. One way of measuring the degree of deviation from equilibrium PPP rates is with real exchange rates:

Definition 4.4.3: *[Real exchange rates, (Mankiw 2008)]*

Let $FX_t^{H/L}$ denote the nominal exchange rate between two countries, and let P_t^H and P_t^L denote domestic and foreign prices respectively. Then the real exchange rate is defined as

$$RX_t^{H/L} = \frac{FX_t^{H/L} P_t^H}{P_t^L} \tag{4.2}$$

While nominal exchange rates measure the amount of foreign currency that can be obtained with a single unit of homeland currency, real exchange rates focus on real goods and services instead of currency. They measure the amount of goods and services that a single unit of currency can buy locally, versus the amount of goods and services that could be obtained abroad. PPP hence is given whenever real exchange rates are equal to one, and real exchange rates other than one will measure the degree to which equilibrium is not fulfilled. And, given that nominal exchange rates deviate from PPP rates, there exist two ways how equilibrium could be restored: either exchange rates need to adapt, or local prices of goods.

Using Equation (4.2), nominal exchange rates can be split into two components: one part reflecting the real exchange rate, and another part that represents differences in nominal prices:

$$\mathrm{FX}_t^{H/L} = \mathrm{RX}_t^{H/L} \frac{P_t^L}{P_t^H}$$

This way, also changes in nominal exchange rates can be decomposed:

$$\frac{\mathrm{FX}_t^{H/L}}{\mathrm{FX}_{t-1}^{H/L}} = \frac{\mathrm{RX}_t^{H/L}\frac{P_t^L}{P_t^H}}{\mathrm{RX}_{t-1}^{H/L}\frac{P_{t-1}^L}{P_{t-1}^H}} = \frac{\mathrm{RX}_t^{H/L}}{\mathrm{RX}_{t-1}^{H/L}}\frac{P_t^L}{P_{t-1}^L}\frac{P_{t-1}^H}{P_t^H} = \frac{\mathrm{RX}_t^{H/L}}{\mathrm{RX}_{t-1}^{H/L}}\frac{I_t^L}{I_t^H}$$

Thereby I_t^H and I_t^L denote changes in price levels, commonly known as inflation rates. Changes in nominal exchange rates hence can be related to two separate sources: changes in real exchanges rates and differences in inflation rates.

As it turns out, it is fundamentally different whether exchange rate changes have been caused by inflation differentials or real exchange rate movements. As we will see, any changes due to different inflation rates can generally neither be exploited nor hedged. As an example, let's assume that an investor is living in country A with inflation rate equal to 10%. Furthermore, let country B be a country with zero inflation and constant real exchange rate to country A. Due to inflation in country A, currency of country B will appreciate:

$$\frac{\mathrm{FX}_t^{H/L}}{\mathrm{FX}_{t-1}^{H/L}} = \frac{\mathrm{RX}_t^{H/L}}{\mathrm{RX}_{t-1}^{H/L}}\frac{I_t^L}{I_t^H} = 1\cdot\frac{1}{1.1} = 0.909$$

$$\Rightarrow \quad \frac{\mathrm{FX}_t^{L/H}}{\mathrm{FX}_{t-1}^{L/H}} = 1.1$$

Hence, an investor might want to benefit from the currency appreciation by holding N units of currency B. In homeland currency, the worth of N units of currency B is given by

$$W_t = N\,\mathrm{FX}_t^{L/H}$$

Due to currency movements, the worth of this money will change to the next period:

$$W_{t+1} = N\,\mathrm{FX}_{t+1}^{L/H} = N\,1.1\,\mathrm{FX}_t^{L/H} = 1.1\,W_t$$

Hence, after conversion, the investor will be left with an increase in homeland currency of 10%. However, this increase exactly equals the inflation rate, so that in real terms the investor will not be better off at all. An on paper worthwhile investment did become worthless in real terms.

A similar argument also holds for sometimes deceitfully promising expected returns in high-inflation countries. Although these returns might actually realize in foreign currency, it is simply not possible to also obtain them in domestic currency. Inflation differentials in general will cause adverse currency movements, thereby diminishing realized domestic returns.

As nominal exchange rates can be deceitful, investment decisions involving foreign currency better should be judged in real terms: what is the real return in domestic currency, and how will the real exchange rate change. In the short run, prices are sticky, so that nominal and real exchange rate changes are equal. In the long run, however, they can differ widely due to inflation differentials. In fact, a major part of nominal exchange rate changes can be explained through differences in inflation rates (Dimson, Marsh, and Staunton 2002). Real exchange rates themselves are comparatively stable in the long run. For example, (Dimson, Marsh, and Staunton 2002) find that the most extreme real exchange rate changes against the USD have been South Africa, with a geometric mean of -1.3% over the full 101 years, and Australia with -0.6%. Hence, real returns of equities in all countries considered are almost the same for local or U.S. investors over the full sample.

Still, real exchange rates fluctuate significantly in the short run, with annual standard deviations in the range of equity markets these days. This is mainly a result of countries in the world following a regime of floating exchange rates nowadays, which led to increasing volatility of real exchange rates during the last decades (Perold and Schulman 1988), (Catao 2007), (Dimson, Marsh, and Staunton 2002).

Even though exchange rates themselves have a high level of volatility, this volatility only partially adds to asset price volatility of foreign assets. For example, (Dimson, Marsh, and Staunton 2002) find that although the UK/USD exchange rate had an annual volatility of 12% over the period from 1900 to 2000, it only increases asset price volatility of UK assets from 20% to 23% when converted to USD. The reason for this is that UK/USD exchange rates are not strongly correlated with UK stock prices, such that a major part of the additional volatility of exchange rates gets diversified away. Similarly to the results of Figure (4.5), however, exchange rate flucuations comparatively add much more volatility for low-risk asset classes like bonds than for equities (Perold and Schulman 1988).

Summing up, we have seen that exchange rates can influence investments quite substantially. Currency conversion can change risk-return profiles of individual assets, correlations between assets and add an additional source of uncertainty. Investors might find it desirable to eliminate this additional source of risk, and they can do so by hedging against currency movements through forward contracts. A forward contract between two counterparties determines an exchange of assets at a future point in time. With this, investors can already today determine an exchange rate for which they will exchange a fixed amount of currency at some point in the future. Forward contracts are set up in such a way that the initial value of the contract is zero, and thus no immediate cash-flow occurs. The future rate at which currencies will be exchanged is the variable quantity that needs to be specified such that the initial value of the contract is zero. It's fair value can be derived by no-arbitrage arguments.

Definition 4.4.4: *[Forward rate, (Hull 2009)]*

Let i_H and i_L denote the risk-free interest rates with maturity $T-t$ in homeland and foreign currency respectively. Let further $FX_t^{H/L}$ denote the spot exchange rate between both currencies. Then the forward rate between both currencies is given by

$$FWD_t^{H/L}(T) = FX_t^{H/L}\frac{(1+i_L)}{(1+i_H)}$$

This value can be derived from the assumption that two investments with exactly equal risks and payoffs also need to have the same price. In this case, there are two ways of investing money into risk-free assets, and both need to have the same compounding. The first one is by simply investing money into domestic risk-free assets at domestic risk-free interest rates. The second one is by exchanging homeland currency into foreign currency, investing it into risk-free assets abroad, and converting the money back to domestic currency. Through the use of forward contracts the future exchange rate at which foreign currency is converted back into domestic terms can already be determined today. This way there is no uncertainty left, and the final payoff of a risk-free investment abroad is already known in advance. With given domestic and foreign risk-free rates, it is the forward rate that needs to be set such that both investments are equal. When it is set accordingly, one generally speaks of covered interest rate parity.

Since forward contracts do lock in future exchange rates, they effectively eliminate uncertainty. This makes them a well-suited instrument to hedge against currency risks. Another desirable property is that forwards do not require any cash-flow upfront, as their initial value is always equal to zero. This, however, does not mean that it is the current spot rate that gets locked in for future transactions. Quite the opposite is true: whenever interest rates in both countries differ, the locked-in forward rate will differ from the current spot rate. Hence, even when future spot rates do not deviate from current spot rates, the forward position still will close with non-zero value. Only when future spot rates equal the forward rate, no settlement payment needs to be done. Depending on whether the forward rate is higher or lower than current spot rates, the final return of a hedged foreign investment will also be higher or lower than what the foreign investment achieves in local currency. The local currency return in a foreign country is something that is generally not achievable for investors of other countries, and the difference of returns is called hedge return (Charles and Bosse 2014).

So what has to be expected for hedge returns? Do they diminish expected asset returns? Hedge returns are ultimately driven by whether forward rates sell at a discount or at a premium. When hedging currencies, one mostly uses forwards with maturities in the range of one or several months, in order to be able to constantly adjust the amount of currency hedged to the usually fluctuating amount of foreign exposure. Hence, forward rates themselves are driven by interest rate differentials on the short end of the yield curve, which are mainly subject to central banks' inflation targeting. In the long run, this will adjust returns of foreign investments similar to what unhedged investors would experience due to inflation differentials (Charles and Bosse 2014). This would mean that hedging basically eliminates the risk of changes in real exchange rates only. Nominal exchange rate changes that are driven by inflation differentials are basically inevitable for investors in the long run, even with hedging. This will guarantee that foreign investors get real returns similar to domestic investors. And, in particular, for bonds this will mean that international bond returns will more closely align to domestic bond returns.

Even when nominal exchange rates can not be hedged fully, hedging against real exchange rate changes might already be worthwhile. A popular view is that expected returns of foreign investments are not affected, while risk can be reduced substantially (Perold and Schulman 1988). This way risk-adjusted returns can be improved, making hedging a promising approach.

Still, however, one would have to decide on the best way of how hedging should be employed (Black 1995). The easiest approach might be to link the hedging decision to individual assets, and either hedge an asset fully or not at all. A more sophisticated approach would also allow partial hedges: due to possible diversification effects between local returns and exchange rate changes, the variance minimizing hedge ratio might differ from a full hedge. In particular, optimal hedge ratios of bonds and equities should be quite different. In general, however, there is no reason why hedging decisions should be made on an asset-wise level: eventually, the quantity of interest is risk on a portfolio level. The global best solution hence is to simultaneously determine optimal portfolio weights and optimal hedge ratio.

Although all of these arguments sound convincing, there also exist arguments against hedging. A first such argument is that hedging actually might reduce expected returns due to Siegel's paradox (Sinn 1989). Although exchange rates are usually perceived as symmetric risk that has an expected return of zero, the distribution of currency returns for any given investor is actually asymmetric. This best can be seen by an example.

Example [Siegel's paradox]

Let's assume for simplicity that the EUR/USD exchange rate is equal to one:

$$\mathrm{FX}_t^{\mathrm{EUR/USD}} = 1$$

Let's also assume that exchange rates are symmetric in the following sense: the likelihood of a EUR appreciation of $x\%$ is equal to the likelihood of a USD appreciation of equal amount. For simplicity, we only assume a single possible value for x, $x = 30\%$, and both EUR and USD appreciation occur with probability of 50%.

Then, from a European perspective, the exchange rate in case of a EUR appreciation will be $\mathrm{FX}_t^{\mathrm{EUR/USD}} = 1.3$. In the case of a USD appreciation, however, the exchange rate will be

$$\mathrm{FX}_t^{\mathrm{EUR/USD}} = \frac{1}{1.3} = 0.769$$

Hence, for the expected value we get

$$\mathbb{E}[\mathrm{FX}_{t+1}^{\mathrm{EUR/USD}}] = 0.5\,[1.3 + 0.769] \approx 1.035 > 1$$

Although only one side can profit from foreign currency exposure at a time, in terms of expected returns both sides are simultaneously better off. Hence, fluctuations in exchange rates themselves are a cause of positive expected returns, such that investors of both sides will benefit. Still, the magnitude of this effect is rather limited, and positive expected returns are hard to prove empirically (Sinn 1989).

There is another point that can be made in favor of foreign exposure, but it depends on the economic importance of the actual currencies involved. (Bundesbank 2010) finds that some currencies play a special role for global investors during times of global financial turmoil.

They are perceived as “safe haven”: liquid and low-risk when risks are otherwise unclear. For example, U.S. dollar and Swiss franc are found to be such currencies, as they tend to appreciate whenever uncertainty in financial markets is heightened. This makes foreign exchange exposure to the U.S. and Switzerland particularly favorable, as these positions tend to appreciate just at a time when it is needed most. This way losses in individual positions can be somewhat absorbed, providing a cushion against major drawdowns. Hence, from the point of view of a European investor, U.S. dollar exposure might be desirable.

All arguments in favor of hedging were basically based on one fundamental assumption: the portfolio value in terms of homeland currency is the relevant quantity to measure expected returns and risk. For any given level of expected return, fluctuations of the nominal portfolio value should be decreased, and hedging might very well contribute to this goal. However, money is just a means to an end, and ultimately investors want to consume goods and services. The more crucial aspect hence is: how much does the amount of goods and services fluctuate that an investor could buy from the value of his portfolio? Or, in other words: it is not the nominal value of a portfolio that should count, but the purchasing power that goes along with it. And purchasing power should be measured not only nationally, but with regards to goods and services provided in foreign currency as well. One possible way of measuring fluctuations of global purchasing power is through real effective exchange rates. Nominal effective exchange rates are calculated as nominal value of one unit of currency compared to a basket of foreign currencies. Related to that, PPP conversion factors can be used to measure how a basket of goods and services bought domestically compares to a similar basket of goods bought abroad. The idea is that a basket of domestic goods and services can be exchanged into domestic currency, which in turn can be converted into some basket of international currencies, which then could be used to buy some fraction or multiple of similar baskets of goods and services abroad. Building on this idea, real effective exchange rates measure how domestic purchasing power compares to global purchasing power over time. When homeland currency appreciates in terms of real effective exchange rates, vacation abroad gets comparatively cheaper, as well as imports of tradeable goods. This makes investors better off, even when the portfolio value itself remains unchanged in nominal terms. Hence, even when a portfolio is fully hedged against any nominal currency exposure, investors still are exposed to fluctuations in their global purchasing power. In some sense, there is currency risk even when there is no nominal currency exposure.

For example, UK citizens did experience this phenomenon in the aftermath of the Brexit referendum in June 2016. With the decision to leave the European Union the British pound immediately did depreciate against almost all foreign currencies. With only domestic cash on some bank account, a UK citizen would not have been affected in terms of the nominal value of his account. However, he actually would have lost quite substantially in terms of real purchasing power due the change in exchange rates. On the other hand, a UK citizen with unhedged foreign assets in his portfolio would have experienced a different picture. Due to the depreciation of GBP, keeping all other things fixed all foreign positions would have increased in value. In reality, foreign positions did not remain constant, but did mostly drop in local terms. Still, the exchange rate crash was dramatic enough to more than offset these losses, such that portfolios with sufficiently high foreign exposure did increase in GBP. This, in turn, did dampen the effect of diminished global purchasing power. Hence, foreign exposure can be a safety net against dramatic depreciation of local currency, which usually occurs when a country has to suffer from political or economical problems. In other words: foreign exposure can be beneficial exactly in those times where it is required most.

To get a more general picture on this perspective, Figure (4.7) shows real effective exchange rates over time for four different currencies. Real effective exchange rate data is put together by Bank for International Settlements and was retrieved from FRED (Federal Reserve Bank of St. Louis). From the figure one can see that there have been substantial shifts in global purchasing power for different currencies. For example, global purchasing power did decrease approximately 20% for UK citizens in 2016, and 40% for Japanese citizens over the full period from 1994 to mid 2017. Without any foreign exposure, this fully amounts to the changes in global purchasing power that citizens would have gotten. Through investments with foreign currency exposure, however, one can diversify across purchasing powers of multiple currencies, thereby reducing the risk of sudden negative developments. This can be seen by the green line, which is a simple average of all four effective exchange rates.

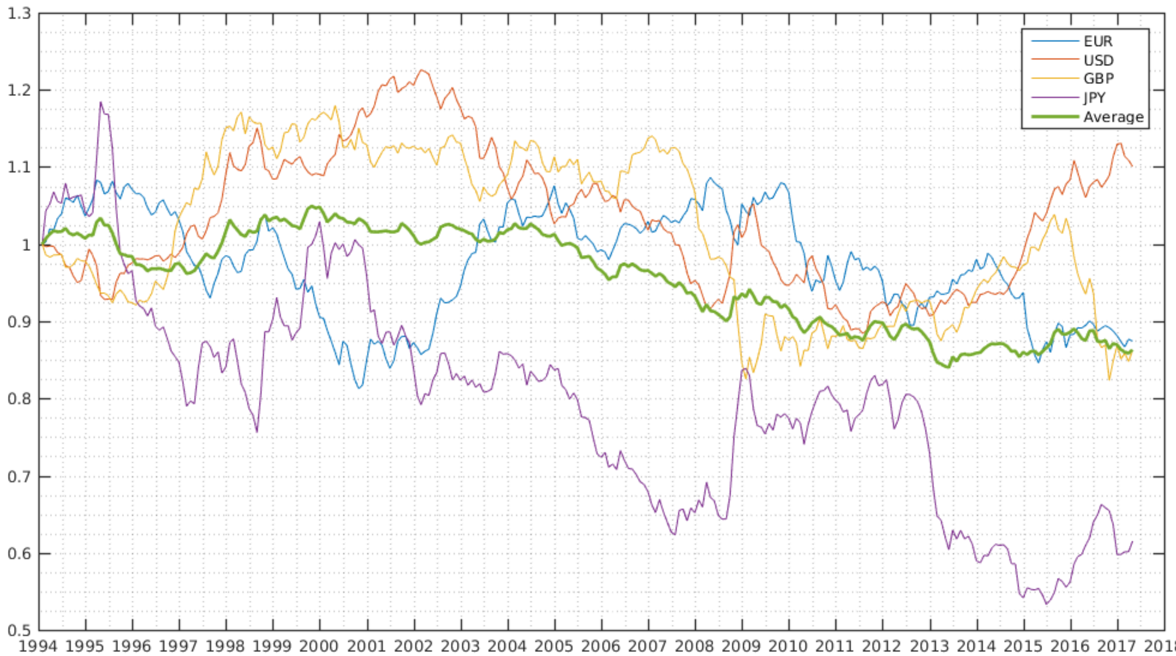

Figure 4.7: Real effective exchange rates for selected regions

These diversification effects that can be achieved through exposure to changes in global purchasing power of multiple countries, however, can easily be overlooked when focusing on nominal portfolio values only. Let's just emphasize this with a simple example. Let's say a European investor has 10000 EUR on a no interest-bearing bank account in year 2000. Without any changes to this investment, the nominal of this portfolio will be constant over time, and the investor will still have exactly 10000 EUR on his bank account in mid 2017. Let's now assume that the investor splits his wealth equally among a European and a U.S. bank account, each in local currency. Through the currency exposure introduced by holding USD, the nominal overall wealth in EUR is not constant anymore, but fluctuates with the

EUR/USD exchange rate. This is shown as blue line in Figure (4.8), where the nominal portfolio value drops to almost 8000 EUR in 2008. On first sight, hence, currency exposure seems to be undesirable, as it just adds a source of risk that is not compensated (apart from Siegel's paradox at least). Comparing the EUR only and EUR/USD cash portfolios in terms of global purchasing power, however, the picture changes. The EUR only portfolio now suddenly fluctuates quite substantially when real effective exchange rate changes are taken into account, which is shown as yellow line. Compared to that, the EUR/USD cash portfolio, which is diversified across two currencies, fluctuates less in terms of global purchasing power. It is shown in red.

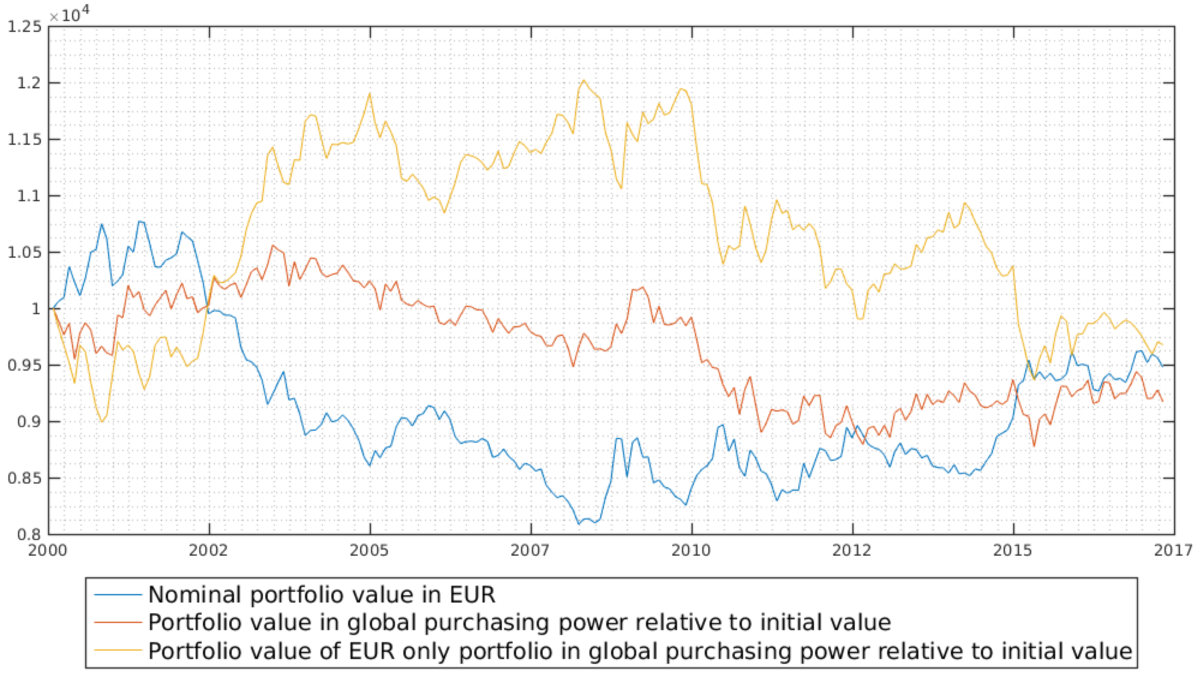

Figure 4.8: Portfolio value changes measured in nominal and global purchasing power

Summing up, considering the currency risk of either a nominal or a purchasing power perspective will generally lead to different conclusions with regards to hedging. So the question is: which of both measures should be the quantity of interest? There is no general answer to this question, however, as it also depends on the investor's consumption behavior. If an investor spends almost all of his wealth on either non-tradeable goods or domestic services (like going to the hairdresser), nominal values should be of primary interest to him. If, however, he spends a huge amount of wealth on traveling, global purchasing power could become the more relevant quantity. Quite generally, however, one can say that a sole focus on nominal terms only gives an incomplete picture. In addition, the only difference between hedged and unhedged returns should primarily lie in changes of real exchange rates in the long run, and real exchange rates have proven to be rather stable over periods of multiple decades. Given that we have a long-term investment horizon in mind, we will hence work with unhedged foreign assets in the following.

