University of St. Gallen
Graduate School of Business,
Economics, Law and Social Sciences

Master's Thesis

Risk-Based Portfolio Construction for Retail Investors

Supervisor:

Prof. Ph.D. Enrico De Giorgi

Co-Supervisor:

Prof. Dr. Matthias Fengler

Submitted by:

Ferid Dupljak
ferid.dupljak@student.unisg.ch
12-611-612

February 10, 2020

Abstract

Several studies find that the performance of retail investors is poor due to psychological pitfalls which lead to suboptimal investor behavior. In this thesis, I exclude human behavior from the asset allocation decision and let the numbers do the work. I implement and investigate various risk-based portfolio construction techniques that rest on mathematical and statistical concepts. Using daily data (01/1999 - 12/2019) of six indexes that cover different asset classes and are investible through ETFs, I conduct out-of-sample backtests for yearly, quarterly, monthly and weekly rebalancings. The portfolio moments are estimated using robust methods and a novel multivariate worst-case mixture copula approach is investigated for CVaR-based portfolio construction techniques. The empirical analysis shows that complexity pays off: All risk-based portfolio optimizations can effectively reduce the portfolio risk compared to a naive 1/N portfolio. Moreover, using multidimensional mixture copulas in a worst-case framework leads to more conservative asset allocations for the Minimum CVaR approach and to higher returns for the Minimum CVaR Concentration approach, respectively.

Zusammenfassung

Zahlreiche Studien belegen, dass psychologische Faktoren der Grund sind, warum sich Privatinvestoren suboptimal verhalten und dadurch eine schlechte Performance in den Finanzmärkten erzielen. In dieser These schliesse ich die menschliche Komponente von der Anlageentscheidung aus und überlasse den Zahlen die Arbeit. Ich implementiere und untersuche unterschiedliche risikobasierte Portfoliokonstruktionstechniken, welche auf mathematischen und statistischen Konzepten basieren. Mithilfe täglicher Finanzdaten (01/1999 - 12/2019) von sechs unterschiedlichen Anlageklassen-Indizes, welche über ETFs zugänglich sind, führe ich Out-Of-Sample Backtests für jährliche, quartalsweise, monatliche und wöchentliche Rebalancings durch. Die Portfoliomomente werden durch robuste Methoden geschätzt und ein neuartiger multidimensionaler Worst-Case Mixture-Copula Ansatz wird für CVaR-basierte Portfoliokonstruktionstechniken analysiert. Die empirische Analyse zeigt, dass sich die Komplexität auszahlt: Alle risikobasierten Portfoliooptimierungen erzielen tiefere Portfoliorisiken im Vergleich zum simplen 1/N Portfolio. Zudem zeigt sich, dass die Nutzung von multidimensionalen Mixture Copulas im Worst-Case Framework bei Minimierung des CVaR zu konservativeren Portfolioallokationen führt. Wird hingegen der Beitrag zum CVaR optimiert, führt der Worst-Case Mixture Copula Ansatz zu deutlich höheren Renditen.

Contents

List of Tables

List of Figures

1 Introduction

With unprecedented access to financial markets, data and real-time news, retail investors seem to be in a promising position to take charge of their financial destiny without relying on costly services provided by professionals in the industry. Recent developments in financial technology as well as the emergence of financial start-ups have challenged traditional business models and have democratized the financial markets by lowering costs and barriers. Moreover, relatively young financial products such as exchange-traded funds (ETFs) provide retail investors with an easy and inexpensive access to a broad spectrum of asset classes, geographies and strategies. Despite the levelled playing fields between Wall Street and Main Street, retail investors continue to be bad at investing. Studies show that retail investors are subject to numerous psychological pitfalls that make them behave suboptimally. This, in turn, leads to the constant underperformance of retail investors.

This thesis investigates various portfolio construction techniques that are based on mathematics, statistics and econometrics. The goal is to bring retail investors back to the track of optimal behavior by providing them with systematic investment strategies that might help to avoid the psychological pitfalls associated with retail investing. Markowitz was one of the pioneers in the field of portfolio optimization. In the 1950s, he proposed the mean-variance framework, which is nowadays widely accepted and part of the Modern Portfolio Theory. Since then, the amount and complexity of portfolio construction techniques have increased substantially. This thesis provides the necessary knowledge to gain access to sophisticated risk-based portfolio construction techniques that can be implemented on a home computer. The covered risk measures are portfolio variance, semivariance, Conditional Value-at-Risk and Worst-Case Conditional Value-at-Risk. In order to investigate the performance of each technique, historical data comprising 21 years of daily index returns is used to conduct out-of-sample backtests for yearly, quarterly, monthly and weekly portfolio rebalancings.

The thesis is structured as follows. Section 2 outlines the role of retail investors and provides an overview of exchange-traded funds. Section 3 explains the asset selection criteria and investigates the selected data. Section 4 provides detailed information regarding various risk-based portfolio construction techniques, their potential estimation errors as well as their practical implementation. Section 5 explains the out-of-sample backtesting procedure and reports the empirical results. Finally, section 6 closes with conclusions.

2 Retail Investors

2.1 The Role of Retail Investors

2.1.1 Wall Street vs. Main Street

A retail investor (also referred to as private investor or individual investor) is a non-professional investor who trades securities through any kind of personal brokerage account. The retail investor generally trades in fairly small amounts compared to institutional investors like pension funds, banks or insurance companies. Because of the relatively small trading amounts, one might expect that the role of retail investors is negligible in the global markets. This tends to be true in these days but was not always the case. According to Evans (2009, p. 1105), retail investors owned over 90% of the stocks of U.S. corporations in 1950. Today, retail investors own less than 30% and represent a very small percentage of U.S. trading volume: less than 2% of the trading volume on the New York Stock Exchange can be attributed to retail investors. The financial markets are clearly dominated by institutional investors. However, even though the retail investors' proportion of ownership has declined over the past 70 years, the absolute dollar amount of retail investment has increased over time. This observation might be the result of the fact that personal finance and retail investment still remain highly relevant on a micro level because of individual goals such as maintaining the purchasing power, building financial security and preparing for retirement. Dealing with those aspects has become much easier for retail investors. Today, retail investors have unprecedented access to financial education, real-time information and financial markets worldwide. Developments in financial technology and the emergence of financial start-ups have challenged the traditional industry. The price competition among brokers has increased and many of today's online brokers do not even impose any minimum deposit or investment amounts, thereby democratizing the financial markets by lowering the existing costs and barriers. There has been a race to zero in the brokerage industry since 2013 when the start-up Robinhood implemented free trading on its platform. Today, this race is over and the competition has intensified with many electronic brokers such as TD Ameritrade, Schwab, Interactive Brokers and Firstrade offering commission-free trading for stocks, exchange traded funds or other securities – all to the immense benefit of retail investors (Fitzgerald, 2019). In addition, the electronic trading offered by brokers enables retail investors to manage their portfolios anywhere and anytime, which has levelled the playing fields between Wall Street and Main Street. Besides the latest developments in financial technology, the associated increase in flexibility and the reduction in trading costs, retail investors might even be in a better position than institutional investors due to several reasons (Mauzy, 2015): First, retail investors can act independently and focus their entire resources on their own portfolio without having to serve multiple clients at the same time. Second, retail investors can buy and sell securities with almost no market impact and there are no restrictions imposed on their portfolios by laws or clients. The only restriction might be that short-selling is more difficult to access compared to an institutional investor. Third, the alignment of incentives is optimal meaning that there are no managers trying to maximize their end of year bonus. Therefore, a retail investor can focus on maximizing his own profit over the long run. And last but not least, retail investors do not suffer from pressure to report short-term performance. There is no benchmarking against competitors, which might increase the pressure and therefore incentivize aggressive investing.

2.1.2 Psychological Hurdles of Retail Investment

"Investor behavior is not simply buying and selling at the wrong time, it is the psychological traps, triggers and misconceptions that cause investors to act irrationally. That irrationality leads to buying and selling at the wrong time, which leads to underperformance" (Dalbar, 2018, p. 7).

Albeit having many advantages over institutional investors, retail investors seem to perform poorly in a constant fashion. Extensive studies of investor behavior have shown that the track record of retail investment is not encouraging at all. Dalbar (2018), a leading financial services research firm, shows that the average retail investor constantly earns less than mutual fund performance reports would suggest. Their analysis covers data of more than 20 years up to December 31, 2017. Dalbar points out that retail equity investors are constantly underperforming the S&P 500 index over the time horizons of 20 years, 10 years, 5 years, 3 years and 12 months (p. 6). The constant underperformance of retail investors is also found in portfolios of fixed income investors. According to Dalbar, the main reason for the underperformance of retail investors is their own psychology (ibid.). Retail investors are highly exposed to behavioral biases, which lead to poor investment decision-making.