4.5 Results

4.5.1 Universe data

As we decided to forego any currency hedging, the first step is to convert all assets into Euro prices first. These Euro price trajectories are shown in the left part of Figure (4.9). In the right part, we can see estimated annualized sample moments for each asset's historic return series. Moments are estimated based on discrete returns, and scaled to annualized moments according to the formulas provided in Section (1.2.3). As can be seen from estimated volatilities, assets can be divided into roughly three risk classes. The first group consists of European bonds, which form the group of assets with lowest risk. Then, with slightly more risk, all other bonds form a second group. Thereby U.S. government and corporate bonds have elevated risks as compared to European ones mainly due to the additional exchange rate risk that is introduced when prices have to be converted into Euro. Exchange rate risk is also part of the explanation for the elevated annualized volatility of emerging market government bonds. The index comprises debt issued by more than 30 different emerging market countries. However, these bonds are not denominated in local currency, but they are all denominated in U.S. dollar. Hence, there is no diversification between multiple currencies, but bond prices are fully exposed to changes in EUR/USD exchange rates. In addition to exchange rate risk, emerging market bonds also have higher levels of credit risk that can further contribute to volatility. Last but not least, all equities form the group of high-risk assets together with real estate and commodities.

Over the period of roughly 18 years, government emerging market bonds did achieve the highest performance, with almost 12% annualized return. Each single Euro invested into this index would have resulted in roughly 6.5 Euro after the end of the period. Only slightly worse, equity emerging markets, equity Asia and real estate indices all have annualized returns between 11% and 12%. In contrast to that, equity markets of developed countries did achieve inferior performances, with roughly 8.5% annually for U.S. equities and annualized returns between 5% and 7% for European, German and Japanese equity indices.

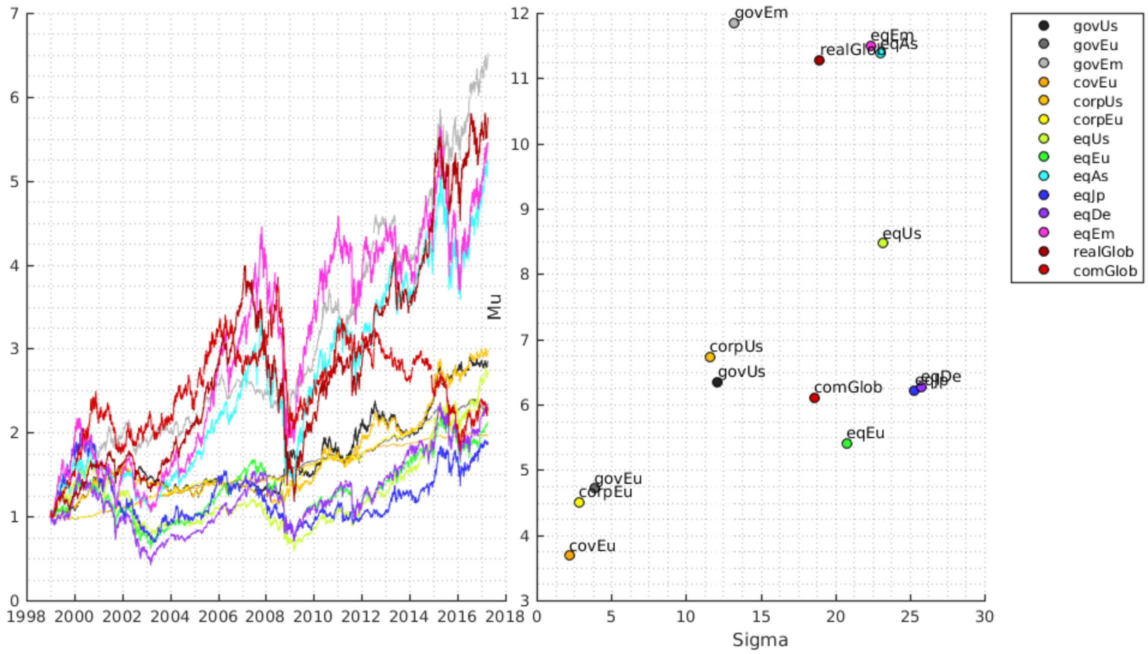

Figure 4.9: Normalized price trajectories (EUR) and estimated annualized sample moments

Due to globally decreasing interest rates, all bond indices did achieve remarkably good returns over the full period. U.S. government and corporate bond indices achieved the highest performances with values above 6%, as they also get an additional boost by a slightly decreasing EUR/USD exchange rate. European government and corporate bonds still achieved more than 4%, while covered bonds did grow less than 4% annually as minimum risk asset.

A different way of measuring risk is shown in Figure (4.10), where maximum drawdowns are shown for individual assets. Assets again can be categorized in the same three groups of risk. Maximum drawdowns of European bonds are below 10%, while they are between 20% and 30% for all other bonds. In contrast, all high-risk assets did suffer maximum drawdowns above 50%, with German equity even losing more than 70%.

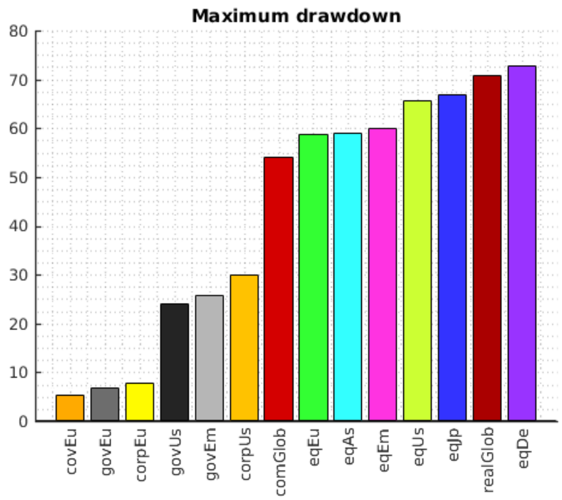

Figure 4.10: Maximum drawdowns

Maximum drawdowns can easily build up over durations of one or more years. Hence, a well suited dynamic asset management strategy should be able to circumvent some of these losses by rebalancing into other assets eventually. The most difficult situations are when losses occur sudden, because it is hard to identify major drawdowns on short notice. Hence, to get a better feeling for risks on medium time scales, Figure (4.11) shows yearly returns for all assets. As can be seen, losses of a single calendar year can already be devastating, with negative returns of more than 30% for all high-risk assets except Japanese equities in 2008. This also underlines the strong links that exist between individual markets. The global financial crisis of 2008 did negatively affect all high-risk assets simultaneously, such that pure equity or high-risk portfolios did not really benefit from global diversification. Only bonds did provide the opportunity to circumvent losses, with U.S. government bonds achieving the single best yearly return in 2008 with a return of 25.3%. Besides all the risks that come with high-risk assets, of course, there also exists huge upward potential. For example, Japanese and emerging market equities did achieve returns of more than 80% in a single year in 1999. This was an exceptionally good year for almost all assets, with returns larger than 25% for all high-risk assets, and only European bonds incurring losses of 2%-3%. Keep in mind, however, that this exceptionally good year will not be part of the backtest period. Backtests of strategies only start in June 2000, as we need a sufficiently long history to get a first reliable estimate for individual asset moments. Hence, our backtest period will directly start with the rough years related to a bursting dot-com bubble.

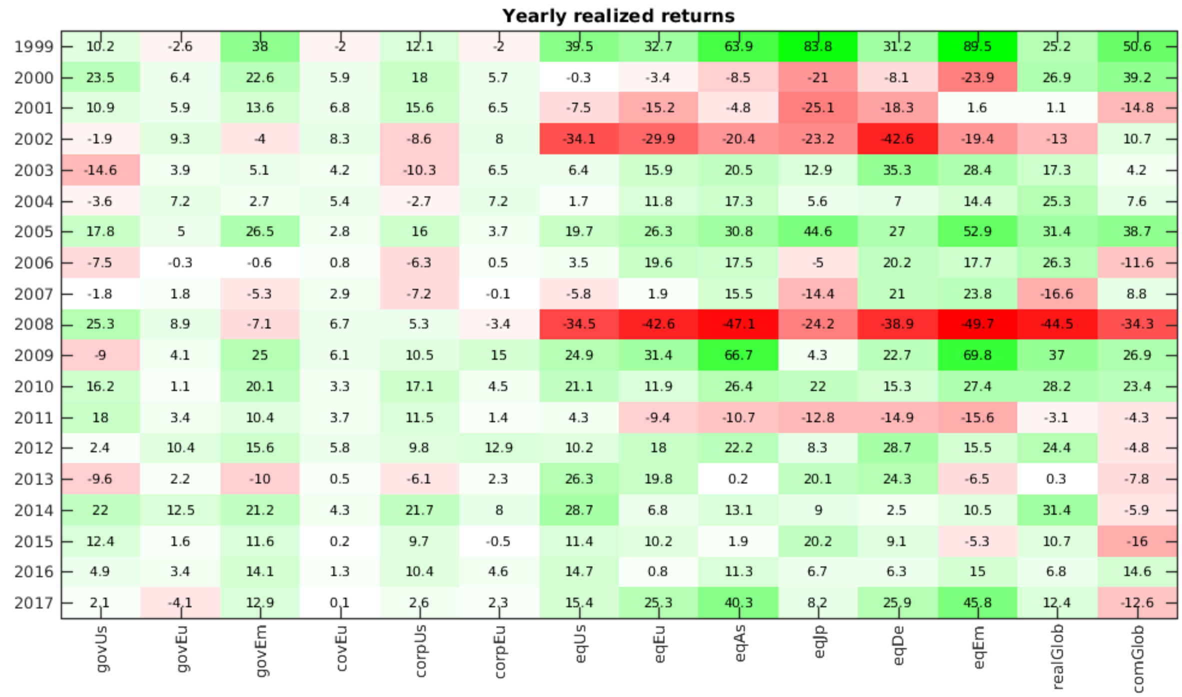

Figure 4.11: Yearly realized returns of individual universe assets

4.5.2 Strategy results

We now will look at backtest results of multiple dynamic trading strategies that we have tested. When setting up strategies to be tested, we used the modular structure of strategies that we defined in Chapters (2) and (3). Strategies hence can be composed of several components:

- a single-period strategy
- endurance filters
- cost-reduction filters

As single-period strategies we will test diversification-aware portfolio selection strategies for two different levels of diversification targets: 0.6 and 0.8. As benchmark, we will also backtest Markowitz mean-variance portfolio selection with short-selling constraints. In order to keep trading volumes and hence also trading costs within reasonable bounds, each single-period strategy can be adapted with possibly several endurance or cost-reduction filters. A full list of all dynamic strategies tested can be found in Section (6.0.3) in the Appendix.

As pointed out in Section (3.1), we do not test a single risk target per strategy only, but we do set up a full ramp of different risk categories instead. This ramp is made up of 15 daily σ-targets, which we obtain through the following simple definition:

$$\sigma_i = \frac{i}{100\sqrt{250}},\ i = 1, \ldots, 15$$

This way individual σ-categories can roughly be interpreted as representing levels of annual volatility: category i should very roughly have annual percentage volatility equal to i. Note, however, that this is not at all an explicit target for individual categories, and their respective levels of true annual volatility will generally deviate from these specifications. Here, we basically use these specified σ-targets as a simple way to span some grid of categories with increasing levels of risk. Once we really want to target a specific level of annual volatility, we could simply take that category which is closest to the target in terms of realized annualized volatility. In other words, daily category targets σ_i can be seen as calibration parameters, and simply testing enough of them will lead to some strategy that has the desired true annual volatility level. Of course, this only works for annual volatility levels that are within a reasonable range and hence also can be obtained.

As the main motivation for any kind of weight filters is the reduction of trading costs, we now need to define some way of measuring the amount of trading. We will use the following measure:

Definition 4.5.1: ***[Turnover]***

Let TV_{t_i} denote the volume associated with rebalancing at time t_i. Let further AUM_{t_i} denote the assets under management immediately before the execution of rebalancing. Then we define ***turnover*** *related to n trading events by*

$$TO = \sum_{i=1}^{n} \frac{1}{2} \frac{TV_{t_i}}{AUM_{t_i}}$$

In our backtest setting we do not have multiple clients per category, but only look at backtest results for a single representative client. Hence, AUM_{t_i} simply equals the portfolio value at time t_i. Factor $\frac{1}{2}$ normalizes trading such that a full sale and re-investment of the portfolio will map to turnover equal to one. A full rebalancing of the portfolio will require trading volume equal to AUM during sales, and the same trading volume again when re-investing the money again. Hence, trading volume is two times AUM, and turnover will be equal to one.

Now that we have a way to measure the amount of trading, we still need to define a way to map turnover to trading costs. Therefore, we will make the following assumptions on the amount of direct and indirect transaction costs:

- variable transaction costs: $c_v = 0.1$ %
- average bid-ask spread: $c_{ba} = 0.2$ %
- fees custodian: $c_f = 0.05$ %

Variable transaction costs and indirect costs from bid-ask spreads are proportionally to trading volume. Note that we only calculate with half of bid-ask spreads, as we only lose the difference to mid prices on average. Fees to be paid to the custodian usually vary with regards to assets under management, i.e. with regards to portfolio values for our case. Altogether, with overall annual trading volume TV, annual trading costs will be given in absolute terms by

$$\mathrm{TC}_{abs} = \mathrm{TV}(0.5c_{ba} + c_v) + c_f \mathrm{AUM}$$

In relative terms, trading costs hence will reduce annual gross returns by an amount equal to

$$\mathrm{TC}_{rel} = \frac{\mathrm{TC}_{abs}}{\mathrm{AUM}} = \mathrm{TO}(0.5c_{ba} + c_v) + c_f$$

Now we are finally able to calculate net returns for all tested strategies, which is the primary yardstick that we will use for the comparison of performances. However, returns should not be evaluated without also taking into account the risks that were taken, and hence our objective will be large risk-adjusted returns net of trading costs. In addition to that, however, we also take trading volume itself into account, as we do not want to select strategies that try to exploit investment opportunities too aggressively. This should hopefully further reduce the risk of being surprised with substantially different results out-of-sample. Hence, of all strategies that we tested, we will directly exclude all strategies from further examination where at least one single category did exceed an average annual turnover of four. For all remaining strategies Figure (4.12) shows risk-return profiles, with gross returns shown on the left, and net returns shown on the right. Thereby each single line consists of 15 points that represent the full ramp of categories. Yellow lines represent strategies that use Markowitz with short-selling constraints for single-period portfolio selection, while blue and red lines represent diversification-aware strategies with diversification levels of 0.6 and 0.8 respectively. The thick black line represents the dynamic strategy that we will further examine in the following section.

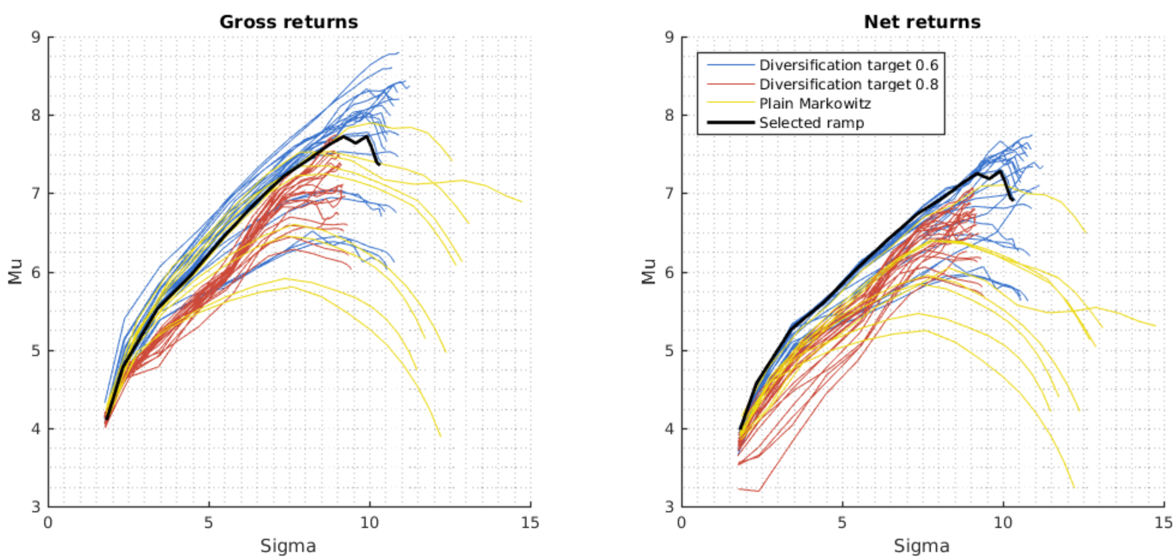

Figure 4.12: Realized annualized returns and volatilities for different dynamic strategies

As can be seen, the best strategies in terms of gross returns are built on diversification-aware target weights with diversification level 0.6. Next, there are some dynamic strategies based on Markowitz target weights that achieve almost similar results. However, all Markowitz curves tend to point downwards from medium risk categories onwards, such that higher levels of risk are not compensated with higher returns in these categories. Results for the best diversification-aware target weights with diversification level 0.8 are slightly inferior to both other approaches. However, individual ramps are generally pointing upwards, and almost all curves are rather close together, indicating even better robustness of results. Even though all dynamic strategies use the same values for daily σ targets, actually realized levels of volatility differ quite substantially. Strategies based on Markowitz generally have highest levels of volatility, followed by diversification level 0.6 and diversification level 0.8. Hence, with increasing constraints on diversification realized volatility decreases. This should not mean that additional diversification decreases volatility for any given level of expected return, which would mean superior risk-return profiles. It only is another indicator that σ targets should not be understood as something that we actually want to achieve on an annual basis, but that targets merely represent calibration parameters, and that these parameters even have slightly different influence for different single-period strategies.

Looking at the right plot in Figure (4.12) we can see risk-return profiles when returns are measured net of transaction costs. Trading strategies hence get shifted downwards in return space, with larger shifts for trading strategies with higher trading frequencies and higher levels of turnover. This eliminates much of the superiority of the best strategies based on diversification level 0.6, although they still come out as favored strategies after trading costs. From this group of best strategies we also pick our most favored one in order to analyze it in more detail in the following sections. The strategy that we choose is shown as thick black line, and it is very close to the uppermost border of all strategy results. The reason why we did not simply pick the highest risk-return ramp is that we again want to favor strategies with less aggressive trading behavior and lower overall turnover.

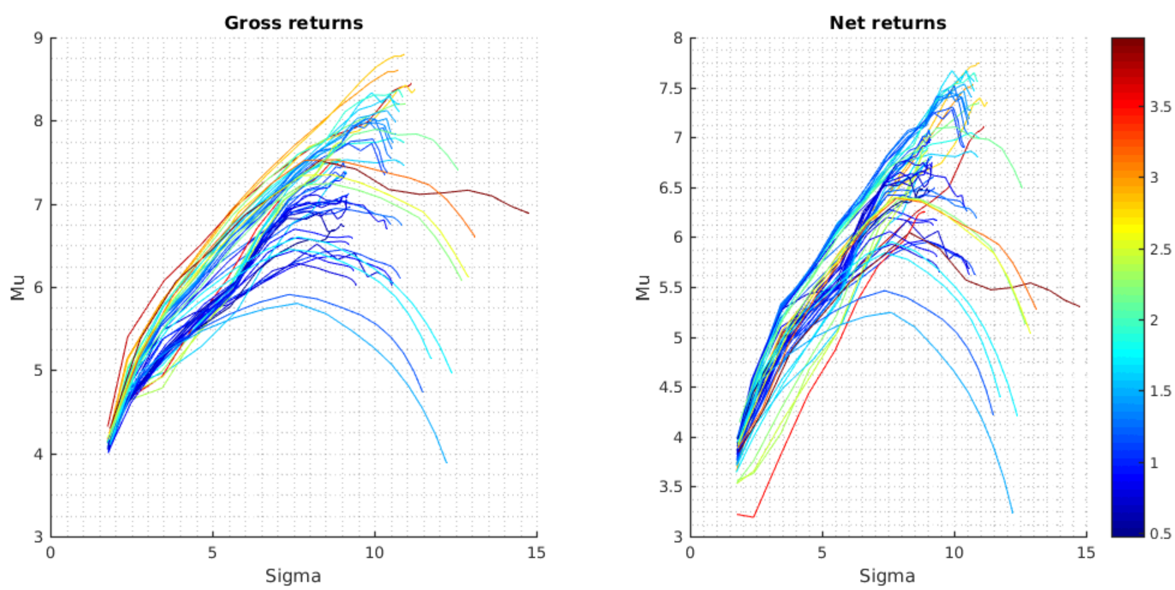

Figure 4.13: Realized maximum annual category turnover

Figure (4.13) further illustrates this choice of a favored strategy, as it sheds light on the annual turnover that is associated with each strategy. For each ramp of risk-return profiles the respective line is colorized with regards to the maximum annual turnover that a single category within this ramp has. As can be seen, maximum annual turnover values are in a range of 1.5 and higher for strategies with best risk-return profiles, but below 1.2 for the favored strategy. Again, higher trading costs are already reflected in returns net of trading costs, but deliberately selecting a more cautious strategy hopefully increases the trustworthiness of backtest results. Simply picking the best performing overall strategy is likely to introduce selection bias and overestimate backtest performances.

4.5.3 Analysis of single selected strategy

From now on we will focus on a single strategy which we want to analyze in more detail. The strategy that we pick consists of the following components:

- single-period portfolio selection based on diversification-aware portfolio selection with target diversification level of 0.6
- moving average persistence filter with window size equal to 40 days
- minimum relative tracking error filter with threshold 0.22
- cost-aware capped relative tracking error with cap at 0.18

Hence, the full strategy is described by

$$\left(\text{Div. aware } (0.6),\, h^{\text{MA}}(40),\, h^{\text{RTEmin}}(0.22),\, h^{\text{RTEcap}}(0.18)\right)$$

From now on we will refer to this strategy as selected strategy. The following table shows some key metrics for the respective backtest results. Thereby μ denotes annualized percentage returns, which are given both gross and net of transaction costs. TO denotes annualized turnover, σ annualized volatility of discrete percentage returns and Max. DD the maximum percentage drawdown that did occur over the full backtest period.