Examples for such biases are loss aversion, overconfidence, mental accounting, anchoring, herding, media response and many more. In line with Hirshleifer (2001), Montier (2007, p. 19) claims that most of these psychological pitfalls can be traced to four common causes: self-deception, heuristic simplification, emotion and social interaction. Being subject to these causes, retail investors tend to exhibit two behaviors that are most detrimental for investment performance: First, the tendency to move into and out of investments too frequently, which makes the investors stack up costs instead of profits. And second, the tendency to time the market (i.e. the attempt to predict the future). Dalbar's data indicates that the average mutual fund investor lacks the patience and stays invested in any fund for not more than four years and is hence unable to execute a long-term strategy (pp. 10-11). In addition to that, investors tend to sell their securities when markets decline or when there is imminent fear of a correction. Montier (2007, p. 97) claims that many investors try to perform on very short term horizons and overtrade as a consequence. The average holding period for a stock on the New York Stock Exchange is only 11 months, which has nothing to do with investment but rather constitutes speculation (p. 297). The International Organization of Securities Commissions [IOSCO] (2019, pp. 6 - 9) supports the argument that poor decision-making of retail investors can be attributed to psychological traits. They claim that retail investors use mental shortcuts (so-called heuristics) to simplify their decision-making since the amount of advice providers and investment products is overwhelming in light of the growing complexity of the global financial systems. According to the IOSCO, it is important to realize that retail investing is not merely about dollars, risk and return. For retail investors, high emotional stakes are involved. The choices they make may determine the quality of life when they retire or their ability to support family members. Such emotionally charged decisions are particularly stressful, which leads investors to rely even more heavily on heuristics to make a decision. This again leads to uninformed, poor decision-making, which can result in unsatisfactory investment performance. In order to deal with the psychological hurdles and the information overload, it might be helpful to implement a long-term investment strategy that does not involve any market timing or stock picking. Nowadays, such a passive investment approach can be put into practice by introducing exchange-traded funds into the portfolios of retail investors.

2.2 ETF-Investments in the Portfolios of Retail Investors

2.2.1 ETFs Briefly Explained

Exchange-traded funds (henceforth abbreviated as ETFs) are hybrid investment vehicles that track an index or a defined basket of assets (Pagano, Sánchez Serrano & Zechner, 2019, p. 4). They are continuously traded on liquid secondary markets and combine features of open-end as well as closed-end funds. ETFs belong to the most successful financial innovations of the early decades of the 2000s: over the past 20 years they have grown greatly in size, diversity, complexity and market significance. According to Fry (2008, pp. 13, 51), the ETF boom taking off in early 2004 has overwhelmed and challenged conventional business models in the financial sector by providing access to a wide array of low-cost products. ETF asset growth has been dramatic since the beginning of the second millenium: In 2001, 85 billion US-Dollars were invested globally in ETFs. By 2003, ETF assets totaled $150 billion and by the end of 2006, the assets reached $422 billion. As of the end of 2018, the aggregate total net assets of ETFs almost reached $5 trillion worldwide with $3.5 trillion located only in the USA (Pagano, Sánchez Serrano & Zechner, 2019, p. 5).

ETFs essentially act like funds but trade like stocks. The investor buys ETF shares on a stock exchange and receives a proportionate share of a pool of stocks, bonds and/or other asset classes (Ferri, 2008, p. 23). Each ETF is managed differently depending on the benchmark it follows. Similarly to open-end mutual funds, ETFs can offer broad diversification in just one product. The all-day trading, however, gives ETFs the leading edge over open-end mutual funds by increasing the investor's flexibility of adjusting his positions anytime the markets are open. To shed light on the differences between ETFs and open-end mutual funds, Ferri (pp. 25-30) outlines the major aspects: First of all, there is a significant difference in how investors acquire shares of each fund. For an open-end fund, shares are directly traded with the fund company who buys and sells the shares. ETF shares, on the other hand, are traded on a stock exchange and are already in the public domain. Besides the differences in trading, the two securities exhibit differences in pricing, too. Open-end mutual funds are priced once per day at their closing net asset value while ETF shares are continuously priced throughout the day whenever the stock market is open. Generally speaking, it is not possible for an investor to trade ETF shares directly with the fund company. However, there is one exception: The only players allowed to trade ETFs directly with the fund company are special institutional investors referred to as 'authorized participants' (APs). APs play a crucial role in the pricing efficiency of ETFs (pp. 35 - 37). The market price for ETF shares is essentially driven by forces of supply and demand. This is the reason why the ETF price can deviate from the prices of the underlying securities in the fund. Such price discrepancy is controlled for and stopped by the aforementioned APs acting on the primary market. The market price of an ETF is kept close to its net asset value by allowing APs to buy and redeem ETF shares from the fund company. Whenever a small price discrepancy between an ETF and its underlying securities occurs, APs step in and exploit the discrepancy by executing risk-free arbitrage. Those arbitrage trades allow APs to exchange individual securities for large blocks of ETF shares (so-called creation units) and vice versa. This mechanism aligns the market price of ETF shares with the fund's true value while the AP can reap a profit. BlackRock (2017, p. 4) claims that there are on average 34 authorized participants for each ETF. Besides the APs, secondary market participants attempt to benefit from the price discrepancies as well. The price pressure resulting from those arbitrage trades leads to the convergence of the prices. This process happens within split seconds and is effective in keeping ETF shares in line with their true value. The provision for arbitrage is what fundamentally makes ETFs differ from closed-end mutual funds.

As provided in the definition, ETFs track an index or a defined basket of assets. The benchmark replication mechanisms that specify how the tracking is implemented divide ETFs into two types: physical ETFs and synthetic ETFs (Meziani, 2016, p. 32-33). Physical ETFs attempt to replicate the benchmark index by holding all (i.e. full replication) or a representative sample (i.e. sampling) of the benchmark securities in their portfolios with weights that closely mimic those in the benchmark. Fully replicating ETFs hold all securities that are constituents in the benchmark. On the one hand, this replication mechanism can be costly considering the potentially large number of securities to be held. On the other hand, it provides clear transparency for the ETF buyer and ensures a precise tracking of the benchmark. However, there might be the case where some constituents of a benchmark are illiquid or subject to large trading costs. In those situations, sampling techniques are applied. This type of optimization filters out less liquid securities that could incur large trading costs. Obviously, this comes at the risk that the chosen sample fails to fully reflect the benchmark's performance resulting in a potentially significant difference between the ETF performance and that of the benchmark it is supposed to track, which is known as 'tracking error'. With the emergence of synthetic replication mechanisms, ETFs have somewhat exchanged their once celebrated feature of being transparent and straight-forward against an increased flexibility and access to various asset classes. Synthetic replication implies the use of derivatives (ibid.). In this case, the fund manager purchases an OTC traded derivative based on the underlying benchmark. This derivative takes the form of a total return swap: The counterparty, which is usually an investment bank, promises to pay the fund manager the exact performance including dividends of the selected benchmark. As for many derivatives, the fund company holds a basket of securities as a collateral for the swap contract. The crux lies in the fact that this basket can be different from the securities held in the benchmark. If the basket of securities achieves a higher return than the benchmark, the fund company pays the difference to the counterparty. In the other case, the fund company gets paid by the counterparty. It becomes clear that only one trade is required between the investment bank and the fund manager to have full exposure to the underlying index instead of several trades as in the case of full replication. This feature goes hand in hand with generally low costs and small tracking errors of synthetic ETFs. Moreover, this concept allows investors to gain exposure to various asset classes and sectors where effectively holding the security might be impossible or hardly practicable owing to liquidity issues. For example in the case of commodity ETFs, synthetic replication is often the only choice available. Due to regulatory differences, synthetic ETFs are more popular in Europe than in the United States. It is important that retail investors are aware of counterparty risks, which are inherent to both physical and synthetic ETFs. This is further discussed in 2.2.3.

2.2.2 Advantages of ETFs

The demand for ETFs has grown markedly in the recent years as both institutional and retail investors have found their numerous advantages appealing. Fry (2008, p. 33) claims that the logical choice for retail investors was and continues to be ETFs.

Intraday Tradability: As outlined previously, an ETF is essentially a mutual fund that trades like a stock on a secondary market. This feature provides all kinds of investors with liquidity and immediate access to various types of asset classes. The trade order flexibility of brokerage accounts allows ETF investors to benefit from making timely investment decisions and placing orders in a variety of ways (Ferri, 2008, p. 60 - 68). Moreover, the liquidity of ETFs leads to lower bid-ask spreads and therefore to lower trading costs.

Low Costs: Fund operating costs are incurred by any kind of managed fund. Costs have historically been very important in forecasting investors' returns. ETFs are able to reduce their expenses since monthly statements, notifications and information about transfers must only be provided to authorized participants and not to every shareholder as in the case of mutual funds. Moreover, the aforementioned feature of intraday tradability on a brokerage account leads to the fact that the brokerage firm is responsible for informing ETF shareholders about the investment's development and not the fund itself. Consequently, ETFs are able to reduce their costs even more. In addition to that, the aforementioned intense price competition among brokerage firms has led to very low commissions for trading ETFs, making them an even more desirable vehicle for investing.

Diversification: ETFs that track the performance of a specified index allow the investor to gain exposure to multiple securities in just one product. This reduces trading costs and increases investment diversification. Moreover, passive investment strategies are easily implemented with index ETFs. Antoniewicz and Heinrichs (2014, p. 4) claim that the rising popularity of passive investing has led to significant outflows in mutual funds, which have flown directly into ETFs. Those outflows can be explained by the results of many researchers who have studied the performance of actively managed mutual funds with the crushing result that only a relatively small number of active funds is able to achieve a persistent outperformance which would justify the management and performance fees of such funds. Besides their straight-forward implementation into passive investment strategies, ETFs have greatly simplified the access to specific markets or asset classes that would otherwise be difficult to reach. An example are foreign markets where investors are required to have local bank accounts or a foreign investor status in order to invest. Investors can overcome such investment obstacles simply by buying an ETF and therewith gaining exposure to global markets, commodities such as gold and energy as well as various dividend and value styles of investing – all in just a few products.

Besides the presented advantages, ETFs exhibit other desirable features such as the distribution of dividends and interest to shareholders, tax efficiency and many more. Those will not be further discussed. The general consent, however, is that the advantageous features of ETFs outweigh their disadvantages and risks.