Category	μ (gross)	μ (net)	TO	σ	Max. DD
1	4.10	3.97	0.17	1.82	3.25
2	4.77	4.57	0.38	2.33	4.88
3	5.52	5.25	0.53	3.40	7.26
4	5.94	5.61	0.72	4.44	7.97
5	6.43	6.04	0.87	5.46	9.66
6	6.87	6.42	1.01	6.45	11.84
7	7.22	6.73	1.08	7.34	13.79
8	7.43	6.94	1.11	8.14	15.27
9	7.61	7.13	1.09	8.75	16.68
10	7.72	7.25	1.05	9.19	16.64
11	7.63	7.18	1.01	9.57	16.74
12	7.72	7.27	0.99	9.91	16.83
13	7.59	7.15	0.99	10.05	17.40
14	7.39	6.93	1.00	10.23	18.53
15	7.35	6.90	1.00	10.33	18.69

As can be seen, lower risk categories generally have lower levels of turnover. From risk category six and upwards all categories have annual turnover approximately equal to one, with some categories' turnover even larger than one. Compared to that, risk category one has only turnover of 0.17, indicating largely reduced trading behavior compared to higher risk categories. One reason for this is that such low-risk categories effectively have less diversified portfolios, as they select from a reduced set of universe assets only. If categories' σ-targets are low enough, most assets are not eligible in terms of risk, as they would introduce more volatility than what is allowed. This effectively leads to less diversified portfolios and, in turn, to lower turnover. Annualized turnover levels then can be translated into trading costs that reduce gross performances to net performances. As can be seen from columns μ (gross) and μ (net), the difference of annualized returns can amount up to approximately 0.5%. Due to lower levels of turnover, the reduction of net returns is smaller again for lower risk categories, and almost negligible for risk categories one and two. Finally, net annualized expected returns and volatilities are the relevant risk-return profiles that the strategy produces, and they can be used to evaluate success of the used dynamic trading strategies.

A different way of evaluation of success can be seen in Figure (4.14), which allows a relative assessment of strategy outcomes with regards to risk-return profiles of all universe assets. Each individual asset's risk-return profile is indicated as black cross, and these individual risk-return profiles basically determine the room for manoeuvre of the dynamic trading strategy. While achieved performances in absolute terms can be subject to selection bias from the actual selection of the universe, comparison to universe assets allows a more genuine assessment of the outcome. In other words, if a dynamic strategy is applied to assets with negative annualized returns only, it is not surprising when the strategy itself also will not generate overly good returns. In contrast, if the individual assets already all achieve particularly well risk-adjusted returns, basically any trading strategy will achieve good results. A more meaningful evaluation of any dynamic strategy hence corrects for this effect. As can be seen, both gross and net outcomes are in the upper left part of available universe assets, indicating a rather satisfying result. This is exactly the region that we would pick with Markowitz when all asset moments were perfectly known and stable over time. Of course, with time-varying asset moments one could further max out these limits in theory. When asset moments are not known, however, it is far from guaranteed that a strategy is able to exploit these time-varying investment opportunities appropriately out-of-sample. Looking at the results in Figure (4.14), however, the set of categories seems to do a satisfying job: through dynamic re-balancing we manage to create new risk-return profiles that not yet have been available at the market, and these outcomes are not dominated by any of the already existing assets. In particular, they complete existing financial market opportunities in the rather sparsely populated domain of medium risk assets.

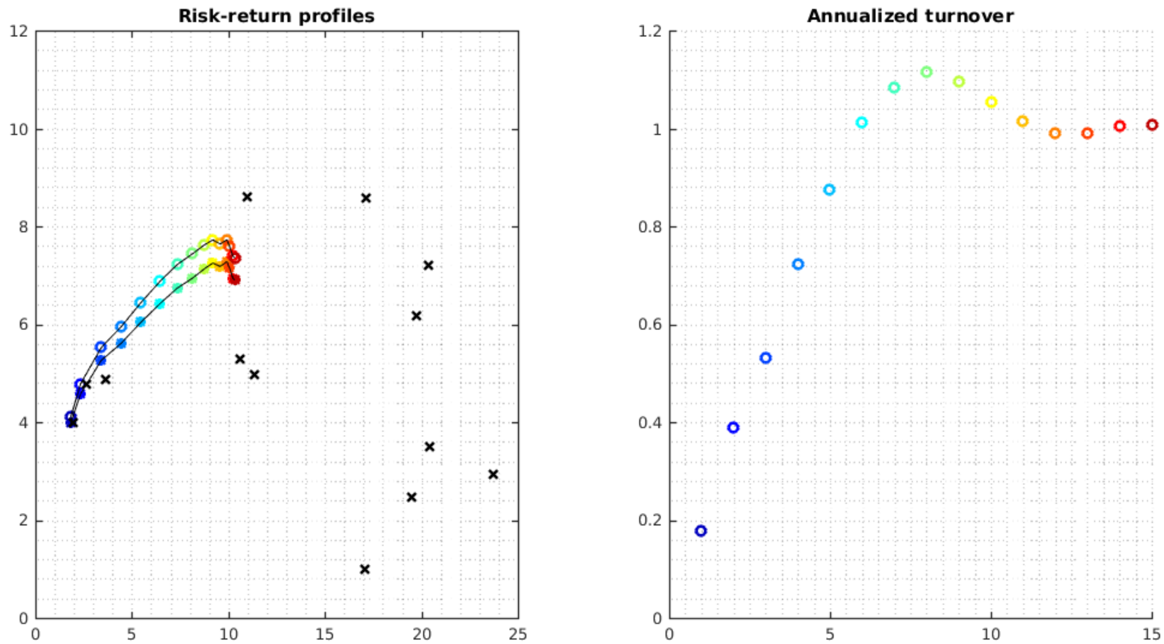

Figure 4.14: Annualized gross and net risk-return profiles and category turnover

The only medium risk assets part of this universe are U.S. and emerging market bonds. Thereby U.S. bonds have higher risk as compared to European ones mainly due to the additional exchange rate risk that a European investor has to face. In this particular sample, exchange rate risk was not significantly compensated in terms of higher returns, and even in the long run it is at least questionable whether exchange rate risk is adequately compensated by financial markets. Hence, bonds denoted in foreign currency might not be a particularly good investment choice in terms of risk-return profiles when they are considered for themselves. Only when they are also subject to other risks that are generally compensated by financial markets, like credit risk for the case of emerging market bonds, they might become a favorable investment. This sparsely populated medium-risk domain, which also includes assets with inadequately compensated risks, is exactly where the selected dynamic strategy fits in nicely. When measured in terms of annualized volatility, all categories of the dynamic portfolio strategy had realized volatility below those of high-risk bonds, while still providing higher returns than even most high-risk assets. When measuring risk in terms of maximum drawdown, dynamic strategy results are also less risky than high-risk bonds. While U.S. and emerging markets bonds did have maximum drawdowns between 25% and 30% for the chosen sample period, maximum drawdowns remain below 20% for all risk categories.

In addition to these final aggregated backtest outcomes, all measures also can be accessed as time series over time. Figure (4.15) shows performance time series and drawdowns over time for all categories. Again, low-risk categories are shown in dark blue, while high risk categories are shown in dark red. As can be seen in the left chart of the figure, performance series are almost perfectly color-sorted since mid of 2005, with higher risk categories associated to better performance. Only in the first quarter of the backtest this sorting is not fulfilled, which is mainly caused by the unfavorable starting date of the backtest period. Starting in June

2000 implies that the initial investment happens almost at the peak of the dot-com bubble, such that portfolios immediately have to struggle with harsh financial conditions during their first years. This holds all the more for portfolios of higher risk categories, as they generally have higher exposure to equity markets. Hence, low-risk categories end up outperforming higher risk categories in the first part of the backtest.

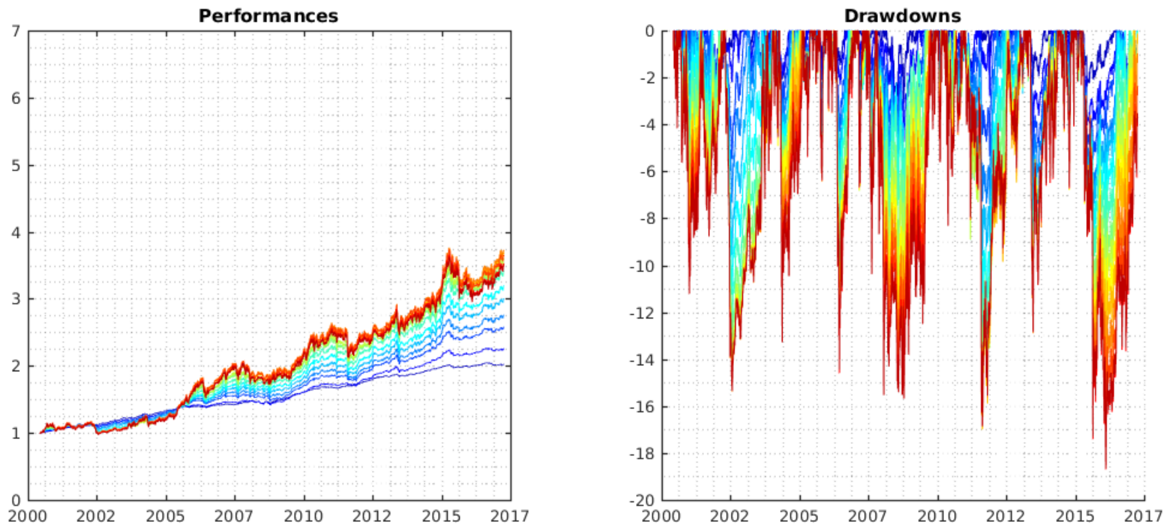

Figure 4.15: Strategy performances and drawdowns over time

Although there is a fairly stable upward trend for all categories, a couple of events did have lasting negative impact on performances. This best can be seen in the right part of Figure (4.15), which shows percentage drawdowns for individual risk categories over time. In particular, three major events did cause lasting drawdowns for portfolio values: the bursting of the dot-com bubble, the Global Financial Crisis and the European sovereign debt crisis. An alternative way of seeing this is by looking at yearly realized returns of risk categories, as shown in Figure (4.16). From sixteen full years that are part of the backtest and two partial years (2000 and 2017), only three years have negative returns for some strategies: 2002, 2008 and 2011. And even in these years, the worst drop in portfolio values is rather moderate with most adverse return equal to -7.2%. Compared to that, yearly returns of individual assets are orders of magnitudes larger, and extreme values amount up to more than -40% (see Figure (4.11)).

Another important aspect that can be seen from the chart is that yearly returns across different risk categories are consistent. In other words, neighboring risk categories tend to have quite similar returns, and yearly returns across the full range of categories are almost sorted in each year. Either high risk categories are performing best, and yearly returns increase steadily from low to high risk categories, or low-risk categories are best and realized returns decrease over categories. Only few years and categories do deviate from this pattern. This result is promising, as it indicates that returns are robust with regards to small perturbations. Slight changes in risk targets will also cause small deviations in outcomes only.

Figure 4.16: Yearly realized returns of individual risk categories

4.5.4 Asset contributions

For now, we only did look at aggregated results on a portfolio level, and not at individual components of the portfolio. Taking on a more disaggregated perspective, however, helps to get a deeper understanding about the portfolio decisions made by the algorithm: what assets did we end up investing into, and what was these individual assets' impact on performance and risk?

As first step into a more disaggregated perspective, the left part of Figure (4.17) shows average portfolio weights for each risk category. This is something that even might be required for regulatory purposes, as the ratio of bonds and equities is sometimes used to classify a portfolio with regards to risk. As can be seen, asset weights largely change over categories, with the fraction of high-risk assets increasing for higher risk categories. European covered bonds, for example, on average comprise more than 70% of the portfolio for risk category one, and its relative portfolio share decreases to roughly 5% only for risk category fifteen. In contrast, high-risk assets make up less than 10% on average in risk category one, but more than 60% in high-risk categories. But also within asset classes, fractions do differ between individual assets. For example, European and Japanese equities make up only small portfolio fractions for all risk categories, while other equity markets like Germany, Asia and emerging markets do get considerably more weight on average.

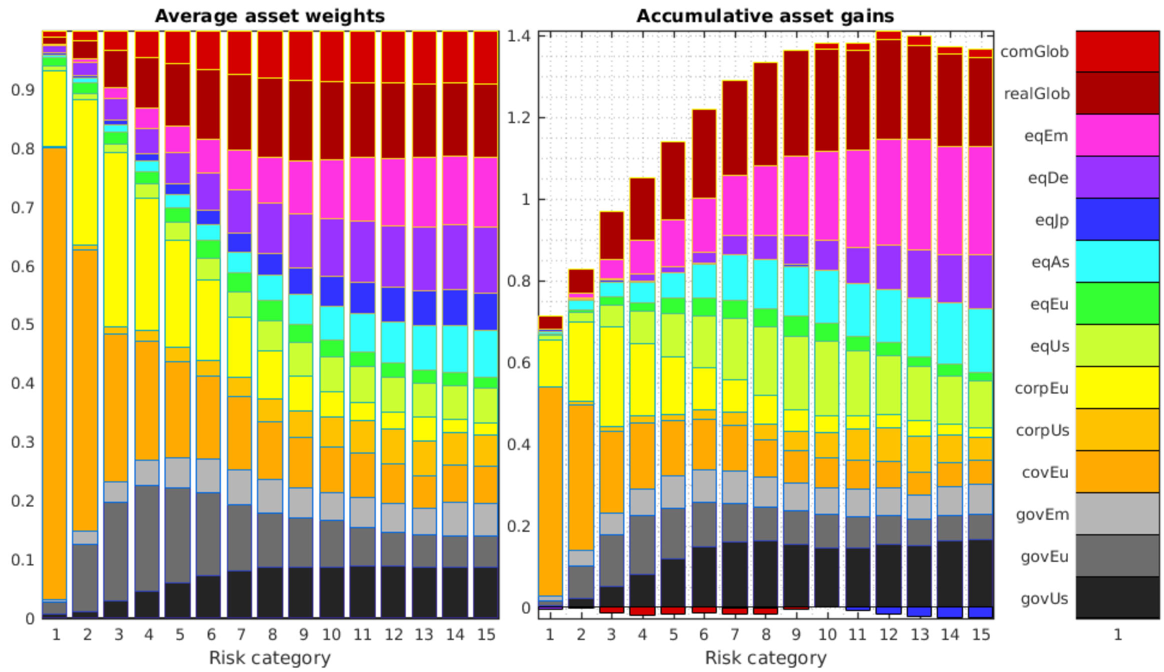

Figure 4.17: Average asset weights and asset performance contributions

Now that we know how much each asset has been chosen by the strategy on average, we want to evaluate the effects of each asset on the portfolio. Just because a certain asset has a large average portfolio weight does not mean that it also did contribute to portfolio performance proportionally much. Hence, we ideally want to evaluate how much each individual asset has contributed to the overall portfolio performance. It is generally not possible, however, to linearly split up a final portfolio performance into its components due to the way that compounding works.

Let's illustrate this with a simple example. Let's assume that a portfolio consists of asset A only in period t, and of asset B only in period $t+1$. For the two-period portfolio return we hence get

$$\begin{aligned} r_{t:t+1,P} &= (1 + r_{t,P})(1 + r_{t+1,P}) \\ &= 1 + r_{t,P} + r_{t+1,P} + r_{t,P}r_{t+1,P} \\ &= 1 + r_{t,A} + r_{t+1,B} + r_{t,A}r_{t+1,B} \end{aligned}$$

Due to the interaction term $r_{t,A}r_{t+1,B}$, it is generally not possible to uniquely allocate the effects of compounding to a single asset. An alternative way of allocating portfolio gains to individual assets would be to look at gains in absolute terms. Instead of decomposing a portfolio return, which is a relative quantity, we can identify the origin of each individual unit of money of gains that we made in absolute terms. In other words, when the portfolio value did increase by 100 Euro, we can associate each of the 100 Euro of gain in portfolio value to a unique asset. Let's extend the previous example by filling in some more numbers. Let the initial portfolio value be $P_t = 100$, and the individual asset returns be $r_{t,A} = 10\%$ and

$r_{t+1,B} = 20\%$ respectively. Then the portfolio value did increase from 100 to 110 in the first period, and from 110 to 132 in the second period. The overall gain of 32 Euro hence could be split up into 10 Euro price gain from asset A and 22 Euro price gain from asset B.

The problem with this way of allocating gains to individual assets is, however, that it weighs gains and losses more the higher the portfolio value is. This can easily be seen by a slight modification of the previous example. Let now individual asset returns both be equal, with $r_{t,A} = 10\%$ and $r_{t+1,B} = 10\%$. Then asset A generates gains of 10 Euro, while asset B generates 11 Euro. Hence, asset B generates more wealth in absolute terms even though both assets had the same return, only due to the fact that the portfolio value is larger in period $t+1$ than in period t. Individual asset returns hence get "portfolio value weighted" in a sense that returns will lead to higher gains and losses in absolute terms when the portfolio is large. This would not be a huge problem if portfolio weights were rather stable over time, because then all individual assets would be subject to low and high portfolio value periods in a similar way. However, in a setting with highly dynamic asset weights we want to avoid the distortion introduced by the "portfolio value weighted" approach.

Hence, we will choose a slightly different approach to measure performance contribution. The idea of the approach is quite simple: by artificially keeping portfolio values stable, we effectively eliminate "portfolio value weighting" as well as compounding effects. By fixing the portfolio value to a value of one each day, we can split up daily portfolio gains and losses to individual assets, and then aggregate over all days. The portfolio gain or loss in a certain period we will determine by

$$Gain = P_{t+1} - P_t = r_{t,P} P_t = \left(\sum_{i=1}^{d} w_{t,i} r_{t,i} \right) P_t$$

Setting the portfolio value equal to one, the formula can be simplified and reformulated

$$\begin{aligned} Gain = P_{t+1} - 1 &= \sum_{i=1}^{d} w_{t,i} r_{t,i} \\ &= \sum_{i=1}^{d} w_{t,i} r_i + w_{t,i} - w_{t,i} \\ &= \sum_{i=1}^{d} \left(w_{t,i}(1 + r_{t,i}) - w_{t,i} \right) \end{aligned}$$

The gain on a one Euro portfolio hence can be calculated as sum of the individual asset gains, which are computed as the difference between final value and initial value. If the portfolio value is equal to one, then the value of asset i in t is equal to its weight $w_{t,i}$. One period later, this absolute value has either increased or decreased depending on the asset return $r_{t,i}$, to a final value equal to $w_{t,i}(1 + r_{t,i})$. Aggregating over all individual periods hence gives a proxy for the wealth that a certain asset generates over time. Due to steady re-scaling of the portfolio value it does not matter anymore whether this wealth was generated at the beginning of the backtest period, where portfolio values are generally smaller, or at the end of the backtest, where portfolio values are already higher due to accumulated price gains. In turn, however, the interpretation of final asset contributions that represent accumulated gain by each individual asset is quite evolved. For example, if an asset did generate accumulated gains of 0.2 over the full backtest, then with a backtest period of roughly 17 years it did generate $0.2/17 \approx 0.011$ gains per year. Given that the average asset weight was 20%, this could be interpreted as a return of $0.011/0.2 \approx 0.059$. However, we are not really using asset contributions as a proxy for absolutely achieved asset returns, but we will only use them for relative comparison: which assets did contribute the most, and which assets did contribute the least.

Looking at Figure (4.17) again, the right part shows performance contributions of all individual assets for all categories of risk. As can be seen, overall cumulative asset gains increase with increasing category, only slightly dipping for highest risk categories. In addition to overall category results, the graphics now also sheds light on how overall gains can be split across the individual assets. If all assets had equal returns at all points in time, asset contributions would simply reflect average asset weights. However, asset returns obviously differ, such that contributions will generally deviate from average asset weights. For risk category one, almost all wealth is created by bonds, and only a tiny fraction by high-risk assets. This does not come as a surprise, as the portfolio also consists of bonds almost exclusively, as can be seen from average portfolio weights. For category fifteen, however, bonds only make up one third of average portfolio weights, but the wealth that is generated by these bonds is still roughly 70% of the wealth that is created by bonds in category one. The reason for this is that the composition of bonds changes over categories, with weights of U.S. and emerging market bonds increasing for higher categories. By gradually switching to higher-yield bonds, overall bond weights can be largely reduced without significantly reducing overall bond gains. Similarly, also some high-risk assets do have comparatively higher contributions, like for example U.S., Asian and emerging market equities. This is also something that we would expect for long backtest periods, as high-risk assets should be compensated with higher expected returns and hence they should contribute more strongly to increasing wealth. In this rather short backtest period of almost 17 years, this expected compensation for risk does not yet materialize for some of the high-risk assets, however. For example, Japanese equities have even slightly negative contributions to portfolios, such that they effectively diminish wealth. Similarly unsuccessful have been commodity investments that either did contribute only very little to wealth, or even did have small negative overall contributions for some categories. Small or even negative asset contributions indicate that investment decisions were not successful for these assets. If these negative asset developments had been predicted correctly, they would have been avoided, thereby reducing average asset weights for these assets. In other words, negative contributions should be a strong indicator for wrongly estimated asset moments. Following this conclusion, the exponentially weighted moment estimator that was used for the backtest seems to do a particularly bad job for equity Japan as well as commodities.

4.5.5 Weight filter impact

As laid out in Section (3.3), a substantial part of the dynamic trading strategies that were backtested is the usage of weight filters that help to reduce investment opportunities to those that are worthwhile even after subtraction of trading costs. These weight filters gradually modify portfolio target weights, and hence also influence portfolio properties. Hence, we will now analyze the impact of weight filters in more details, using the selected dynamic strategy with its persistence and cost-reduction filters as an example. Figure (4.18) shows the original single-period target weights that we would use in the absence of trading costs for three selected risk categories over time.