2.2.3 Risks of ETFs

Tracking Error and Counterparty Risk: The tracking error is the difference between the ETF performance and that of the benchmark it is supposed to track. In passive investment strategies, the performance of an investor is measured by their ability to minimize the tracking error compared to the benchmark. Hence, tracking errors can constitute costs to investors. There are two reasons for the existence of tracking errors: First, ETFs hold cash positions while indexes do not. Therefore, a certain amount of tracking error has to be expected (Ferri, 2008, p. 70). Second, the timing of dividends is very difficult for the fund manager. Stocks go ex-dividend one day and pay out the dividend on some other day. Indexes, however, assume that the dividend is reinvested on the same day the company went ex-dividend. Because of these discrepancies, ETFs cannot perfectly track a targeted index, which leads to tracking errors. Due to the existence of counterparty risks, the selected replication method can lead to tracking errors as well. Physical ETFs are subject to counterparty risk because they can engage in security lending (lending out their securities against a fixed fee). By having this additional source of income, physical ETFs can mitigate the tracking error and lower their total expense ratio (TER). However, this feature is associated with an additional risk layer. Synthetic ETFs are exposed to the risk of a defaulting counterparty because they engage in derivative contracts. Of course, in both cases a collateral is required. Nevertheless, researchers and regulators raise concerns that the increasing counterparty risk in ETFs might cause systemic issues. According to Ben-David, Franzoni and Moussawi (2017), "past evidence shows that assets with a long chain of intermediaries and counterparties may cause or exacerbate financial shocks due to risk exposure along the chain of financial intermediaries" (p. 184). Moreover, Blocher and Whaley (2016, p. 24) claim that revenues from securities lending are substitutes for the fund's expense ratio, which indicates that ETFs can fine-tune the ratio of revenues between direct fees and securities lending. They find that the asset weights within an ETF can significantly diverge from those of the underlying benchmark in the sense that securities – which are more profitable to lend – are over-weighted. Those deviations in asset weights can in turn lead to larger tracking errors. The majority of today's ETFs do not disclose any lending-related activities and do not provide information about counterparties and collaterals, which is the reason why Blocher and Whaley (2016) call for more transparency.

Intraday Liquidity can backfire: Bhattacharya et al. (2017, pp. 1244 - 1245) show that retail investors who invest in ETFs perform worse than those who stick with traditional, less liquid mutual funds. They claim that the advantageous flexibility of ETF trading encourages retail investors to attempt to time the market. It is the retail investors' poor timing ability that is adversely affecting their portfolio returns. This finding is in line with the behavioral hurdles of retail investing outlined in section 2.1.2.

Systemic Risk: Due to their continuous trading, ETFs are able to replicate baskets of less liquid assets in the form of tradable ETF shares. Pagano, Sánchez Serrano and Zechner (2019, p. 7, 29) claim that this liquidity transformation can be the source of potential frictions. ETFs that are more liquid than their underlying securities are particularly attractive to short-term investors, who might take large short-term directional bets on entire asset classes. In case of amplified ETF price swings, leveraged investors with material ETF positions such as banks or hedge funds may become insolvent, which could lead to a chain reaction with systemic risk implications. Moreover, they claim that the co-movement between securities and their respective indices increases when they are included in ETF portfolios (p. 19). This, in turn, can increase the likelihood of investors experiencing capital losses simultaneously, which could lead to a disruption in the financial system with waves of insolvencies and synchronized fire sales. This risk is particularly pronounced when the ETFs' constituent securities are illiquid as well as during periods of financial market stress.

Risks during Bear Markets: ETF prices diverge from their underlying security prices due to forces of supply and demand. Pagano, Sánchez Serrano and Zechner (2019, p. 26) investigate ETF prices and claim that the arbitrage mechanisms, which try to avoid such decoupling, may operate imperfectly. Especially for illiquid underlying securities and during bear markets, the efficient price discovery of ETF shares might be troubled. During bear markets, the divergence of asset prices can result from authorized participants and other arbitrageurs not having the incentives and/or capacity to realign the ETF prices with those of the underlying securities.

APs may therefore not be willing to engage in arbitrage, especially when the underlying securities are illiquid. Marshall, Nguyen and Visaltanachoti (2017, p. 5) claim that the liquidity of an ETF is positively correlated with the liquidity of its underlying securities: A larger liquidity of the underlying securities implies a greater ability for the arbitrageurs to engage in arbitrage trades, which leads to a larger liquidity for the ETF. Petajisto (2017, p. 24) extends this view and finds that the illiquidity of the underlying assets amplifies the ETF price discrepancy, which is more difficult to resolve due to the illiquidity. The behavior of arbitrageurs is further investigated by Pan and Zeng (2019, pp. 26 - 29) who present evidence that the trading volume of authorized participants who deal with illiquid constituents significantly declines when market volatility is high. Ben-David, Franzoni and Moussawi (2017, pp. 183 - 185) point out that during several episodes in recent years, ETFs have experienced illiquidity issues during times of market turbulence. Examples are the Flash Crash of May 6, 2010 as well as the market turbulence of June 20, 2013. Various market participants have exited the market once signs of extreme volatility and illiquidity appeared. Price discovery of ETFs no longer took place and the decoupling of ETF share prices with respect to their underlying securities became reality. Hence, the efficient price discovery of ETF shares crucially depends on the presence of agents who facilitate arbitrage. Moreover, ETFs might provide a false sense of liquidity: They are very liquid in a regular trading environment. As soon as the financial markets experience hefty turbulence, however, they tend to experience a dry-up in liquidity because authorized participants and other arbitrageurs abstain from engaging in arbitrage activity. Besides potential liquidity issues, ETFs' diversification benefits might be scrutinized during bear markets as well. Neves, Fernandes and Martins (2019, pp. 154 - 156) find evidence of contagion effects between index ETFs. They claim that the benefits of diversification through ETFs are limited, particularly in times of a financial crisis. In line with many researchers, their results indicate a high correlation between markets during turbulent times, which leads to a decrease in the diversification benefits that ETFs provide. They conclude that the 22 stock ETFs they considered are not a great vehicle for diversifying an investor's portfolio despite representing the majority of global indexes. One must remark, however, that their results are based solely on stock indexes and that they do not harness the diversification across asset classes.