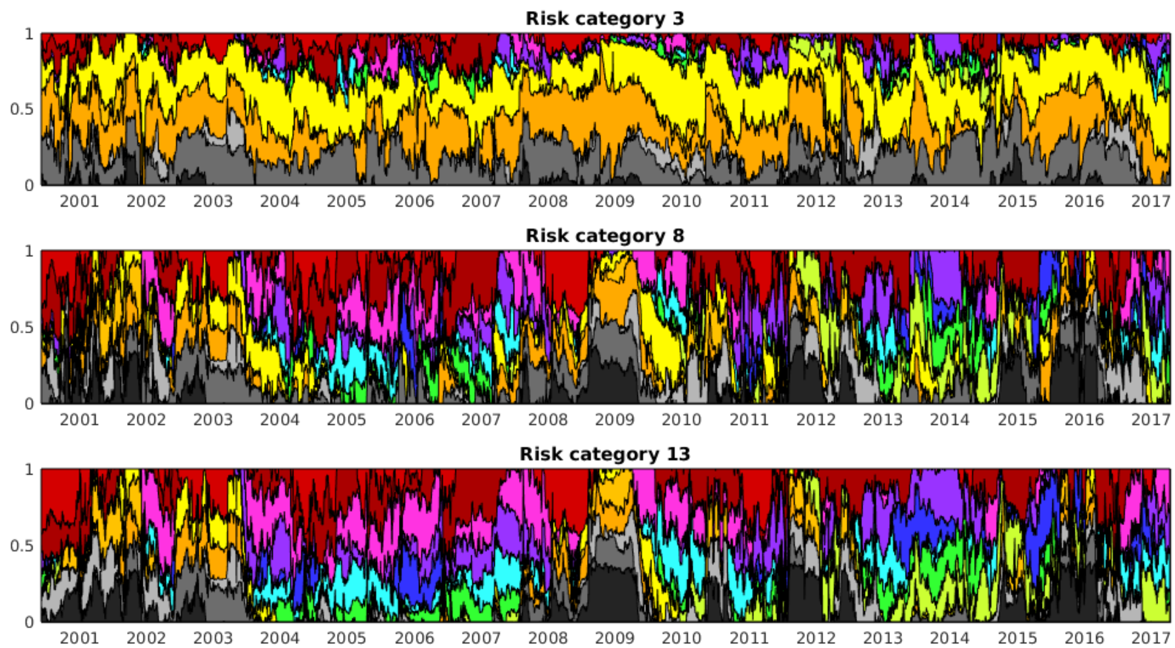

Figure 4.18: Single-period target weights over time

As can be seen, target weights for risk category three are more stable than for both other risk categories. This, however, is mainly caused by the fact that low-risk categories have less diversified portfolios, as they generally can invest in low-risk assets due to risk constraints only. Higher risk categories are subject to this effect to a lesser extent, such that they have larger scope when it comes to the selection of portfolio constituents. The reduced scope of risk category three also is reflected in the fact that portfolio weights are rather stable with regards to weights of different asset classes. Bonds account for at least 50% of the portfolio at each point in time, with aggregated weight fluctuating in the range of 60% and 80% most of the time. Only in periods of financial distress the weight on bonds increases even further, with weight above 90% during the Global Financial Crisis in 2009, and reaching even 100% at the end of 2001. On top of these broad shifts in asset weights, however, there still is a huge amount of high-frequency fluctuations that drive portfolio weights on a daily basis. As already mentioned, risk categories eight and thirteen have more room to manoeuvre, and hence dominating broad shifts in asset classes are harder to pin down. One major shift that can be determined is highly similar to risk category three, but conducted even more pronounced: bond weights increase substantially during financial distress. Both categories of higher risk also reach aggregated bond weights of 100% during 2001, and almost equally high levels during 2009 and occasionally in 2011 and 2016. In-between, however, portfolios are predominantly invested in high-risk assets, with almost 100% during 2004 and 2007. So while risk category three could be described as predominantly a bond portfolio that elevates bond weights even further during financial crisis, the other two categories can be described as predominantly high-risk portfolios that substantially re-balance into bond portfolios during volatile market periods. Again, broad re-balancing patterns are superimposed by high-frequency weight fluctuations that would cause tremendous trading costs if realized without further modifications.

Now that we know the original single-period target weights, let's examine how they are modified when weight filters are applied. The chain of filters that is used for the selected dynamic strategy is:

- persistence filter: $h^{\mathrm{MA}}(40)$
- cost-reduction filters: $h^{\mathrm{RTEmin}}(0.22)$ and $h^{\mathrm{RTEcap}}(0.18)$

Figure (4.19) shows target weights after the application of a moving average persistence filter. As can be seen, a large amount of high-frequency fluctuations in target weights is immediately filtered out. As a reminder, persistence weight filters should filter out any investment opportunities that are expected to hold only temporarily. In the simple implementation used here, the moving average of past weights will almost ignore target weight changes that did occur only very recently, and new patterns in target weights first need to persist for some time until they become effective.

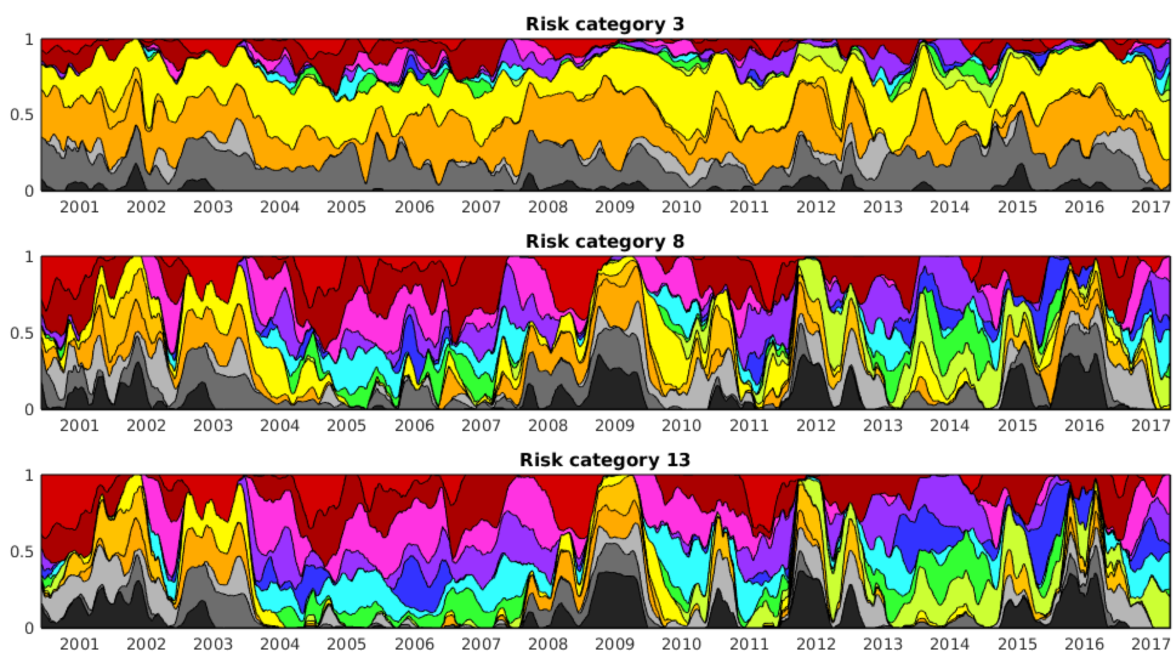

Figure 4.19: Target weights after persistence filter

While persistence filters should correct for endurance of investment opportunities, cost-reduction filters try to realize given targets as cost-efficiently as possible. In other words, they should solve the tradeoff between a preference of staying as close to portfolio targets as possible and the costs of trading that ideally should be minimized. For the selected strategy we did use two cost-reduction filters which both are based on relative tracking error as measure of portfolio similarity. Taken together, these cost-reduction filters basically determine a range of allowed relative tracking errors. If tracking errors pass the upper bound of allowed values, trading will be triggered, and the portfolio will be brought at least as close to the desired target as the value of the lower bound of relative tracking errors. As the portfolio adjustment is done with minimal costs, the portfolio will effectively hit the lower bound exactly, as any exceedance of the required portfolio similarity would incur additional trading costs. Final target weights after application of cost-reduction filters are shown in Figure (4.20). As can be seen, cost-reduction filters introduce some degree of discreteness into adjustment steps, as trading does not occur on a continuous basis anymore. If portfolios are close enough to targets, no trading will occur at all. Only if portfolios do deviate too much they will be re-balanced into the right direction again. Hence, portfolio deviation generally accumulates over time and then is resolved in certain irregular adjustment steps. This is what causes the step-wise patterns for target weights, especially for higher risk categories. Trading events thus are not equidistant anymore, but trading frequencies increase in times where faster reactions are required. For example, from 2005 to mid 2006 almost no trade events did occur for risk category thirteen, while multiple small adjustment steps are visible at the end of 2009, where bond exposure is reduced again.

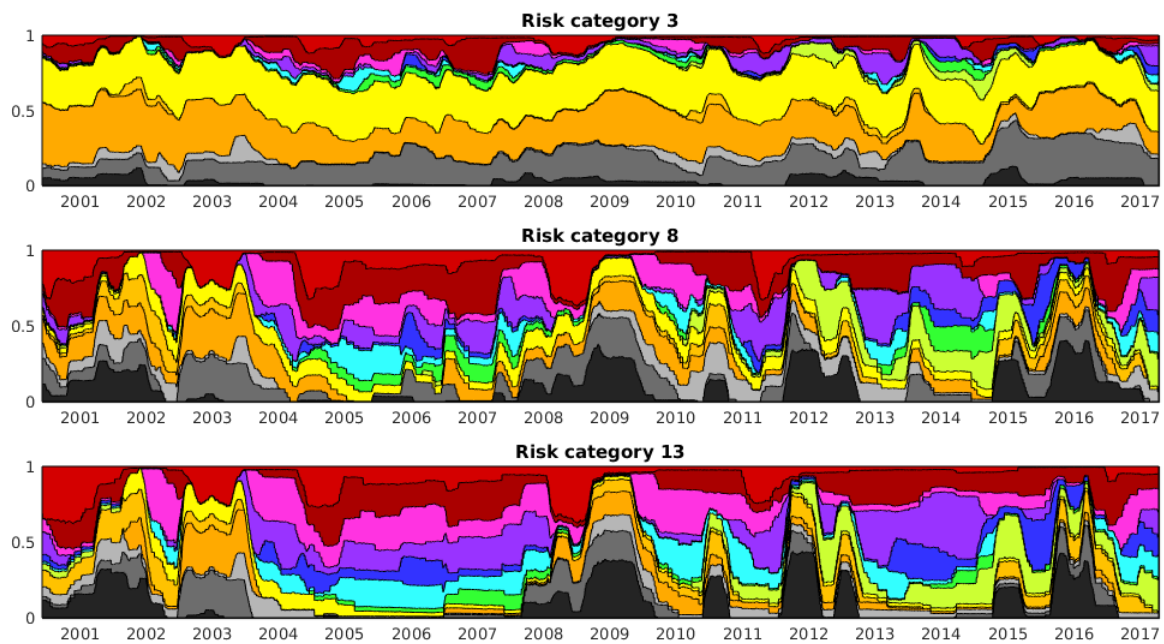

Figure 4.20: Final dynamic weights after cost reduction filter

As we have already pointed out before, differences in weights are a necessary requirement for diverging portfolio behavior, but is not guaranteed that portfolio properties really differ also. Two portfolios could be dramatically different in weights but still reflect the same risk factors and hence behave very similarly. Hence, although we have seen the effects of weight filters on portfolio weights already, we now want to examine the consequences of cost-reduction filters for portfolio characteristics. If cost-reduction filters work very well, then we should basically see no differences in portfolio characteristics. Figure (4.21) shows portfolio risk and return properties for all target weights before and after the application of cost-reduction filters. While differences in weights are directly observable, portfolio properties are never observable but have to be estimated. Hence, the analysis of cost-reduction filters will necessarily depend on the estimation method that is used for portfolio moments. Here we use two different estimators: first, we use conditional asset moments that were also used to determine single-period portfolio targets. And second, we use unconditional sample moments to get a more long-term perspective on how portfolios would differ. Conditional moments are shown in both graphics on the left, while unconditional moments are shown to the right. "Original" in titles refers to target portfolios before application of cost-reduction filters and associated graphics are shown in the upper row. Portfolio properties after cost-reduction filters are shown below. Portfolio moments are for daily discrete returns and not annualized, and shown for risk category twelve.

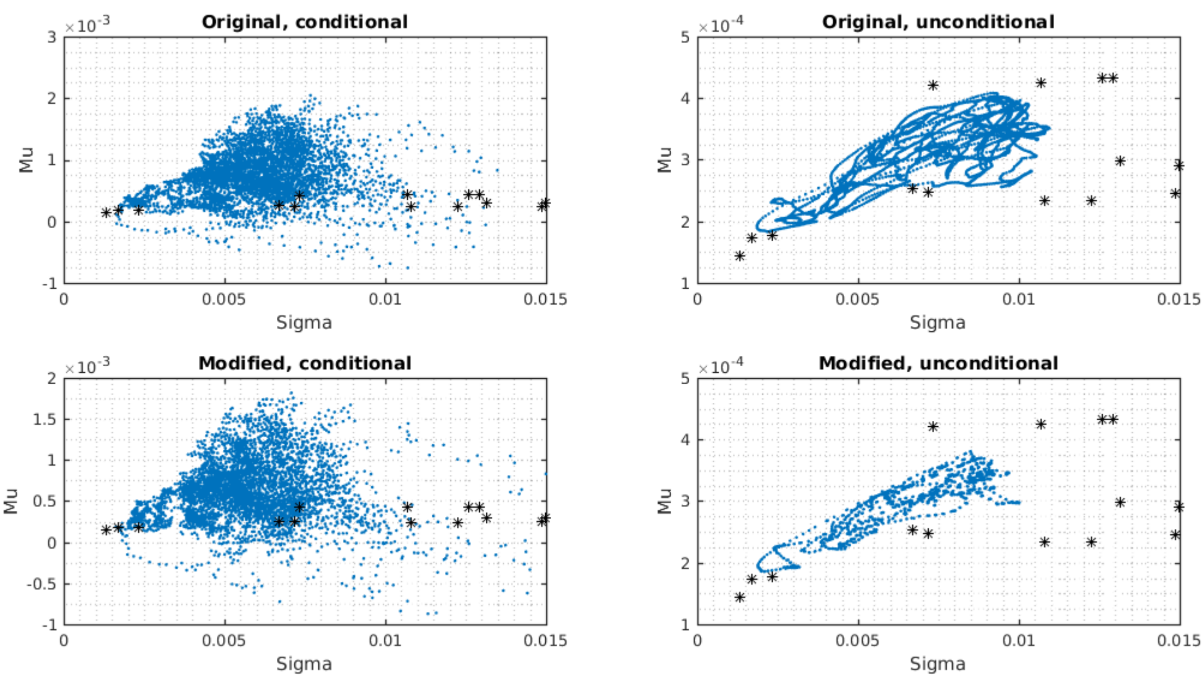

Figure 4.21: Changes of portfolio characteristics induced by cost-reduction filters

As can be seen in charts with conditional moments, portfolios do not represent original single-period target weights anymore, but are either target weights after application of persistence filters or after all weight filters respectively. Single-period target portfolio moments would generally match σ-targets in terms of portfolio volatility, but this is not the case anymore when portfolio weights already have been modified through weight filters and estimated portfolio moments then are distributed all over the place. A somewhat clearer pattern only can be seen when portfolio moments are computed from unconditional values, as it is the case for both plots on the right. Still, portfolio moments change quite heavily over time, which is also expected as the dynamic trading strategy can pick either equity or bond dominated portfolios depending on the financial environment. Comparing the right top chart with the chart below one can see how portfolio moments thin out, as portfolios are not allowed to continuously change anymore, but only change in discrete steps. On first sight, however, there is no sign that portfolio moments are too heavily affected by cost-reduction filters.

Let's now inspect the effects of cost-reduction filters more quantitatively with measures for turnover reduction and portfolio similarity. The following table contains values for all risk categories. Column two and three measure annualized turnover per category with (Final TO) and without cost-reduction filters (TO w/o CRF). By not fully realizing original target weights anymore, cost-reduction filters highly reduce annualized turnover values, and values in the range of 2.5 get diminished to turnover values close to one. Filters hence seem to do a good job when it comes to reduction of trading costs. Columns three and four measure the other important dimension of cost-reduction filters, which is the distance to the desired portfolio targets. While column three (Avg. RTE) measures the average relative tracking error to the target weights before application of cost-reduction filters, column four measures this distance in terms of turnover distance. By construction, average relative tracking errors are between 0.22 and 0.18, the ranges determined by both filters. Turnover distance measures how similar portfolios are in terms of weights, without taking into consideration whether weight differences even cause significant changes in portfolio properties. As can be seen, a rather stable deviation in terms of relative tracking errors leads to different levels of turnover distances. Portfolio weights have to be kept more closely to target weights for low-risk categories in order to achieve the same level of portfolio similarity than for high-risk categories. This finding, however, is strongly related to the fact that filters were built on relative tracking errors instead of (absolute) tracking errors, and relative tracking errors represent lower absolute tracking errors for low-risk categories.

Category	Final TO	TO w/o CRF	Avg. RTE	Avg. TO distance
1	0.17	0.91	0.20	0.15
2	0.38	1.71	0.19	0.19
3	0.53	1.82	0.20	0.18
4	0.72	2.25	0.20	0.22
5	0.87	2.56	0.20	0.25
6	1.01	2.76	0.20	0.26
7	1.08	2.85	0.19	0.26
8	1.11	2.84	0.19	0.25
9	1.09	2.80	0.19	0.25
10	1.05	2.73	0.19	0.26
11	1.01	2.65	0.19	0.26
12	0.99	2.59	0.19	0.26
13	0.99	2.54	0.19	0.26
14	1.00	2.51	0.19	0.26
15	1.00	2.50	0.19	0.26

Summing up, the chain of weight filters that was used as part of the selected dynamic risk management strategy exert an important influence on final portfolio decisions. While trading costs can be strongly diminished, portfolio characteristics themselves are changed only moderately. In this chapter we also analyzed the impact of all used cost-reduction filters in more detail. However, any analysis of individual parts of the full chain of filters only should be interpreted with care - any weight filter will have different impacts on weights depending on the actual combination that it is used in. For example, a cost-reduction filter could highly reduce trading costs when it is applied as first component in a series of filters. When used as last filter of some series, however, it might have almost no influence anymore, if weights are already hardly fluctuating anymore.

4.5.6 Comparison to static portfolio weights

In Figure (4.12) we can see that final annualized strategy returns and volatilities are close to what the efficient frontier would be for the full sample period. This is already a quite satisfying outcome, as we do almost reach the efficient frontier out-of-sample although we obviously did not use the true asset moments. However, we now want to further determine the origin of this favorable outcome, in order to distinguish between two possible different sources. First, the way we estimate moments and construct our portfolios in some way could just point us closely to the true optimal portfolio weights. In other words, the portfolio weights that we pick on average in each category are already close to the efficient frontier. And second, besides just picking adequate portfolio weights on average, the dynamic nature of the strategy itself adds some additional value to the outcome. In order to be able to differentiate between both components, we now will compare our dynamic strategies with fixed weights strategies that simply invest into the average realized weights of dynamic strategies. Hence, if a dynamic strategy would end up with 20% weight on U.S. equity on average, then we would compare it with a static portfolio with constant asset weight of 20% for U.S. equity. Figure (4.22) shows backtest results for these static strategies.

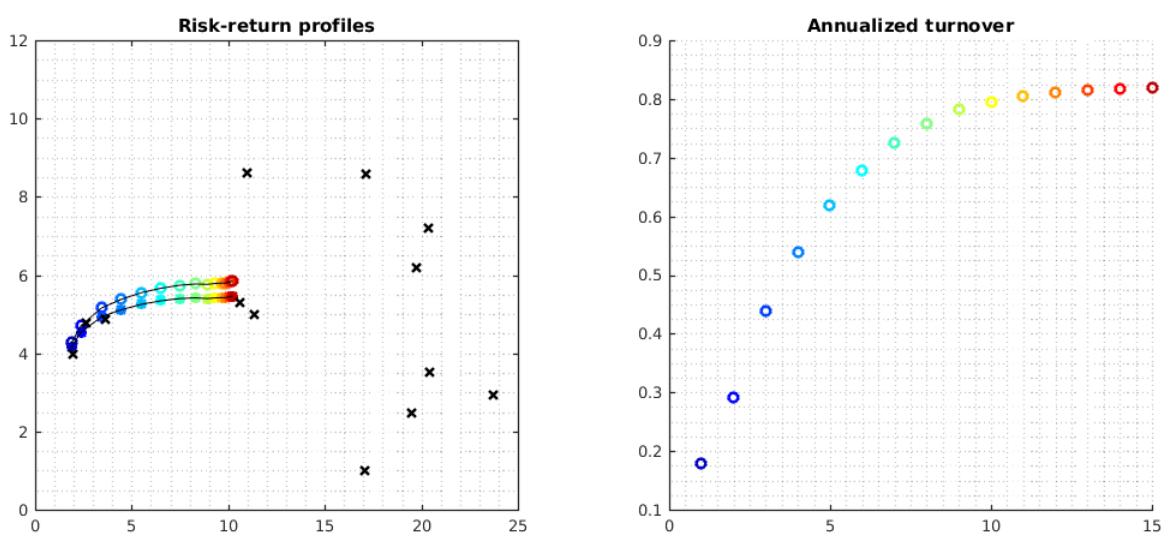

Figure 4.22: Annualized risk-return profiles and turnover of static portfolio strategies

As can be seen, static portfolio strategies do not have equally favorable outcomes anymore, but their risk-return profiles are considerably below the efficient frontier. Annualized turnover rates are still rather high for these static portfolio strategies, as we do offset weight changes due to market moves on a daily basis. This obviously could be made tremendously more cost-efficient with application of some cost-reduction filters, but we will anyways concentrate on gross returns only for this exercise. Comparing outcomes of dynamic and static strategies, dynamic strategies achieve significantly higher annualized gross returns than static strategies with similar levels of volatility. On top of that, however, dynamic strategies also are superior when considered through a different perspective on risk, namely maximum drawdown. This can be seen in Table (4.5), where gross annualized returns, volatilities and maximum drawdowns are shown for each category and both dynamic and static risk strategies. Maximum drawdowns of dynamic strategies are roughly half of those that are generated by static weight strategies. Hence, the dynamic risk management approach that is part of the selected dynamic trading strategy clearly pays off.

Category	Static weights			Dynamic weights		
	μ	σ	Max. DD	μ	σ	Max. DD
1	4.27	1.88	2.97	4.10	1.81	3.11
2	4.70	2.37	5.49	4.74	2.33	4.73
3	5.16	3.44	11.48	5.54	3.41	7.20
4	5.37	4.46	15.77	5.94	4.44	7.95
5	5.54	5.50	19.93	6.45	5.47	9.65
6	5.66	6.52	23.51	6.83	6.46	11.98
7	5.73	7.47	27.02	7.23	7.35	13.82
8	5.77	8.29	29.75	7.42	8.12	15.28
9	5.76	8.92	31.74	7.55	8.75	16.55
10	5.78	9.31	32.76	7.70	9.17	16.49
11	5.79	9.66	33.58	7.71	9.58	17.00
12	5.79	9.87	34.28	7.77	9.87	16.81
13	5.81	10.02	34.93	7.64	10.07	17.56
14	5.83	10.17	35.18	7.42	10.25	18.69
15	5.83	10.23	35.33	7.38	10.33	18.65

Table 4.5: Backtest key metrics for dynamic and static portfolio weight strategies

Besides the overall metrics shown in the table, performances and drawdowns of all categories of static portfolio weight strategies also can be seen over time in Figure (6.5) in the Appendix. To get a more detailed insight into why dynamic strategies mitigate drawdowns, we will now analyze drawdowns of static and dynamic strategies of risk category eight. Thereby the upper plot of Figure (4.23) shows drawdowns for both strategies, with drawdowns of the dynamic strategy shown in blue. The second chart shows corresponding portfolio weights over time for the dynamic case. For the static case a similar chart would just show the same weights at each point in time. Chart three and four show daily percentage returns for dynamic and static strategies respectively.

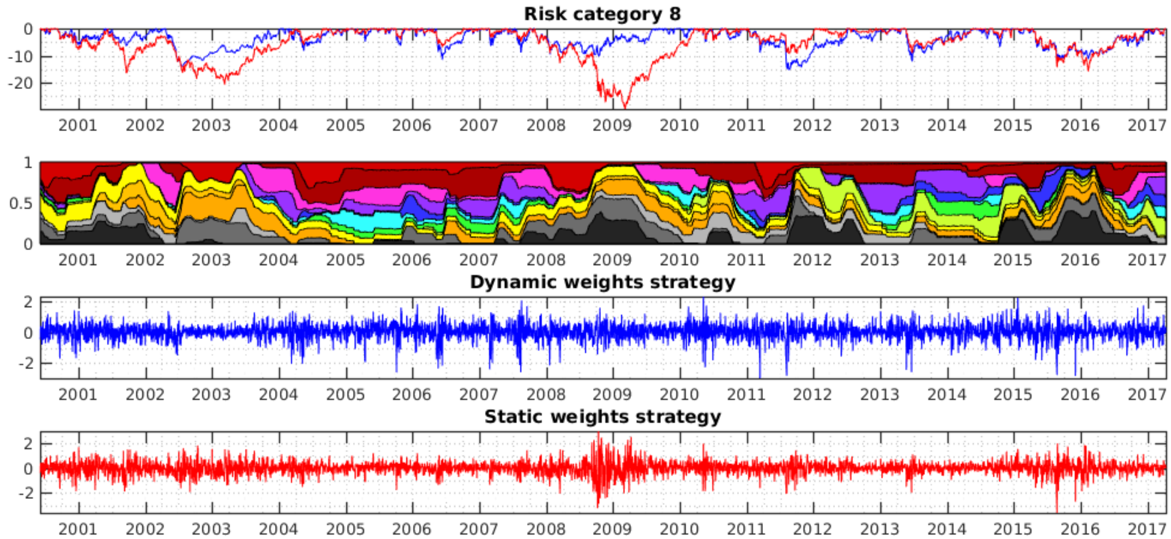

Figure 4.23: Drawdowns, daily strategy returns and weights over time

As can be seen, drawdowns of both strategies generally tend to coincide over time, with deviations only when drawdowns of static portfolio weights heavily increase. Without the possibility for adjustments, portfolio drawdowns of static weight strategies will reflect asset drawdowns unrelievedly. In contrast to that, dynamic strategies are able to adapt market exposure, and hence potentially preserve portfolios from more severe losses. This is exactly what can be seen in backtests, where portfolio weights are largely adapted in periods of major market downturns. The strategy reduces exposure to high-risk assets by increasingly shifting portfolio weights into bond markets. This portfolio re-structuring effectively avoids further losses, and drawdowns of dynamic weights suddenly start to deviate from static ones. Actively shifting portfolios into low-risk assets hence proves to be an effective risk management approach. Similar effects also can be seen for all other risk categories, and results are shown in Figure (6.6) and Figure (6.7) in the Appendix for categories three and thirteen respectively.

Daily portfolio returns for the static weight case are just as one would expect. With static weights, portfolio returns simply inherit volatility clusters from asset returns. Given that asset returns are stationary processes with time-varying moments, the same also holds for discrete portfolio returns under quite general conditions as they are just a linear combination of asset returns. In contrast, however, dynamic changes of portfolio weights will complicate things, as individual daily returns suddenly arise from substantially different underlying portfolios. In some periods, dynamic weights are largely dominated by bond investments, such that returns are basically driven by fixed-income risk factors. In particular, bond returns are much less prone to heteroskedasticity than equity returns. If bond returns and asset returns were following different GARCH processes, then dynamic strategy returns would roughly be a mixture between a bond GARCH process and an equity GARCH process, and hence most likely not be a GARCH process themselves. We do not intend to give any formal proof for this argument, but want to leave it at building some intuition for the intricacies that could generally arise for dynamic strategy returns. In the worst case, these intricacies even might shatter stationarity of dynamic portfolio returns.

For now we will not further question stationarity, however, and instead have a straightforward look on unconditional distributions of realized portfolio returns. Figure (4.24) shows quantiles and extreme values for both dynamic (left) and static (right) strategies. On the x-axis the figure shows all fifteen risk categories, while absolute quantiles of logarithmic percentage returns are shown in y-dimension. The black line close to zero in both charts represents the median of realized logarithmic portfolio returns. It is slightly above zero for all risk categories, and hence coincides with the finding that all backtested strategies end up with positive overall returns for the tested period. Red lines represent absolute values of lower return quantiles, and blue lines represent values for upper quantiles. The lowest blue and red lines mark 25% and 75% quantiles respectively. Since red lines are slightly below blue lines, daily returns are not symmetric around zero for these quantiles and positive return quantiles are slightly larger than negative return quantiles. This is not surprising, as already return medians are above zero. Moving further into the tails of return distributions, this effect gradually reverses. While 5% and 95% return quantiles are already fairly symmetric around zero, 1% quantiles are further away from zero than 99% quantiles. This difference between positive returns and negative returns even increases for extreme values, where minimum realized returns have absolute values considerably larger than realized return maxima.

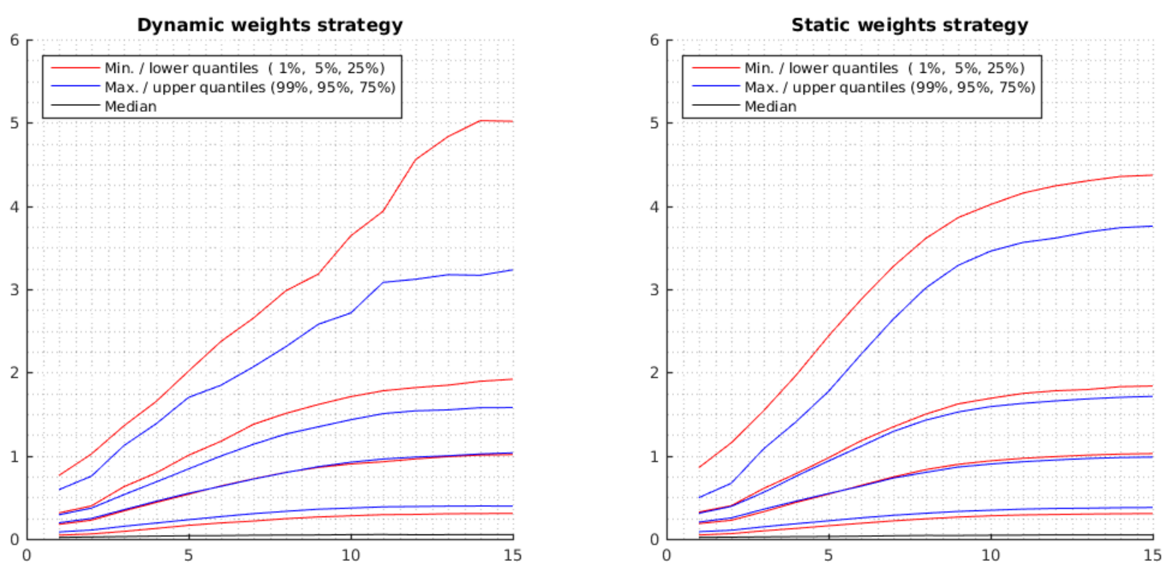

Figure 4.24: Daily return quantiles for static and dynamic portfolio weight strategies

Comparing left and right charts, this skewness of logarithmic return distributions is clearly more pronounced for realized returns of dynamic strategies than for static weight strategies, with more extreme negative returns for dynamic strategies. Interestingly, this pattern has to be related to the time horizon that is considered, as the significantly lower drawdowns of dynamic strategies indicate that dynamic strategies have less downside risk than static strategies for longer time horizons. By actively adapting portfolio weights, dynamic strategies temporarily increase high-risk asset exposure, which in turn increases downside risk on a daily basis. Let's assume for example that dynamic weights did select an equity only portfolio temporarily. Then such a portfolio will be prone more to sudden market downturns than the corresponding average weights portfolio that also is invested into bonds to some degree. After taking the first hits, however, dynamic strategies will re-balance the portfolio into low-risk assets, thereby ultimately reducing downside risks measured on longer time scales.

A similar conclusion we also can derive from estimated realized strategy volatilities shown in Figure (4.25). Both charts show estimated conditional volatilities for logarithmic percentage returns. For each category we thereby estimate a single GARCH process and use it to retrieve realized volatilities. Comparison of dynamic weight volatilities in the upper chart and static weight volatilities in the lower chart show that dynamic weights lead to more medium spikes, while maximum volatility levels are larger for static weights. This again indicates that short-term dynamic portfolios can be more exposed to market risks when portfolios put excess weight on risky assets, and hence produce more spikes in the first place. But due to the ability to adapt, largest volatility spikes can be avoided, as they usually do not build up instantaneously but over medium time horizons.

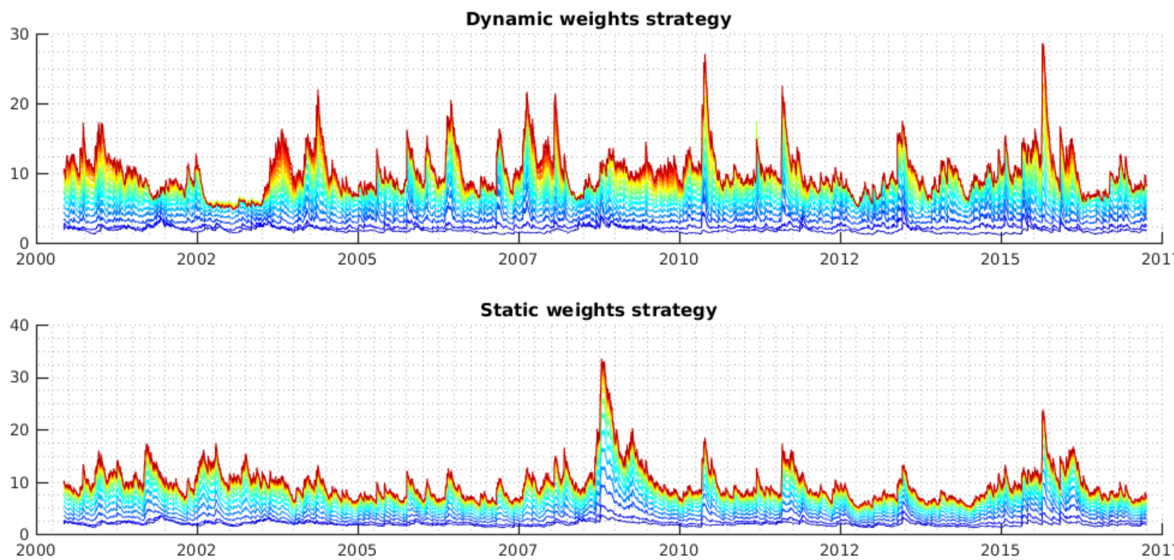

Figure 4.25: Estimated conditional volatilities over time

To provide a more complete picture, we want to again point out a problem that we already mentioned before. Dynamic strategies potentially could represent substantially different portfolios at different points in time. In some periods they are largely dominated by fixed-income risks, while they are dominated by equity risks in other periods. This way, final strategy return series could become non-stationary, but at the very least they generally should be arising from different processes with different patterns of heteroskedasticity. Hence, estimating a single GARCH process for the full return series should be a simplification, and a more complex volatility estimation should take this into account. One way of doing this is by retrieving individual daily volatility estimates from different data generating processes. For example, for each day one could hold portfolio weights fix and use them to generate a synthetic series of historic portfolio returns. This will generally differ from true realized historic portfolio returns, as realized returns arise from different portfolios. Once these synthetic returns are determined, they could be used to estimate a GARCH model for the specific day and portfolio, which in turn could be used to retrieve concurrent volatility. This is something that we also tested, and estimated volatility paths with daily re-estimation of GARCH processes did not produce any significant visual differences, except for a short period starting at the end of 2002. As a response to market developments portfolios were heavily invested into bonds during that period, which did lead to numeric issues during the estimation of GARCH parameters. In other words, when portfolios are low-risk fixed-income portfolios in local currency, they simply might not have volatility clusters, and hence estimation of GARCH parameters could become a problem. There do exist simple workarounds for this problem, like for example using a fallback that treats synthetic returns as *iid* and estimates volatility as sample standard deviation. Results do not really differ visually, so we will not further complicate things at this step and leave it at the simple way of estimating volatilities.

4.5.7 Risk contributions

In Section (4.5.4) we have already analyzed how individual assets contribute to the overall performance of the portfolio. However, returns should never be considered without also taking risk into perspective. This also holds for asset contributions. An asset that generates moderate portfolio gains, but tremendously increases portfolio risk might not be a worthwhile investment in the end. Hence, we now want to split up overall portfolio risks onto individual assets. Thereby we will follow the definition of marginal risk contributions given in (T. Roncalli 2013). First step of the approach is to derive the sensitivity of portfolio volatility with regards to individual asset weights. This corresponds to the derivative of portfolio volatility with regards to asset weights, and it is given by

$$\begin{aligned}
\frac{\partial \sigma(w)}{\partial w} &= \frac{\partial \sqrt{w' \Sigma_R w}}{\partial w} \\
&= \frac{1}{2}\left(\sqrt{w' \Sigma_R w}\right)^{-1} (2 \Sigma_R w) \\
&= \frac{\Sigma_R w}{\sqrt{w' \Sigma_R w}}
\end{aligned}$$

Using this formula for marginal risks, marginal risk contributions are defined as

$$\mathcal{RC}_i = w_i \frac{(\Sigma_R w)_i}{\sqrt{w' \Sigma_R w}}$$

Defining marginal risk contributions that way has the advantageous feature that risk contributions fulfill the allocation property: the sum of individual risk contributions adds up to the overall portfolio volatility. This can be seen from the following equations:

$$
\begin{aligned}
\sum_{i=1}^{d} \mathcal{RC}_i &= \sum_{i=1}^{d} w_i \frac{(\Sigma_R w)_i}{\sqrt{w'\Sigma_R w}} \\
&= w' \frac{\Sigma_R w}{\sqrt{w'\Sigma_R w}} \\
&= \sqrt{w'\Sigma_R w} \\
&= \sigma_P
\end{aligned}
$$

As individual risk contributions depend on covariance matrix Σ_R, and covariance matrices are modeled as varying over time, we also get time-varying risk contributions. Using estimates of conditional covariance matrices we can split up the conditional portfolio volatilities shown in Figure (4.25) into asset risk contributions for each historic date. The results are visualized in Figure (4.26) for backtest results of risk category eight and discrete percentage returns. Due to diversification effects, risk contributions of individual assets can even become negative, and this is actually the case for some assets at specific points in time. Negative risk contributions are stacked below zero, while positive risk contributions are stacked above.

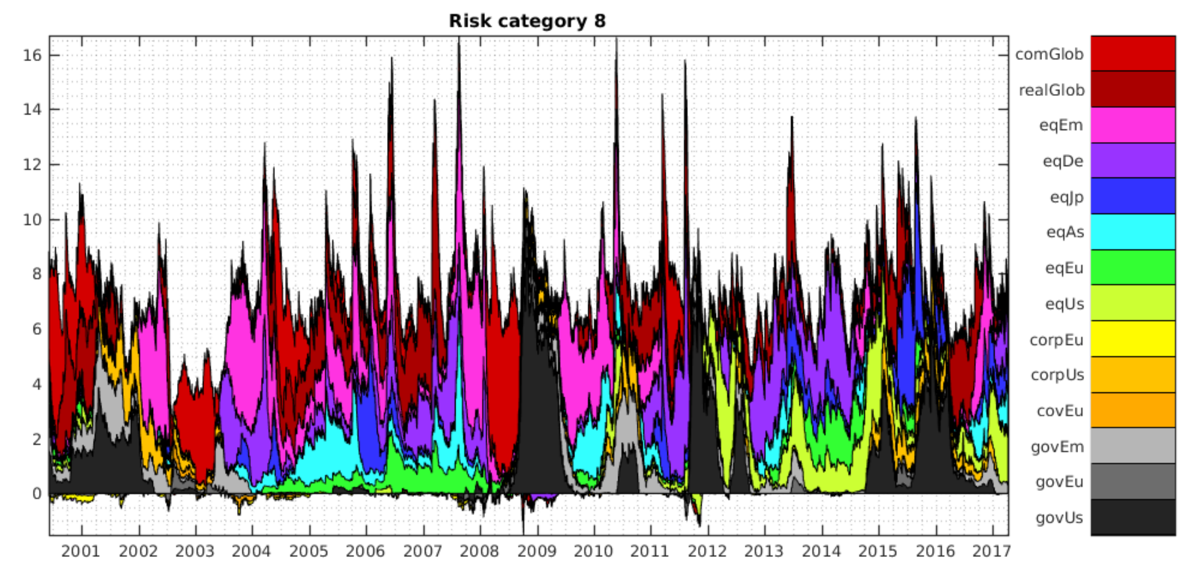

Figure 4.26: Volatility contributions over time

Although aggregated levels of volatility now are estimated based on covariances, and with regards to discrete instead of logarithmic returns, the general pattern of portfolio volatility over time strongly resembles the values estimated by GARCH in Figure (4.25). For example, the heavy reallocation into bonds starting mid 2002 again can be clearly seen as a dip in overall volatility. Looking at individual asset risk contributions, however, we now can identify assets with comparatively low or high risk-adjusted contributions. From the figure we can see that European bonds almost do not contribute to portfolio risk at all. Even at beginning of 2003, when European bonds did make up roughly 70% of portfolio weights, they contribute less than 0.5% to annual volatility. In contrast, global commodities, which made up the remaining

30% of portfolio weights, are contributing another roughly 3.5% to annual volatility. When considered over the full period, commodities do contribute a substantial share to portfolio risk at multiple points in time, although these investments in fact slightly decrease wealth over the full period as we have seen from asset performance contributions.

Figure (4.27) just shows a slightly different perspective on asset risk contributions. Individual risk contributions thereby have been normalized such that they add up to one, which makes them more easily comparable to asset weights. Again, we can clearly see that asset weights only are a proxy to how much risk certain assets contribute. Especially European bonds do almost never add to portfolio volatility substantially, even though they make up approximately 30% of portfolio weights on average. In contrast, U.S. government bonds clearly contribute to portfolio risk in some periods. The reason for this is that U.S. government bonds, as opposed to European bonds, come with exchange rate risk on top of the interest rate risk that is common to all bonds.

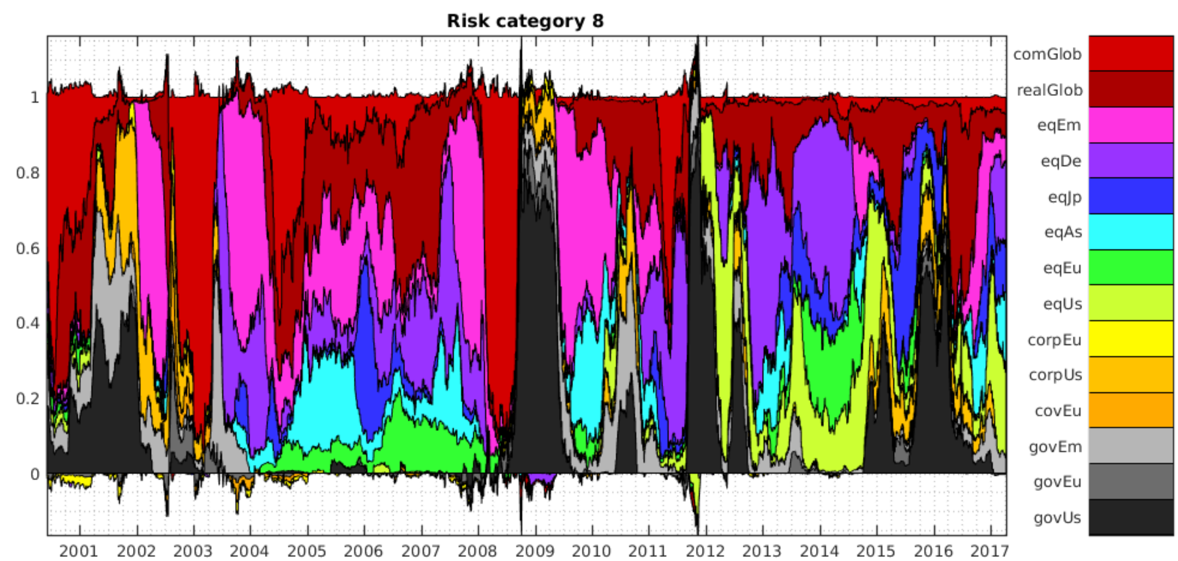

Figure 4.27: Relative volatility contributions over time

4.5.8 Client dispersion

Traditionally, in asset management investment strategies are made available to multiple clients through fund solutions. All assets under management are jointly invested into a single portfolio, and clients participate at portfolio gains and losses according to their respective share of wealth. This way, however, there is no way to advantageously take into account any particular client circumstances, and all clients are treated equally. Clients, however, generally differ with regards to three important characteristics: cash-flows, inception dates and tax conditions. Only when clients are treated individually, inception date and concrete tax conditions could be used to optimize after tax returns. For example, long-term clients that already did build up taxable gains on individual positions could realize taxes in order to fully exploit open tax allowances. This is not possible in a setting of global fund solutions.

Furthermore, individual treatment of clients also can be a matter of fairness. Any new client should initially get the currently best portfolio weights, and not become part of a global portfolio with potentially suboptimal portfolio weights and actual re-balancing requirements due to market moves. Similarly, if a client makes regular deposits they also should be put to use to actively re-balance his individual portfolio towards a preferred portfolio target, and not disappear as negligible contribution to a single joint global portfolio of all clients. For all these reasons, all proposed dynamic trading strategies ideally should be implemented in a customized setting without global fund solution, but with individual portfolio per client.

In such a customized portfolio management environment, individual portfolios will naturally deviate at each point in time, even for clients within the exact same risk category. Depending on inception date, portfolio size, tax situation and cash-flows clients will have slightly different portfolio weights, and differences in weights are a potential source of differences in performances. In other words, even when two clients are part of the same risk category, they could achieve different returns during an equal holding period. This, in turn, might be an undesired consequence of individualized portfolio management approaches, as it might require additional explanations to clients when they compare portfolio performances. For example, tax optimization should lead to optimized after tax performances for some clients, but it ideally should not cause gross returns to differ widely. Hence, we now will examine all risk categories with regards to whether they cause dispersion of client performances.