3 Data

3.1 Asset Selection

In order to investigate the performance of different portfolio optimization techniques, real-world financial data is required. Owing to the above-mentioned advantages and features of ETFs, an ETF portfolio is constructed and analyzed for an international retail investor. Because of the recent growth of ETFs, historical data is rather limited and often does not exceed 10 - 15 years of track record. In order to extend this time horizon, index data of the respective ETFs is used as a proxy for the historical ETF performance. The data on daily prices of the selected indexes was obtained from Bloomberg. Because the perspective of an international investor is taken, all data is based on US dollars. The interest rate on three-month Treasury bills is used as a proxy for risk-free rates. The respective data was gathered from the database of the Federal Reserve Bank of St. Louis, where daily data is available. The ETFs and the corresponding indexes were selected according to multiple criteria outlined below, which should ensure a realistic and practicable retail investor portfolio. Table 1 displays the resulting asset selection.

(i) *Broad Diversification:* The investor is assumed to seek exposure in various asset classes including stocks, bonds, real estate and commodities. Moreover, the investor seeks global exposure including both developed and emerging markets.

(ii) *Denomination in USD*: As already mentioned, the investor is assumed to be globally oriented. Therefore, the investments are based on USD.

(iii) *Practicability:* Index data is only considered if an investible ETF exists, which tracks the respective index performance.

(iv) *Track Record:* In order to obtain a representative set of historical observations, index data is used as proxy for the rather limited ETF data. Moreover, indexes with a longer history are preferred.

(v) *Low Costs:* As mentioned in subsection 2.2.2, costs can be a good predictor of future returns. Hence, the investor is assumed to prefer ETFs with a low TER. Because the portfolio will be rebalanced periodically, it is crucial to choose ETFs with high volumes and low bid-ask-spreads, which are hidden costs of a retail investor's ETF portfolio.

	US Equity	World Equity	EM Equity	World Bond	US Real Estate	Commodity
Index Name	Russell 2000 TR Index	FTSE All-World ex US Index	MSCI EM IMI USD Net Index	BloomBarc Global-Agg TR Index USD	Dow Jones US Select REIT Index	Bloomberg Commodity Index
Index Ticker	RU20INTR Index	FTAW02 Index	MIMUEMRN Index	LEGATRUU Index	DWRTFT Index	BCOMTR Index
ETF Name	iShares Russell 2000 ETF	Vanguard FTSE All-World ex-US ETF	iShares Core MSCI EM ETF	iShares Global Agg Bond ETF	Schwab U.S. REIT ETF	iShares Diversified Commodity Swap ETF
ETF Ticker	IWM	VEU	IEMG	AGGG	SCHH	ICOM
Inception Date	05/26/2000	03/08/2007	10/22/2012	11/23/2017	01/13/2011	07/20/2017
Strategic Focus	Small-Cap	Large-Cap	Large-Cap	IG* or higher	Broad Market	Broad Market
Replication Method	Physical	Physical	Physical	Physical	Physical	Synthetic
Replication Strategy	Full	Full	Sampling	Sampling	Full	Swap-based
Total Expense Ratio	0.19%	0.09%	0.14%	0.10%	0.07%	0.19%
Total Assets ($ Mio.)	48'070	26'090	61'970	3'230	6'040	1'510
Securities Lending	Yes	Yes	Yes	Yes	Yes	Yes

Table 1: Asset Selection Overview

* IG: Investment Grade

Source: Bloomberg, as of December 31, 2019.

3.2 Data Description and Stylized Facts

After preparing the data sample for the subsequent analyses, it consists of daily simple returns of the indexes described in Table 1. The raw data is based on adjusted closing prices of the indexes. The sample period comprises 21 years of financial data: It begins in January 5, 1999 and ends in December 31, 2019. The total number of simple returns for the sample period is 5246. Table 2 provides a first overview of the time series at hand. The total number of simple returns is equal among the indexes since any trading day, for which there are missing values in either of the indexes, is removed from the sample. The reason for such missing values is the difference in trading days between the various indexes.