In our simplified setting of fractional weights without client cash-flows and taxes the only source of client dispersion is different starting dates. Depending on the actual starting date, subsequent rebalancing events will generally be at different points in time, as trading events are neither set to occur on specific dates upfront nor arranged in equidistant time intervals. Hence, portfolio weights will generally differ for clients of different portfolio inception dates. In order to quantify the magnitude of this effect, we analyze backtests for 35 clients per risk category. Starting July 1st, 2000, a client will be added to each risk category every six months. Figure (4.28) shows final portfolio weights for each client of selected risk categories, depending on the date of his initial portfolio investment shown on the y-axis.

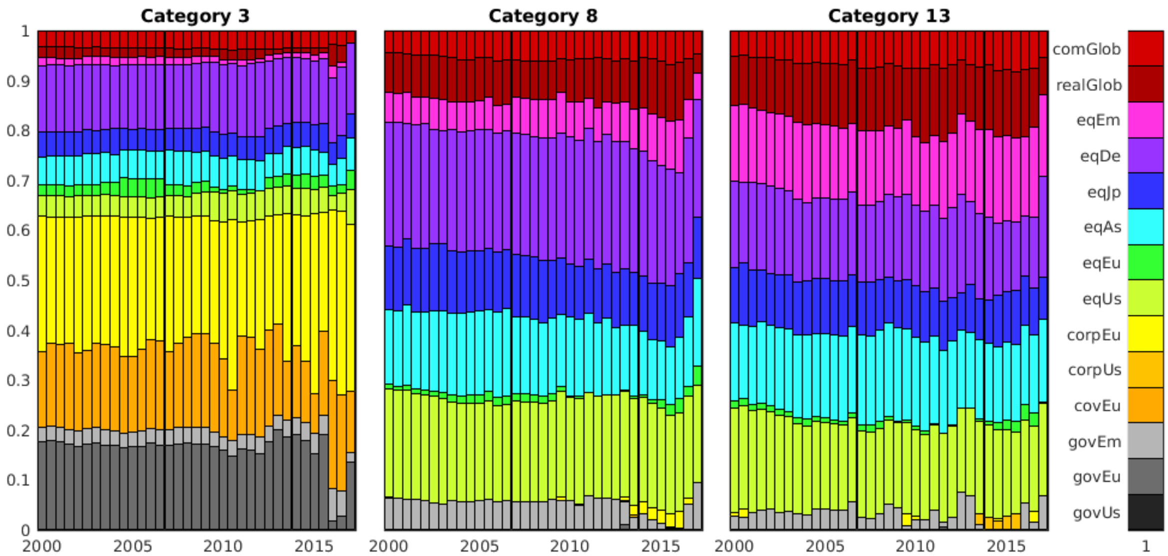

Figure 4.28: Final portfolio weights with regards to portfolio inception dates

As can be seen from the figure, final portfolio weights by and large are similar for all clients within a given risk category. Nevertheless, one can still identify some clients with slightly different portfolio weights at the end of the backtest period. For example, for risk category three there are two outliers among the last three clients who have substantially less weight for European bonds. Similarly, for categories eight and thirteen there exist small ranges of subsequent portfolio start dates from 2013 to 2015 where clients do end up with slightly different portfolio weights at the end of the backtest period. A slightly different perspective on the differences that result for final weights are also shown in Figure (6.8) in the Appendix. It shows Box plots for individual asset weights for all selected risk categories such that differences of individual asset positions can be seen more clearly.

Differences in weights arise although all clients actually have the same target portfolio weights at each individual day, only because these target weights are not fully realized due to different cost-reduction filters. It might be the case that these differences do only exist for a couple of days, and outlier clients are just about to be re-balanced the next days. However, we do not get an estimate for the duration of these differences from this chart.

Another way to examine portfolio dispersion is by using aggregated distance measures on portfolio weights. Figure (4.29) shows for each pair of clients with different starting dates the aggregated distance between final portfolio weights. Thereby aggregated distance is simply measured as sum of absolute differences of asset weights. Dark blue squares represent tiny portfolio distances, yellow squares medium distances and dark red squares large differences in portfolio weights.

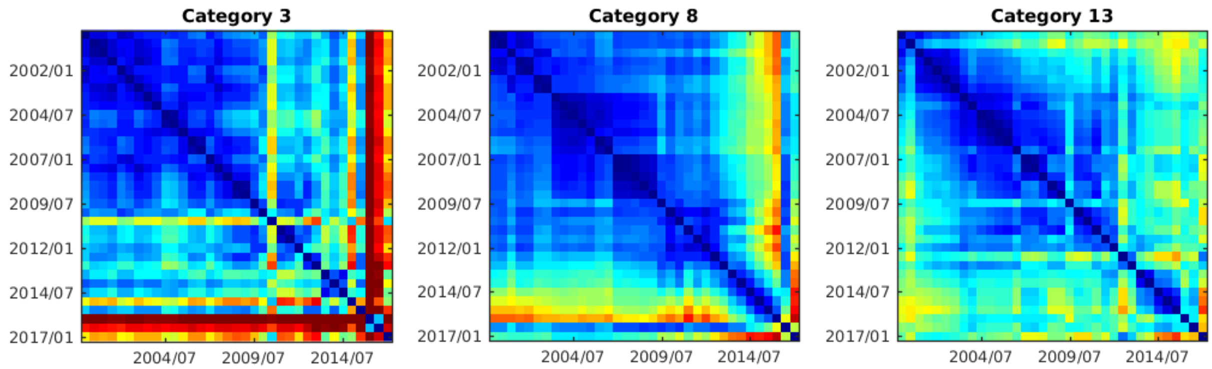

Figure 4.29: Final portfolio weight differences for clients with different starting dates

Similar to the previous results, this chart again shows that certain client portfolios will deviate from an average representative final portfolio allocation. These outliers have large distances to the majority of other clients, with usually small differences to a few neighboring clients. For example, the third and second last clients of risk category three are almost similar, but are largely different when compared with all other clients. Additionally, we can see that portfolio similarity might be something that strengthens with increasing passage of time. For all three categories shown large blueish blocks emerge in upper left direction. Hence, portfolio weight dispersion might be largely induced by the initial portfolio allocation, and the effect generally washes out over time.

Ultimately, however, as we have already argued at several places before, it is not weight differences that matter, but differences in portfolio behavior. Hence, we will now check whether portfolio performances differ over time. Whenever multiple portfolios within the same risk category share an overlapping holding period, they ideally should achieve equal performances for this overlapping time period. In other words, if a client starts in 2005, and another one starts in 2010, and both are observed until the end of the backtest period, then both ideally should have a similar performance from 2010 until the end. In this setting, where clients do neither terminate their account nor leave the risk category prematurely due to any other reason, they can only overlap in their final part of the holding period. Hence, performance dispersion best can be detected when all performance series are pinned to an equal value on the right end. Performance then does not aggregate daily returns chronologically, but backwards over time. We will denote this from now on as backward-looking performance. Figure (4.30) shows backward-looking performances for all backtested clients for risk categories three, eight and thirteen.

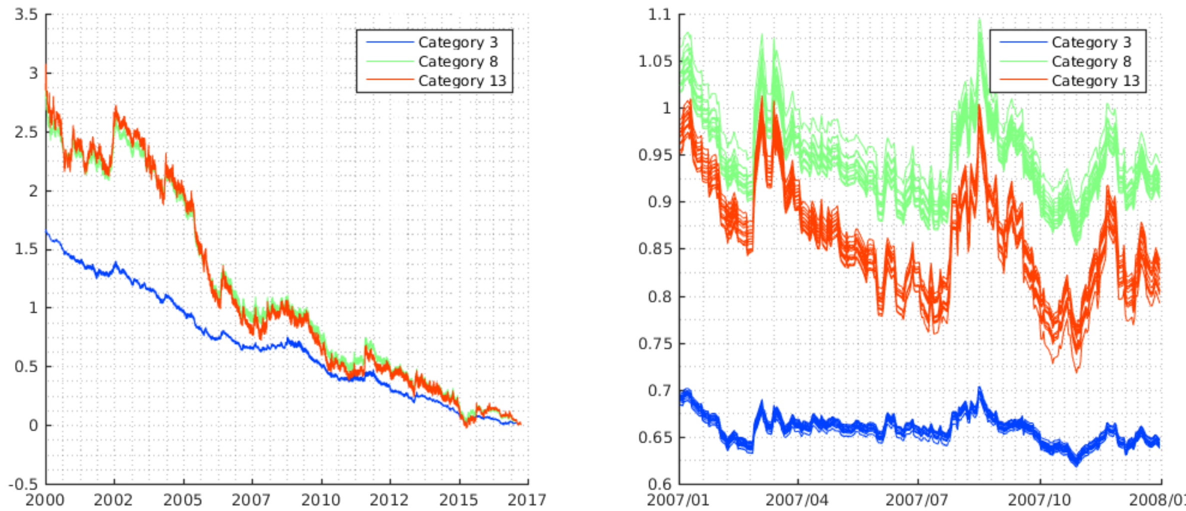

Figure 4.30: Backward-looking portfolio performances for different clients

As can be seen in the left chart, without magnification there basically appear to be three lines only, one per risk category. In reality, however, each of these three colored lines consists of up to 35 lines, one for each backtested client. Clients that only start belated will be represented by a line that extends over a part of the full period only. A client that starts beginning of 2005 hence will be reflected as a line from 2005 until the end, with end point fixed at value zero. Backward-looking performances then can be interpreted as follows: if the backward-looking performance has a value of 2 in 2005, then the performance of the strategy from 2005 until the end is 200%. Moving instead further along time to year 2010, the backward-looking performance will only represent the period from 2010 until the end, and hence it represents a shorter period in time. Backward-looking performances hence are generally decreasing with time, at least when daily returns are positive on average.

Using backward-looking performances, client dispersion can be read-off quite conveniently. At any point in time, the difference between highest and lowest backward-looking performance directly represents performance dispersion. For this, however, we need to zoom in on certain sub-periods, because otherwise client dispersion is already too small to be visible. This is what is shown in the right chart, where the considered time period is restricted to year 2007. For example, backward-looking performances of risk category thirteen at begin of July 2007 are between 0.82 and 0.86. In other words, all clients that already existed at beginning of 2007 had a performance between 82% and 86% until the backtest ends in 2017. Hence, performance dispersion is 4% over a time period of almost ten years. Similar results we also get for risk category eight, while performance dispersion for risk category three is even smaller. With such small deviations in performances we can conclude that client dispersion is not a problem for the selected dynamic allocation strategies. Portfolio performances almost do not differ at all points in time, such that portfolio behavior generally seems to be consistent across all clients even though portfolio weights themselves differ. A possible explanation for this pleasant result could be the fact that cost-reduction filters directly control tracking errors, and hence take a more sophisticated view on portfolio deviations than simply measuring deviation of portfolio weights.

Last but not least, backtests for a full cohort of clients with different starting dates allow another robustness check for dynamic trading strategies. Ideally higher risks should also be rewarded with higher returns, such that higher risk categories should achieve better annualized returns than low-risk categories for sufficiently long investment periods. This is examined by Figure (4.31) which shows final annualized backtest returns for all risk categories and clients with different starting dates. Dark red colors again represent high risk categories, while low-risk categories are marked in blue.

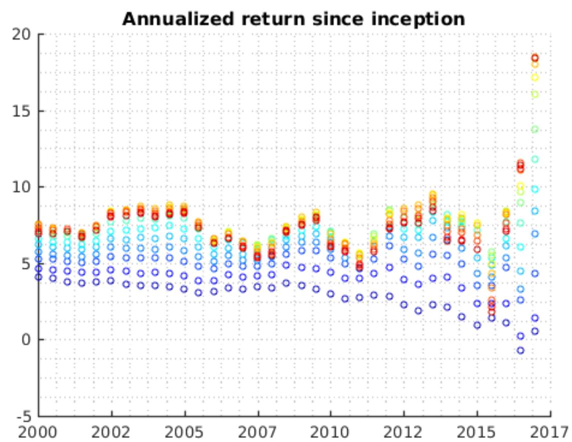

Figure 4.31: Annualized realized portfolio returns for different starting dates

As can be seen, for all clients starting as early as up to 2005 annualized returns show a neatly aligned color gradient. High risk categories show up on top, meaning that higher risks actually are rewarded. Looking at clients starting later in time, from 2006 to 2013, the highest risk categories start to slightly lack behind, and categories around number twelve achieve highest annualized returns. Except from this permutation of top risk categories the correct ranking of strategies still holds. Only for inception dates between 2013 to 2015 the ranking of strategies becomes really flaky. This, however, is not unexpected, as strategy volatilities are much too high to be representative for such short periods.

Overall, the results of Figure (4.31) are quite promising. Neighboring clients of equal risk categories are generally close in terms of annualized returns, which means that strategies are reasonably stable with regards to changing inception dates. Especially when investment periods are sufficiently long, all categories seem to stabilize around reasonably well performances.

Chapter 5

Conclusion

Although mean-variance portfolio optimization still can be seen as the gold standard of quantitative portfolio management, it usually is perceived to be flawed in practical applications. Its main deficiency is that it builds on known asset return moments, although inputs have to be estimated in real life. Straightforward application hence tends to maximize errors: assets that appear to be better due to estimation errors will be overweight in portfolios, which leads to suboptimal portfolio choices. In other words, Markowitz portfolio selection is too powerful for the quality of its inputs (Scherer 2002).

Generally, the quality of any quantitative asset allocation strategy basically depends on two components: first, inputs should be modeled as precisely as possible. Obviously, if one could perfectly predict the market, there is no need for excessively complicated strategies to exploit this capability. Simply investing in the asset with highest return would do the job. And if only asset return moments are fully known, Markowitz portfolio optimization can be used to make optimal portfolio choices, at least for mean-variance investors. Whenever market predictions are highly imprecise, however, an additional component becomes important: being able to also recognize the level of estimation errors around optimization inputs. Hence, while increased precision of inputs obviously leads to better portfolio choices, increasing the resilience to faulty inputs also does.

Therefore, we propose a diversification-aware approach that allows to conduct mean-variance portfolio selection under constraints on diversification. This way one can reduce the sensitivity to inputs as compared to traditional mean-variance portfolio selection. The minimum required level of diversification serves as single parameter to control sensitivity to inputs. Setting the required diversification level low enough, the solution will coincide with the efficient frontier portfolios of traditional Markowitz optimization. On the other end, for the limit case of maximally high required diversification, portfolios will coincide with the equal weights portfolio, making estimated asset return moments meaningless. The required diversification level hence should reflect how much trust one has into the inputs. Having a single parameter to determine sensitivity to inputs keeps the amount of manual decisions to a minimum, and the convex nature of the optimization problem allows fast and robust solving.

For multi-period applications, the diversification-aware approach is combined with a dynamic risk management component that tries to keep portfolio risks stable despite the time-varying nature of asset moments. As response to increasing portfolio volatility, the risk management component triggers portfolio re-balancing into assets with lower volatility. This way, portfolio downside risks like maximum drawdowns are reduced, making the approach promising even beyond a perspective on mean-variance utility only. Frequent portfolio re-balancing stabilizes portfolio risks and hence prevents persistent exceedance of target risks during lasting crisis periods.

In a setting with transaction costs, frequent re-balancing can quickly become costly. The dynamic trading strategy hence is extended with an additional component that shall guarantee a cost-efficient implementation. Thereby the underlying idea is that we get a series of optimal target portfolios from single-period optimizations, which then will be iteratively adapted through weight filters. These weight filters shall realize a series of desired portfolio properties over time in a slightly different but less costly way. In a world without trading costs, we would fully realize single-period target portfolios in each period. With trading costs, however, we will use weight filters in order to modify original target portfolio weights such that they become cost-efficient but still sufficiently similar alternatives to the original targets. Or, as extreme case, target portfolios are modified so much that they are set equal to current portfolio weights, such that no trading will occur at all.

A particularly powerful combination of weight filters is the combination of minimum RTE filter and cost-aware capped RTE. The minimum RTE filter will trigger trading whenever current portfolio properties are not sufficiently similar to target portfolio properties anymore. And, when trading is triggered, one will not just re-balance current portfolio weights to target weights completely, but cost-aware capped RTE optimization will align the portfolio with the target portfolio only up to some required degree of relative tracking error. Effectively, final portfolios hence will always remain within certain bounds of relative tracking error, and this will be realized in a cost-efficient manner. In addition, tracking error optimization is formulated as convex optimization problem, such that it allows fast and robust solving.

In an empirical application we show that the dynamic strategy produces promising allocation decisions. The empirical application is based on globally diversified portfolios with investments into multiple asset classes. Diversification effects inherent to multi-asset portfolios can improve existing risk-return tradeoffs, and these potential benefits become even more pronounced in a setting of dynamic and risk-aware portfolio management. Provided adequate detection of time-varying risks, risks can be bypassed by shifting exposure into comparatively calm markets during times of financial turmoil. In a backtest period of approximately 17 years, we achieve out-of-sample risk-return profiles close to the true, in-sample efficient frontier. Even more remarkable, results hold although asset return moments are obtained from an unreflected application of exponentially weighted sample moments to assets of multiple asset classes. This estimator is predestined to generate poor portfolio optimization inputs, as it is backward-looking only and treats highly heterogeneous assets equally. Deliberately using an estimator with large imprecision will allow to examine the robustness of the allocation strategy with regards to faulty inputs. Still, backtest results are promising and hence indicate that diversification-aware portfolio choices have the potential to effectively deal with faulty inputs. In addition, dynamic risk management helps to considerably reduce drawdowns as compared to static portfolio weights, as the particularly cost-efficient implementation provides the dynamic risk management approach with enough scope for action. Hence, results are promising even beyond a perspective on mean-variance utility only. Comparisons of dynamic and static realizations of on average equal portfolio weights further emphasize the particularly promising outcomes of dynamic risk management.

Furthermore, dynamic strategies are set up with particular focus on real life feasibility, and hence reflect some aspects that go beyond the usual academic requirements. In order to address a broad range of clients, the strategy is built to be scalable with regards to realized risk levels: the strategy allows to match client's risk attitudes for a broad spectrum of levels of risk aversion. Furthermore, within each risk category clients are managed individually, in order to allow customization with regards to client-specific tax situations in a subsequent layer. In such a setting, portfolio allocations need to also keep client performance dispersion within reasonable bounds. However, both an examination of convergences of client portfolio weights and of dispersion of client performances indicate that this is fulfilled for the proposed dynamic allocation strategies.

Global investments into multi-asset portfolios, however, do open the door to currency risk. Hence, in the context of risk-aware portfolio management we also analyzed the effects of currency risk, particularly in order to make an informed decision on whether to hedge currency risks or not. Fully hedging against nominal exchange rate changes is not possible anyways, as inflation differentials have to be beared in the long run in either case. Either they will materialize in the form of exchange rate changes for the unhedged case, or they will slip in in the form of hedge returns when currencies are hedged. In the long run, hedging hence is primarily concerned with changes in real exchange rates, and they have proven to be rather stable over periods of multiple decades. Although purchasing power parity is not a reliable assumption in reality, it still constitutes a reasonable anchor that affects real exchange rates. Furthermore, measuring currency risk solely in nominal terms is showing an incomplete picture. Even a pure cash portfolio in homeland currency is not free from currency risk, as the currency might depreciate globally, thereby diminishing global purchasing power. Hence, diversifying over currencies might provide a good insurance against domestic economic and financial crises. This is something that U.K. investors had to experience during the recent GBP depreciation that came with the Brexit. Having exposure to multiple currencies will generally reduce such fluctuations in terms of global purchasing power, particularly when some of the foreign currency exposure is against currencies that generally are perceived to be "safe haven" currencies. For all these reasons, currency risks were left unhedged in the empirical application.

Based on the results of this thesis, as well as on other existing literature, there are a couple of directions were future research might further improve quantitative portfolio management. First and foremost, of course, any improvements in the modeling of asset returns should pay off in basically any framework. The more that we know about the true set of investment possibilities, the closer we will get to true optimal portfolio choices. In particular, with the exponentially weighted sample moments that were used to model asset moments in this thesis, there exists plenty of room for improvements. For example, GARCH models might be a better way to estimate conditional volatility, and further external information might be used to come up with more forward-looking estimates of expected returns (Faber 2012). Furthermore, asset specific patterns could be taken into account more explicitly. While exponentially weighted sample moments do enforce a similar structure for the estimators of all individual asset, assets should actually be treated differently, especially in a setting of multi-asset portfolios. For example, bond prices and commodity prices in many ways behave differently than equities, and this definitely should be taken into account.

Furthermore, several refinements of the diversification-aware single-period strategy are thinkable. Most encouraging seems to be any research that points into the direction of more explicit and more differentiated consideration of uncertainty. For example, one might try to associate different assets with different levels of uncertainty. This way the portfolio could be tilted less aggressively to promising assets when they are associated with higher uncertainty. Possible ways to implement this would be by either enforcing diversification on sub-groups of assets, or by including linear bounds on certain asset weights. In addition, sensitivity to inputs is determined by the minimum level of diversification that is required, and this diversification level is set externally. This possibly could be improved by making the actual choice of this level endogenous and dynamic. In other words, the level of diversification could be automatically set dependent on an actually estimated level of uncertainty around inputs.

In a multi-period setting, several refinements are thinkable for the dynamic risk management component. First, portfolio selections could be based on risk measures other than volatility, like for example Value-at-Risk. Second, the currently optimal portfolio does not need to be selected on the grounds of current risk estimates only, but it could be made dependent on the actual historic performance of an individual client also. For example, an existing client with currently large drawdown could get a less risky portfolio than a rather new client with low drawdown value. Such path dependent portfolio choices are usually already used in the context of portfolio insurance strategies. And third, instead of calibrating dynamic strategies such that they match desired levels of risk in the long run, one could try to aim for a desired risk target directly.

Besides the improvements on risk management, one also might try to further improve on the cost-efficient realization of target weights. In particular, relative tracking errors may not be the best measure of portfolio similarity yet, and other measures of deviating portfolio behavior could be tried (Rachev, Stoyanov, and Fabozzi 2008). In addition, moving averages of portfolio weights that were used as proxy for the persistence of investment opportunities certainly are a quite heuristic approach. It might pay off to estimate persistence more explicitly.

With a more pronounced focus on real world implementation and practical problems, the perspective on dynamic strategies should be extended to also take additional practical challenges into regard. For example, the universe of assets might not be fixed over long time periods anymore, but individual assets might both be added or removed. This will come with the downside of causing additional trading costs, and hence it should pay off to initiate research on how to optimally conduct changes of the underlying universe. For this task, again, measures of portfolio similarity might be useful in pointing out effective ways of how to resolve the tradeoff between trading costs and portfolio improvements. Another aspect with huge potential impact on final profits and losses is the nature of tax liabilities that clients have to face. Depending on the tax system, as well as the individual client's tax situation, tax payments can substantially affect overall profits, and hence an optimal asset allocation strategy should automatically take tax effects into account. Ideally, tax information should be used to implement a full-fledged tax loss harvesting approach. Besides the desired changes to portfolio allocations that are conducted by the dynamic asset allocation strategy, in reality there usually is another potential source of portfolio changes: the client himself. Client initiated cash-flows, i.e. incoming deposits or desired withdrawals, should be met in a cost-efficient manner. For example, new deposits could be used to only buy assets that are underweight, thereby automatically bringing current portfolio weights closer to target weights. Again, the measure used to evaluate portfolio similarity plays a vital role here. Likewise, it also should influence the actual generation of orders, where fractional target weights can not be exactly realized, but have to be translated into integer valued volumes. In both cases, portfolio distance to target weights should be minimized, with constraints that are given by client cash-flows for the first case, and with integer value constraints for the case of order generation.

Chapter 6

Appendix

Appendix 6.0.1: [**Section (1.2.3): Scaling of moments**]

Proof of Lemma (1.2.3) to determine logarithmic return moments from discrete return moments.

Proof. Let's first reformulate Equation (1.3):

$$\begin{aligned}\mu(r) &= \exp\left(\mu(r^{log}) + \frac{\sigma^2(r^{log})}{2}\right) - 1 \\ \log(\mu(r)+1) &= \mu(r^{log}) + \frac{\sigma^2(r^{log})}{2} \qquad\qquad (\star)\end{aligned}$$

Reformulating Equation (1.4) and using Equation ($\star$):

$$\begin{aligned}\sigma^2(r) &= \left(\exp\left(\sigma^2(r^{log})\right) - 1\right)\exp\left(2\mu(r^{log}) + \sigma^2(r^{log})\right) \qquad \Leftrightarrow \\ \log(\sigma^2(r)) &= \log\left(\exp\left(\sigma^2(r^{log})\right) - 1\right) + 2\mu(r^{log}) + \sigma^2(r^{log}) \\ &= \log\left(\exp\left(\sigma^2(r^{log})\right) - 1\right) + 2\left(\mu(r^{log}) + \frac{\sigma^2(r^{log})}{2}\right) \\ &\overset{(\star)}{=} \log\left(\exp\left(\sigma^2(r^{log})\right) - 1\right) + 2\log(\mu(r)+1) \\ &= \log\left(\exp\left(\sigma^2(r^{log})\right) - 1\right) + \log((\mu(r)+1)^2)\end{aligned}$$

Solving for $\sigma^2(r^{\log})$:

$$
\begin{aligned}
\log\left(\exp\left(\sigma^2(r^{log})\right)-1\right) &= \log(\sigma^2(r)) - \log((\mu(r)+1)^2) \\
&= \log\left(\frac{\sigma^2(r)}{(\mu(r)+1)^2}\right) && \Leftrightarrow \\
\exp\left(\sigma^2(r^{log})\right)-1 &= \frac{\sigma^2(r)}{(\mu(r)+1)^2} && \Leftrightarrow \\
\sigma^2(r^{log}) &= \log\left(\frac{\sigma^2(r)}{(\mu(r)+1)^2}+1\right)
\end{aligned}
$$

Plugging into ($\star$) and solving for $\mu(r^{\log})$:

$$
\begin{aligned}
\mu(r^{log}) + \frac{\sigma^2(r^{log})}{2} &= \log(\mu(r)+1) \\
\mu(r^{log}) &= \log(\mu(r)+1) - \frac{\sigma^2(r^{log})}{2} \\
\mu(r^{log}) &= \log(\mu(r)+1) - \frac{1}{2}\log\left(1+\frac{\sigma^2(r)}{(\mu(r)+1)^2}\right) \\
&= \log(\mu(r)+1) - \log\left(\left(\frac{(\mu(r)+1)^2+\sigma^2(r)}{(\mu(r)+1)^2}\right)^{\frac{1}{2}}\right) \\
&= \log\left((\mu(r)+1)\left(\frac{(\mu(r)+1)^2}{(\mu(r)+1)^2+\sigma^2(r)}\right)^{\frac{1}{2}}\right) \\
&= \log\left(\frac{(\mu(r)+1)^2}{\sqrt{(\mu(r)+1)^2+\sigma^2(r)}}\right)
\end{aligned}
$$

$\square$

Appendix 6.0.2: [**Section (2.2): Fully uninformed market view**]

Reformulation of the variance of a global minimum variance portfolio under fully uninformed market view for Theorem (2.2.2), part 1:

$$
\begin{aligned}
\frac{1}{(\bar{a}^2+(d-1)\bar{a}_{ij})} &= \frac{1}{\left(\frac{1-(d-1)\bar{\sigma}_{ij}\bar{a}_{ij}}{\bar{\sigma}^2}+(d-1)\bar{a}_{ij}\right)} \\
&= \frac{1}{\left(\frac{1-(d-1)\bar{\sigma}_{ij}\bar{a}_{ij}+\bar{\sigma}^2(d-1)\bar{a}_{ij}}{\bar{\sigma}^2}\right)} \\
&= \frac{\bar{\sigma}^2}{1+(d-1)\bar{a}_{ij}(\bar{\sigma}^2-\bar{\sigma}_{ij})} \\
&= \frac{\bar{\sigma}^2}{1-(d-1)(\bar{\sigma}^2-\bar{\sigma}_{ij})\frac{\bar{\sigma}_{ij}}{\bar{\sigma}^4+(d-2)\bar{\sigma}_{ij}\bar{\sigma}^2-\bar{\sigma}_{ij}^2(d-1)}} \\
&= \frac{\bar{\sigma}^2}{\frac{(\bar{\sigma}^4+(d-2)\bar{\sigma}_{ij}\bar{\sigma}^2-\bar{\sigma}_{ij}^2(d-1))-(d-1)(\bar{\sigma}^2-\bar{\sigma}_{ij})\bar{\sigma}_{ij}}{\bar{\sigma}^4+(d-2)\bar{\sigma}_{ij}\bar{\sigma}^2-\bar{\sigma}_{ij}^2(d-1)}} \\
&= \frac{(\bar{\sigma}^4+(d-2)\bar{\sigma}_{ij}\bar{\sigma}^2-\bar{\sigma}_{ij}^2(d-1))\bar{\sigma}^2}{\bar{\sigma}^4+(d-2)\bar{\sigma}_{ij}\bar{\sigma}^2-\bar{\sigma}_{ij}^2(d-1)+(d-1)\bar{\sigma}_{ij}^2-(d-1)\bar{\sigma}^2\bar{\sigma}_{ij}} \\
&= \frac{(\bar{\sigma}^4+(d-2)\bar{\sigma}_{ij}\bar{\sigma}^2-(d-1)\bar{\sigma}_{ij}^2)\bar{\sigma}^2}{\bar{\sigma}^4+(d-2)\bar{\sigma}_{ij}\bar{\sigma}^2-(d-1)\bar{\sigma}^2\bar{\sigma}_{ij}} \\
&= \frac{\bar{\sigma}^2(\bar{\sigma}^4+(d-2)\bar{\sigma}_{ij}\bar{\sigma}^2-(d-1)\bar{\sigma}_{ij}^2)}{\bar{\sigma}^2(\bar{\sigma}^2+(d-2)\bar{\sigma}_{ij}-(d-1)\bar{\sigma}_{ij})} \\
&= \frac{(\bar{\sigma}^4+(d-2)\bar{\sigma}_{ij}\bar{\sigma}^2-(d-1)\bar{\sigma}_{ij}^2)}{\bar{\sigma}^2+\bar{\sigma}_{ij}((d-2)-(d-1))} \\
&= \frac{(\bar{\sigma}^4+(d-2)\bar{\sigma}_{ij}\bar{\sigma}^2-(d-1)\bar{\sigma}_{ij}^2)}{\bar{\sigma}^2-\bar{\sigma}_{ij}} \\
&= \frac{(\bar{\sigma}^4-(d-1)\bar{\sigma}_{ij}^2)}{\bar{\sigma}^2-\bar{\sigma}_{ij}}+\frac{(d-2)\bar{\sigma}_{ij}\bar{\sigma}^2}{\bar{\sigma}^2-\bar{\sigma}_{ij}} \\
&= \frac{(\bar{\sigma}^4-(d-1)\bar{\sigma}_{ij}^2)}{\bar{\sigma}^2-\bar{\sigma}_{ij}}+\frac{\bar{\sigma}^2\bar{\sigma}_{ij}(d-2)}{\bar{\sigma}^2-\bar{\sigma}_{ij}} \\
&= \frac{(\bar{\sigma}^4-(d-1)\bar{\sigma}_{ij}^2)}{\bar{\sigma}^2-\bar{\sigma}_{ij}}+\frac{\bar{\sigma}^2\bar{\sigma}_{ij}((d-1)-1)}{\bar{\sigma}^2-\bar{\sigma}_{ij}} \\
&= \frac{(\bar{\sigma}^4-(d-1)\bar{\sigma}_{ij}^2)}{\bar{\sigma}^2-\bar{\sigma}_{ij}}+\frac{(d-1)\bar{\sigma}_{ij}\bar{\sigma}^2-\bar{\sigma}^2\bar{\sigma}_{ij}}{\bar{\sigma}^2-\bar{\sigma}_{ij}} \\
&= \frac{\bar{\sigma}^4+(d-1)\bar{\sigma}_{ij}\bar{\sigma}^2-\bar{\sigma}^2\bar{\sigma}_{ij}-(d-1)\bar{\sigma}_{ij}^2}{\bar{\sigma}^2-\bar{\sigma}_{ij}} \\
&= \frac{(\bar{\sigma}^2+(d-1)\bar{\sigma}_{ij})(\bar{\sigma}^2-\bar{\sigma}_{ij})}{\bar{\sigma}^2-\bar{\sigma}_{ij}} \\
&= \bar{\sigma}^2+(d-1)\bar{\sigma}_{ij}
\end{aligned}
$$

Appendix 6.0.3: [**Section (4.5.2): Strategy results**]

The following table lists all backtested strategies. Column **Single-period** refers to the single-period strategy that was used as part of the respective dynamic strategy. Possibly chosen persistence filters are given in column **Endurance**, while up to two cost-reduction filters can be given in columns **CR 1** and **CR 2**.

Single-period	Endurance	CR 1	CR 2
Div. aware(0.6)		$h^{\mathsf{FixFreq}}(1)$	
Div. aware(0.6)		$h^{\mathsf{FixFreq}}(20)$	
Div. aware(0.6)		$h^{\mathsf{FixFreq}}(25)$	
Div. aware(0.6)		$h^{\mathsf{FixFreq}}(40)$	
Div. aware(0.6)		$h^{\mathsf{FixFreq}}(5)$	
Div. aware(0.6)	$h^{\mathsf{EWMA}}(0.92)$		
Div. aware(0.6)	$h^{\mathsf{EWMA}}(0.95)$		
Div. aware(0.6)	$h^{\mathsf{EWMA}}(0.995)$		
Div. aware(0.6)	$h^{\mathsf{EWMA}}(0.99)$		
Div. aware(0.6)	$h^{\mathsf{EWMA}}(0.991)$	$h^{\mathsf{FixFreq}}(5)$	$h^{\mathsf{TOmin}}(0.1)$
Div. aware(0.6)	$h^{\mathsf{MA}}(100)$		
Div. aware(0.6)	$h^{\mathsf{MA}}(10)$		
Div. aware(0.6)	$h^{\mathsf{MA}}(150)$		
Div. aware(0.6)	$h^{\mathsf{MA}}(200)$		
Div. aware(0.6)	$h^{\mathsf{MA}}(20)$		
Div. aware(0.6)	$h^{\mathsf{MA}}(40)$		
Div. aware(0.6)	$h^{\mathsf{MA}}(5)$		
Div. aware(0.6)	$h^{\mathsf{MA}}(60)$		
Div. aware(0.6)	$h^{\mathsf{MA}}(80)$		
Div. aware(0.6)	$h^{\mathsf{MA}}(40)$	$h^{\mathsf{RTEcap}}(0.12)$	
Div. aware(0.6)	$h^{\mathsf{MA}}(40)$	$h^{\mathsf{RTEcap}}(0.14)$	
Div. aware(0.6)	$h^{\mathsf{MA}}(40)$	$h^{\mathsf{RTEcap}}(0.18)$	
Div. aware(0.6)	$h^{\mathsf{MA}}(40)$	$h^{\mathsf{RTEcap}}(0.1)$	
Div. aware(0.6)	$h^{\mathsf{MA}}(40)$	$h^{\mathsf{RTEmin}}(0.14)$	$h^{\mathsf{RTEcap}}(0.08)$
Div. aware(0.6)	$h^{\mathsf{MA}}(40)$	$h^{\mathsf{RTEmin}}(0.16)$	$h^{\mathsf{RTEcap}}(0.12)$
Div. aware(0.6)	$h^{\mathsf{MA}}(40)$	$h^{\mathsf{RTEmin}}(0.18)$	$h^{\mathsf{RTEcap}}(0.14)$
Div. aware(0.6)	$h^{\mathsf{MA}}(40)$	$h^{\mathsf{RTEmin}}(0.1)$	$h^{\mathsf{RTEcap}}(0.05)$
Div. aware(0.6)	$h^{\mathsf{MA}}(40)$	$h^{\mathsf{RTEmin}}(0.22)$	$h^{\mathsf{RTEcap}}(0.18)$
Div. aware(0.6)	$h^{\mathsf{MA}}(40)$	$h^{\mathsf{RTEmin}}(0.2)$	$h^{\mathsf{RTEcap}}(0.14)$
Div. aware(0.6)	$h^{\mathsf{MA}}(60)$	$h^{\mathsf{RTEmin}}(0.22)$	$h^{\mathsf{RTEcap}}(0.18)$
Div. aware(0.8)		$h^{\mathsf{FixFreq}}(1)$	
Div. aware(0.8)		$h^{\mathsf{FixFreq}}(20)$	
Div. aware(0.8)		$h^{\mathsf{FixFreq}}(40)$	
Div. aware(0.8)		$h^{\mathsf{FixFreq}}(5)$	
Div. aware(0.8)	$h^{\mathsf{EWMA}}(0.92)$		
Div. aware(0.8)	$h^{\mathsf{EWMA}}(0.95)$		
Div. aware(0.8)	$h^{\mathsf{EWMA}}(0.995)$		

Div. aware(0.8)	$h^{\mathsf{EWMA}}(0.99)$		
Div. aware(0.8)	$h^{\mathsf{MA}}(100)$		
Div. aware(0.8)	$h^{\mathsf{MA}}(10)$		
Div. aware(0.8)	$h^{\mathsf{MA}}(150)$		
Div. aware(0.8)	$h^{\mathsf{MA}}(200)$		
Div. aware(0.8)	$h^{\mathsf{MA}}(20)$		
Div. aware(0.8)	$h^{\mathsf{MA}}(40)$		
Div. aware(0.8)	$h^{\mathsf{MA}}(5)$		
Div. aware(0.8)	$h^{\mathsf{MA}}(60)$		
Div. aware(0.8)	$h^{\mathsf{MA}}(80)$		
Div. aware(0.8)	$h^{\mathsf{MA}}(40)$	$h^{\mathsf{RTEcap}}(0.18)$	
Div. aware(0.8)	$h^{\mathsf{MA}}(40)$	$h^{\mathsf{RTEmin}}(0.06)$	$h^{\mathsf{RTEcap}}(0.03)$
Div. aware(0.8)	$h^{\mathsf{MA}}(40)$	$h^{\mathsf{RTEmin}}(0.08)$	$h^{\mathsf{RTEcap}}(0.04)$
Div. aware(0.8)	$h^{\mathsf{MA}}(40)$	$h^{\mathsf{RTEmin}}(0.12)$	$h^{\mathsf{RTEcap}}(0.08)$
Div. aware(0.8)	$h^{\mathsf{MA}}(40)$	$h^{\mathsf{RTEmin}}(0.1)$	$h^{\mathsf{RTEcap}}(0.05)$
Div. aware(0.8)	$h^{\mathsf{MA}}(60)$	$h^{\mathsf{RTEmin}}(0.06)$	$h^{\mathsf{RTEcap}}(0.03)$
Div. aware(0.8)	$h^{\mathsf{MA}}(60)$	$h^{\mathsf{RTEmin}}(0.08)$	$h^{\mathsf{RTEcap}}(0.04)$
Div. aware(0.8)	$h^{\mathsf{MA}}(60)$	$h^{\mathsf{RTEmin}}(0.12)$	$h^{\mathsf{RTEcap}}(0.08)$
Div. aware(0.8)	$h^{\mathsf{MA}}(60)$	$h^{\mathsf{RTEmin}}(0.1)$	$h^{\mathsf{RTEcap}}(0.05)$
Markowitz		$h^{\mathsf{FixFreq}}(1)$	
Markowitz		$h^{\mathsf{FixFreq}}(20)$	
Markowitz		$h^{\mathsf{FixFreq}}(40)$	
Markowitz		$h^{\mathsf{FixFreq}}(5)$	
Markowitz	$h^{\mathsf{EWMA}}(0.92)$		
Markowitz	$h^{\mathsf{EWMA}}(0.95)$		
Markowitz	$h^{\mathsf{EWMA}}(0.995)$		
Markowitz	$h^{\mathsf{EWMA}}(0.99)$		
Markowitz	$h^{\mathsf{MA}}(100)$		
Markowitz	$h^{\mathsf{MA}}(10)$		
Markowitz	$h^{\mathsf{MA}}(150)$		
Markowitz	$h^{\mathsf{MA}}(200)$		
Markowitz	$h^{\mathsf{MA}}(20)$		
Markowitz	$h^{\mathsf{MA}}(40)$		
Markowitz	$h^{\mathsf{MA}}(5)$		
Markowitz	$h^{\mathsf{MA}}(60)$		
Markowitz	$h^{\mathsf{MA}}(80)$		
Markowitz	$h^{\mathsf{MA}}(40)$	$h^{\mathsf{RTEcap}}(0.2)$	

Appendix 6.0.4: [**Section (4.5.3): Analysis of single selected strategy**]

The following supplementary charts show absolute and relative risk contributions of individual assets for risk categories three and thirteen.