Index	N	Min (%)	Max (%)	Mean	Stdev (%)	Skewness	Kurtosis
US Equity	5246	-11.85	9.26	0	1.46	-0.15	4.75
World Equity	5246	-7.98	7.88	0	1.05	-0.23	7.79
EM Equity	5246	-9.30	10.19	0	1.13	-0.38	7.80
World Bond	5246	-1.94	2.86	0	0.33	0.11	3.11
US Real Estate	5246	-19.76	18.98	0	1.77	0.41	22.55
Commodity	5246	-6.20	5.81	0	1.00	-0.16	2.62

Table 2: Basic Statistics of the Index Return Data

According to Pfaff (2013, pp. 26 - 33), many empirical studies have identified common features among financial market data, which are known as stylized facts. Essentially, the stylized facts are statistical properties that appear to be present in various financial market data sets. Table 2 in combination with supporting analyses in the appendix show that the data at hand exhibits some of these stylized facts.

In line with Tsay (2010, pp. 8 - 10, 21 - 22), the mean of each index return series is basically zero. Figure A.2 in the appendix depicts the return series of all indexes and allows for a qualitative assessment of this stylized fact. The daily mean returns are very close to zero due to the frequency of the data: The general consent is that aggregating the financial data towards lower frequencies (e.g. monthly, quarterly, yearly) results in higher means.

The daily standard deviations lie between 0.33% and 1.77%. The World Bond Index exhibits the lowest while the US Real Estate Index exhibits the highest volatility in the data sample. Further investigation of the return series' volatilities allows a deeper understanding of the data. Figure A.3 portrays the squared return series plotted against the time index. Squared (as well as absolute) returns serve as a proxy for the returns' volatilities. It is a very typical manifestation in financial time series that extreme (positive or negative) returns are observed closely in time, which is known as volatility clustering. Moreover, A.3 shows that the volatilities of the return processes are not constant with respect to time.

The next stylized facts relate to the higher moments of a return distribution, for which the relevant formulas are stated in the appendix. The skewness is the normalized third central moment of a distribution. In other words, the skewness provides information about the asymmetry of the return distribution. Table 2 shows that negative returns are more likely to occur than positive returns for the majority of the indexes. The only exceptions are the World Bond Index and the US Real Estate Index, whose empirical return distributions are positively skewed and therefore not in line with the stylized fact of negative skewness. The kurtosis is the normalized fourth central moment of a distribution and captures its tail thickness. A large kurtosis indicates that extreme realizations (positive and negative) are more likely to occur. A distribution's tail thickness is often related to the tail thickness of a normal distribution by reporting the so-called excess kurtosis (kurtosis minus 3). A positive excess kurtosis indicates the presence of leptokurtosis, which is also known as 'fat tails' and implies that extreme realizations are more likely to occur compared to realizations drawn from a normal distribution. Table 2 shows that all index return series exhibit fat tails except for the Commodity Index. By far the largest excess kurtosis is detected for the US Real Estate Index with a value of 19.55, indicating very fat tails. This finding is reflected when analyzing the extremes in daily returns: the US Real Estate Index has a maximal and minimal daily return of almost ±20% in the sample period. Figures A.2 as well as A.3 show that the daily returns and volatilities of the US Real Estate Index are very extreme during the global financial crisis. This makes intuitively sense given the fact that the US housing bubble has led to a direct involvement of the real estate market during said crisis.

The presence of skewness as well as the excess kurtosis of almost all index return series might point to the next stylized fact. Namely that financial return series – in harsh contrast to the simplifying assumption often used in the industry – are not normally distributed. To further test this claim, QQ-plots are examined in Figure A.4. The empirical quantiles are plotted against those of a normal distribution with estimated mean and variance. The QQ-plots show that the non-normality of the index return series stems from the upper and lower tail of the empirical return distributions. In addition, Jarque-Bera tests are performed for each index return series in order to support the findings quantitatively. The results are reported in Table A.1 and they strongly indicate that the return series of all indexes are non-normal, while the World Bond Indexes' returns are the closest to a normal distribution.

In order to analyze the linear dependence between a return series and its past values, the autocorrelation is investigated. Since this information is not included in Table 2, ACF plots are provided in Figures A.5 and A.6. It is a stylized fact that the absolute or squared returns – which serve as a proxy for volatility – are highly autocorrelated for several lags while non-transformed returns are not. The strong time dependence in volatility can also be found in the data set at hand and is an important feature in financial market modeling, which has led to the emergence of a huge class of conditional heteroskedasticity models. The phenomenon of significant autocorrelations of squared (or absolute) returns for long horizons is referred to as the long memory of volatility. The long information propagation within the data is in line with Cont (2005, p. 3), who claims that the long memory of volatility is stable across various asset classes and time periods.