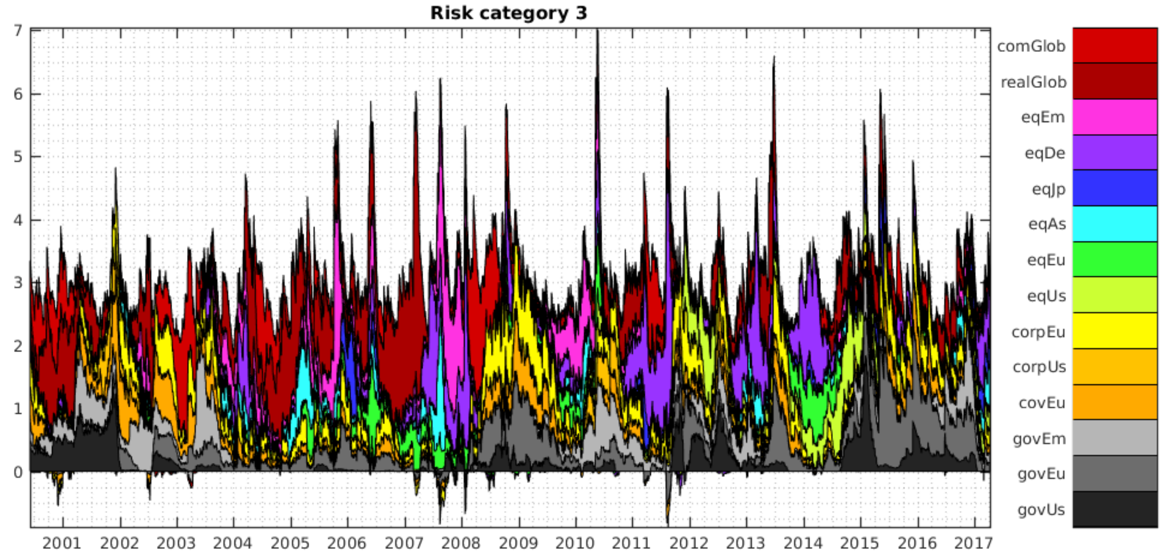

Figure 6.1: Volatility contributions over time

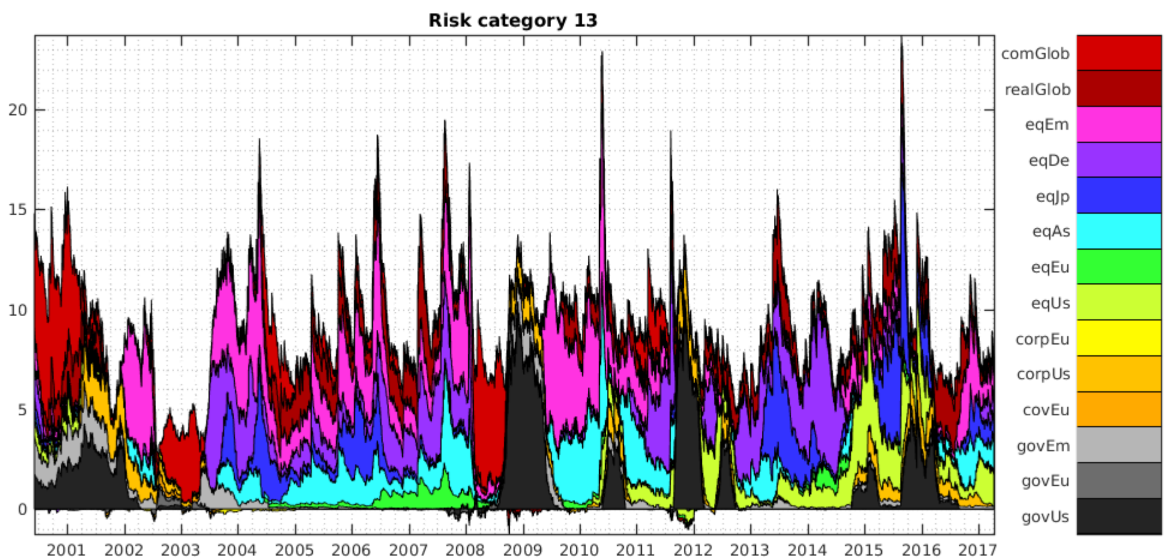

Figure 6.2: Volatility contributions over time

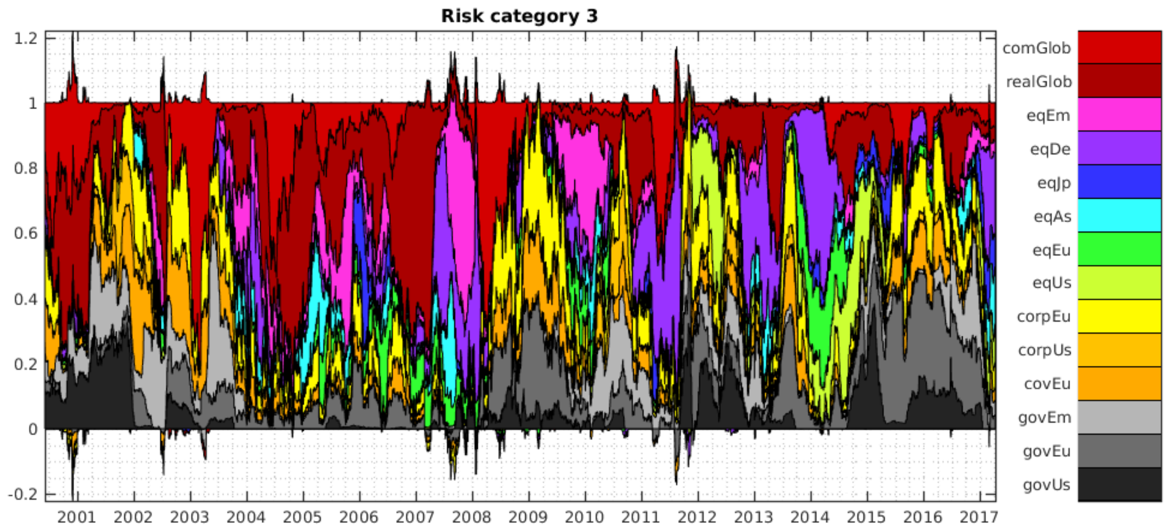

Figure 6.3: Relative volatility contributions over time

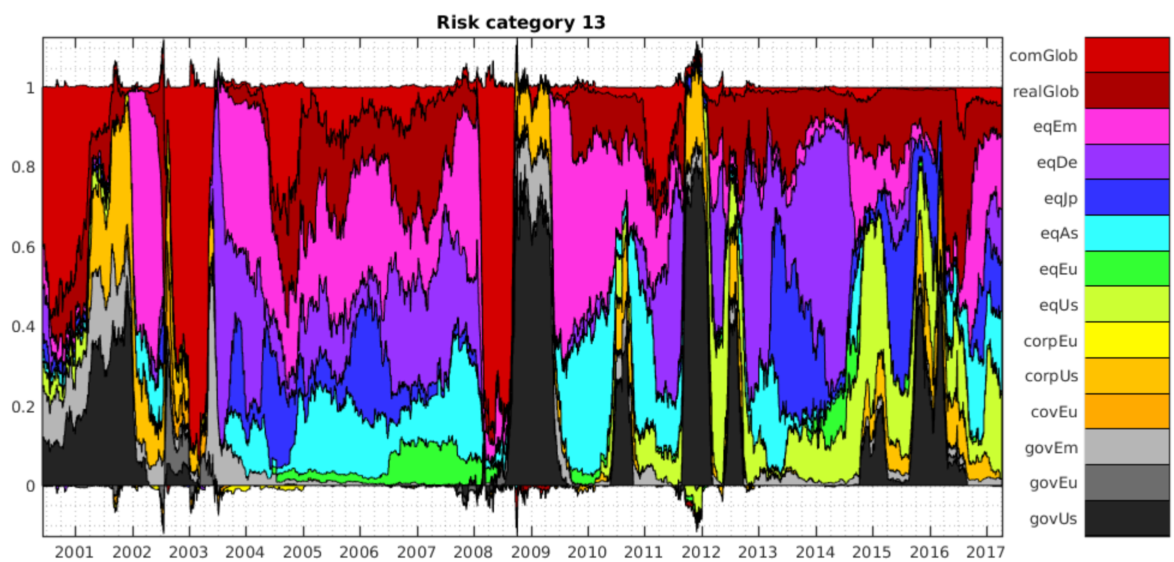

Figure 6.4: Relative volatility contributions over time

Appendix 6.0.5: [**Section (4.5.6): Comparison to static portfolio weights**]

Supplementary charts for the comparison of dynamic strategy results with static portfolio weights equal to realized average weights of dynamic strategies. Performance and drawdown series resulting from static average weights are shown over time in Figure (6.5).

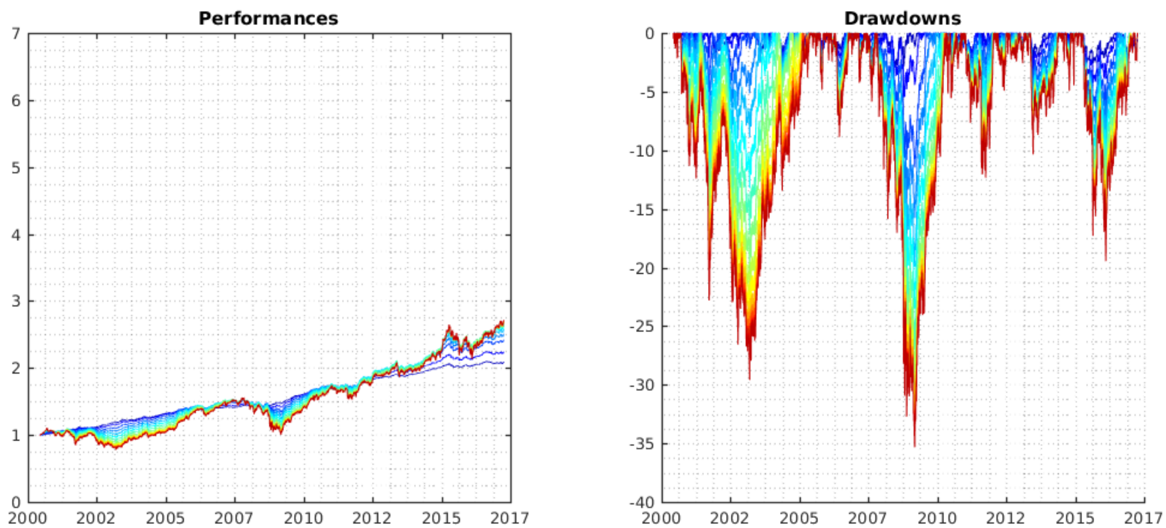

Figure 6.5: Static weight strategy performances and drawdowns

Figure (6.6) and Figure (6.6) show the effects of dynamic portfolio risk management on maximum drawdowns for risk categories three and thirteen similar to Figure (4.23).

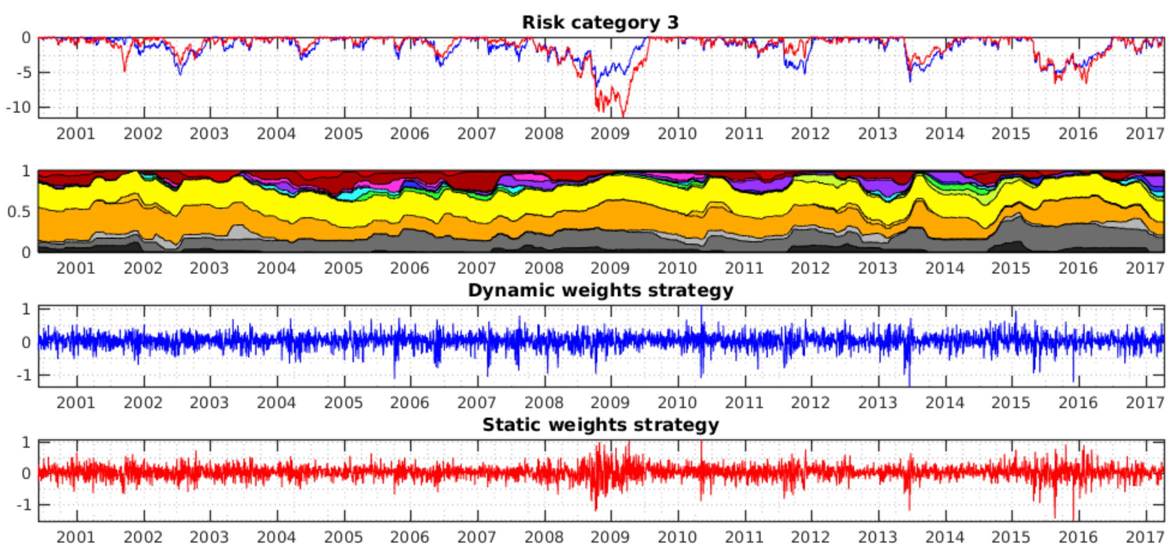

Figure 6.6: Drawdowns, daily strategy returns and weights over time

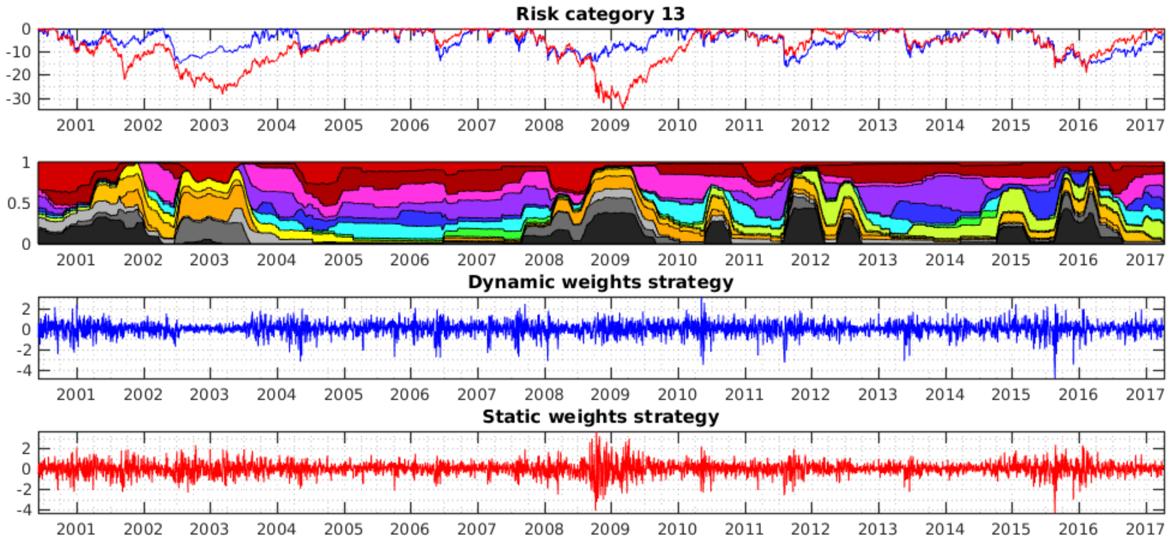

Figure 6.7: Drawdowns, daily strategy returns and weights over time

Appendix 6.0.6: [Section (4.5.8): Client dispersion]

Supplementary chart to analyze client dispersion.

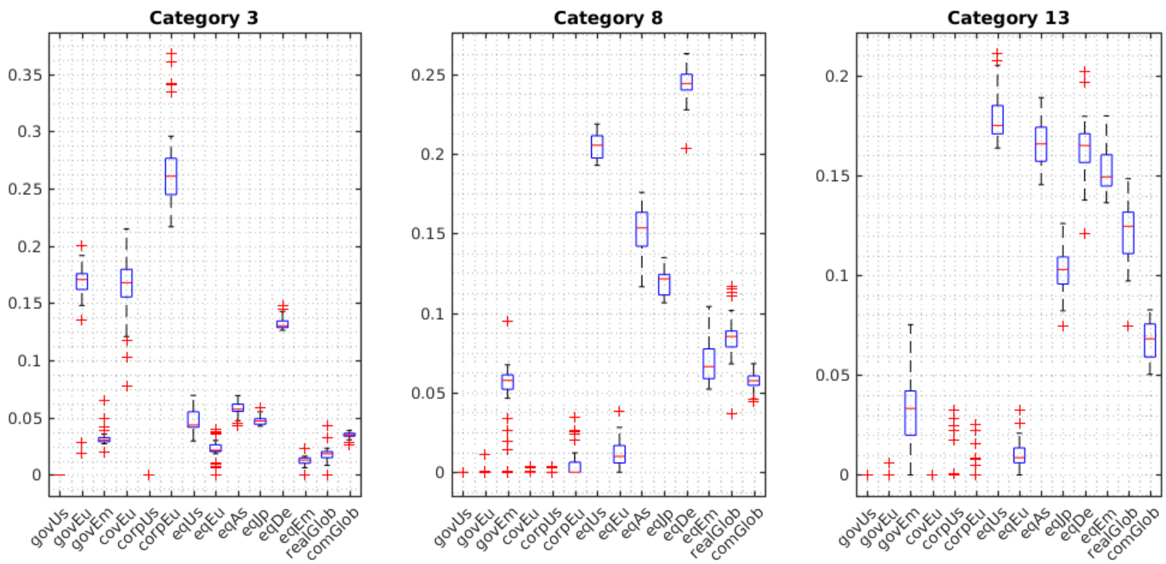

Figure 6.8: Box plots of final asset weights for multiple clients within selected risk categories

References

Ang, Andrew. 2014. *Asset Management: A Systematic Approach to Factor Investing*. Oxford University Press.

Ang, Andrew, Robert j. Hodrick, Yuhang Xing, and Xiaoyan Zhang. 2006. "The Cross-Section of Volatility and Expected Returns." *The Journal of Finance* 61 (1). Blackwell Publishing, Inc.: 259–99.

Avramov, Doron, and Guofu Zhou. 2010. "Bayesian Portfolio Analysis." *Annu. Rev. Financ. Econ.* 2 (1). Annual Reviews: 25–47.

Bae, Gil S. 1999. "Do Stock Prices Fully Reflect Information in Current Earnings, Cash Flows Accruals?: Evidence from Quarterly Data." *Asia-Pacific Journal of Accounting* 6 (2): 161–75.

Banz, Rolf W. 1981. "The Relationship Between Return and Market Value of Common Stocks." *Journal of Financial Economics*, 3–18.

Basu, S. 1977. "Investment Performance of Common Stocks in Relation to Their Price-Earnings Ratios: A Test of the Efficient Market Hypothesis." *Journal of Finance* 32 (3): 663–82.

Black, Fischer. 1995. "Universal Hedging: Optimizing Currency Risk and Reward in International Equity Portfolios." *Financial Analysts Journal* 51 (1): 161–67.

Black, Fischer, and Robert Litterman. 1990. "Asset Allocation: Combining Investor Views with Market Equilibrium." Goldman, Sachs & Co.

Bollerslev, Tim. 1986. "Generalized Autoregressive Conditional Heteroskedasticity." *Journal of Econometrics* 31: 307–27.

Bundesbank, Deutsche. 2010. "Nominal and Real Exchange Rate Movements During the Financial Crisis." Monthly report of the Deutsche Bundesbank. Frankfurt (Main): Deutsche Bundesbank.

Burns, Patrick, Robert F Engle, and Joseph J Mezrich. 1998. "Correlations and Volatilities of Asynchronous Data." *The Journal of Derivatives* 5 (4). Institutional Investor Journals: 7–18.

Carhart, Mark M. 1997. "On Persistence in Mutual Fund Performance." *The Journal of*

Finance 52 (1). Blackwell Publishing Ltd: 57–82.

Catao, Luis A. V. 2007. "Why Real Exchange Rates?" *Finance and Development* 44 (3).

Charles, Thomas, and Paul M. Bosse. 2014. "Understanding the Hedge Return: The Impact of Currency Hedging in Foreign Bonds." Vanguard Research. Vanguard.

Chen, Nai-Fu, Richard Roll, and Stephen Ross. 1986. "Economic Forces and the Stock Market." *The Journal of Business* 59 (3): 383–403.

Cochrane, John H. 2011. "Presidential Address: Discount Rates." *The Journal of Finance* 66 (4). Blackwell Publishing Inc: 1047–1108.

Cornuejols, G., and R. Tütüncü. 2006. *Optimization Methods in Finance*. Mathematics, Finance and Risk. Cambridge University Press.

DeMiguel, Victor, Lorenzo Garlappi, and Raman Uppal. 2007. "Optimal Versus Naive Diversification: How Inefficient Is the 1/N Portfolio Strategy?" *The Review of Financial Studies* 22 (5). Oxford University Press: 1915–53.

Dimson, Elroy, Paul Marsh, and Mike Staunton. 2002. *Triumph of the Optimists: 101 Years of Global Investment Returns*. Princeton University Press.

Faber, Meb. 2012. "Global Value: Building Trading Models with the 10 Year Cape." Vol. 5. Cambria Quantitative Research.

Fama, Eugene, and Kenneth French. 1988. "Dividend Yields and Expected Stock Returns." *Journal of Financial Economics* 22 (1): 3–25.

———. 1993. "Common Risk Factors in the Returns on Stocks and Bonds." *Journal of Financial Economics* 33 (1): 3–56.

Ghaoui, Laurent El, Maksim Oks, and Francois Oustry. 2003. "Worst-Case Value-at-Risk and Robust Portfolio Optimization: A Conic Programming Approach." *Operations Research* 51 (4): 543–56.

Graham, B., and D.L.F. Dodd. 1934. *Security Analysis: The Classic 1934 Edition*. McGraw-Hill Education.

Hull, J. 2009. *Options, Futures and Other Derivatives*. Pearson/Prentice Hall.

Ilmanen, A., and C. Asness. 2011. *Expected Returns: An Investor's Guide to Harvesting Market Rewards*. The Wiley Finance Series. Wiley.

Jegadeesh, Narasimhan, and Sheridan Titman. 1993. "Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency." *The Journal of Finance* 48 (1). Blackwell Publishing Ltd: 65–91.

Johnson, N.L., S. Kotz, and N. Balakrishnan. 1995. *Continuous Univariate Distributions*. Wiley Series in Probability and Mathematical Statistics: Applied Probability and Statistics, v. 2. Wiley & Sons.

Kempf, Alexander, and Christoph Memmel. 2005. "On the estimation of the global minimum variance portfolio." CFR Working Papers 05-02. University of Cologne, Centre for Financial Research (CFR).

Ledoit, Olivier, and Michael Wolf. 2004. "Honey, I Shrunk the Sample Covariance Matrix." *The Journal of Portfolio Management* 30 (4): 110–19.

Levy, H, and Harry M. Markowitz. 1979. "Approximating Expected Utility by a Function of Mean and Variance." *American Economic Review* 69 (3): 308–17.

Lintner, John. 1965. "The Valuation of Risk Assets and the Selection of Risky Investments in Stock Portfolios and Capital Budgets." *The Review of Economics and Statistics* 47 (1): 13–37.

Maillard, Sébastien, Thierry Roncalli, and Jérôme Teïletche. 2010. "On the Properties of Equally Weighted Risk Contribution Portfolios." *The Journal of Portfolio Management* 36 (July): 60–70.

Mandelbrot, Benoît. 1963. "The Variation of Certain Speculative Prices." *The Journal of Business* 36.

Mankiw, N.G. 2008. *Principles of Economics*. Available Titles Coursemate Series, v. 1. Cengage Learning.

Markowitz, Harry M. 1952. "Portfolio Selection." *Journal of Finance* 7 (1): 77–91.

———. 1956. "The Optimization of a Quadratic Function Subject to Linear Constraints." *Naval Research Logistics Quarterly* 3 (1-2). Wiley Subscription Services, Inc., A Wiley Company: 111–33.

———. 1959. *Portfolio Selection: Efficient Diversification of Investments*. Yale University Press.

———. 1990. "Prize Lecture: Foundations of Portfolio Theory." *Nobelprize.org*.

Meese, Richard A., and Kenneth Rogoff. 1983. "Empirical Exchange Rate Models of the Seventies: Do They Fit Out of Sample?" *Journal of International Economics* 14: 3–24.

Michaud, Richard O., and Robert O. Michaud. 2008. *Efficient Asset Management: A Practical Guide to Stock Portfolio Optimization and Asset Allocation*. 2nd ed. Financial Management Association Survey and Synthesis Series. Oxford University Press, USA.

Mossin, Jan. 1966. "Wages, Profits, and the Dynamics of Growth." *The Quarterly Journal of Economics* 80 (3): 376–99.

Perold, André, and Evan C. Schulman. 1988. "The Free Lunch in Currency Hedging: Implications for Investment Policy and Performance." *Financial Analysts Journal* 44 (3): 45–50.

Rachev, Svetlozar T., Stoyan V. Stoyanov, and Frank J. Fabozzi. 2008. *Advanced Stochastic Models, Risk Assessment, and Portfolio Optimization: The Ideal Risk, Uncertainty, and*

Performance Measures. John Wiley; Sons.

Rieder, Rick. 2016. "Navigating today's rolling crises." *Fixed Income Market Outlook*.

"RiskMetrics Technical Document." 1996. J.P. Morgan; Reuters.

Roll, Richard, and John E. Anderson. 1992. *A Mean/Variance Analysis of Tracking Error*. Finance Working Paper. UCLA.

Roncalli, T. 2013. *Introduction to Risk Parity and Budgeting*. Chapman and Hall/Crc Financial Mathematics Series. Taylor & Francis.

Ross, Stephen A. 1976. "The arbitrage theory of capital asset pricing." *Journal of Economic Theory* 13 (3): 341–60.

Scherer, Bernd. 2002. "Portfolio resampling: Review and critique." *Financial Analysts Journal* 58 (6): 98–109.

Sharpe, William. 1964. "Capital Asset Prices: A Theory of Market Equilibrium Under Conditions of Risk." *Journal of Finance* 19 (3): 425–42.

Sinn, Hans-Werner. 1989. "Expected Utility and the Siegel Paradox." *Journal of Economics* 50 (3): 257–68.

Tobin, James. 1958. "Liquidity Preference as Behavior Towards Risk." *Review of Economic Studies* 25 (2): 65–86.

Treynor, Jack L. 1961. "Market Value, Time, and Risk." Unpublished manuscript.

Williams, J.B. 1938. *The Theory of Investment Value*. Harvard University Press.

Würtz, Diethelm, Yohan Chalabi, William Chen, and Andrew Ellis. 2009. *Portfolio Optimization with R/Rmetrics*. Rmetrics.

Eidesstattliche Versicherung

(Siehe Promotionsordnung vom 12. Juli 2011, § 8 Abs. 2 Pkt. 5)

Hiermit erkläre ich an Eides statt, dass die Dissertation von mir selbstständig, ohne unerlaubte Beihilfe angefertigt ist.

München, den 31. August 2017

Christian Groll