Because the selected index return series will be analyzed in a portfolio context, cross-correlations between the assets must be considered as well. In the case of N assets, there are $N*(N-1)/2$ unique asset pairs to be examined. The cross-correlations of each asset pair are investigated and the result is similar to the univariate case: the cross-correlations between two return series are less pronounced and not significant compared to those of the squared return series. Additionally, the asset pairs' rolling correlations are analyzed. Figure A.8 and Table A.2 provide an overview of the linear dependencies within the data set, making it obvious that the rolling correlations are clearly not constant over time.

4 Portfolio Construction

4.1 Preliminary Remarks and Mathematical Definitions

In this chapter, various portfolio construction techniques are explained and investigated. The goal is to provide retail investors with purely data-driven strategies for risk-based portfolio construction, which can be regarded as tools to avoid the behavioral pitfalls of retail investing. In short, the retail investor is encouraged to let the numbers do the work by employing a systematic investment approach based on mathematical and statistical concepts. Besides providing an overview of existing portfolio construction techniques and their implementation approaches, the subsequent sections also discuss estimation errors and their mitigation attempts. To ease the interpretation of the various portfolio construction techniques, some mathematical definitions and notations are required:

Let $(\Omega, \mathscr{F}, \mathbb{P})$ be a given probability space representing financial market outcomes. In total, there are N risky assets with simple return series $R_1, ..., R_N$, which are all random variables on the probability space. Let $\mathbf{R} = (R_1, ..., R_N)$ be an N-dimensional vector time series of the assets' simple returns of length $T : \{r_1, ..., r_T\}$. The boldfaced index $\mathbf{1}$ denotes an N-dimensional vector of ones. In the following subsections, all portfolio optimizations are subject to the same constraints: It is assumed that the retail investor aims to be fully invested. This means that the sum of all asset weights should sum to one. Moreover, due to the difficulty for retail investors to access short sales, all portfolio weights are assumed to be long only. To summarize, the retail investor's portfolio is described by the weight vector

$$\mathbf{w} = (w_1, ..., w_N)\ , \ \text{ such that } \ \mathbf{w} \in \mathbf{W} = \left\{ \mathbf{w} \in \mathbb{R}^N \mid \sum_{i=1}^{N} w_i = 1, \ \ w_i \geq 0, \ \ \forall\, i = 1, ..., N \right\},$$

where $\mathbf{W}$ denotes the set of all shortsale-constrained and fully invested portfolios. The weight vector $\mathbf{w}$ represents the allocation of a retail investor's wealth across the various assets in the sense that w_i describes the fraction of wealth invested in the i-th asset. For any $\mathbf{w} \in \mathbf{W}$, the portfolio return is defined by $R_{\mathbf{w}} = \sum_{i=1}^{N} w_i R_i = \mathbf{w}'\mathbf{R}$. Let $\mu_i = \mathbb{E}_{\mathbb{P}}[R_i]$ denote the expected return of asset i for $i = 1, ..., N$. Moreover, for two risky assets' returns R_i and R_j, $V_{ij} = Cov(R_i, R_j)$ denotes the covariance between the two assets for $i \neq j$. For $i = j$ we have $V_{ij} = Cov(R_i, R_j) = Var(R_i) = \sigma_i^2$, which is the variance of asset i. The N-dimensional vector $\boldsymbol{\mu}$ denotes the expected returns of the risky assets while $\boldsymbol{\Sigma} = (V_{ij})_{1\leq i,j\leq N}$ denotes the asset returns' $N \times N$ variance-covariance matrix. For a portfolio $\mathbf{w} \in \mathbf{W}$, the first two moments of the portfolio return distribution are defined as:

$$m_1 = \mu_P = \mathbb{E}_{\mathbb{P}}[R_{\mathbf{w}}] = \sum_{i=1}^{N} w_i \mathbb{E}_{\mathbb{P}}[R_i] = \mathbf{w}'\boldsymbol{\mu}, \text{ and}$$

$$m_2 = \sigma_P^2 = Var(R_{\mathbf{w}}) = \sum_{i=1}^{N}\sum_{j=1}^{N} w_i w_j Cov(R_i, R_j) = \mathbf{w}'\boldsymbol{\Sigma}\mathbf{w}.$$

In order to provide definitions for the third and fourth moment of the portfolio return distribution, it is computationally convenient to express the portfolio moments as a function of the multivariate moments of the underlying assets' returns using the $N \times N^2$ co-skewness matrix $\boldsymbol{M_3}$ and the $N \times N^3$ co-kurtosis matrix $\boldsymbol{M_4}$:

$$\boldsymbol{M_3} = \mathbb{E}_{\mathbb{P}}[(\mathbf{R}-\boldsymbol{\mu})'(\mathbf{R}-\boldsymbol{\mu}) \otimes (\mathbf{R}-\boldsymbol{\mu})] \qquad \boldsymbol{M_4} = \mathbb{E}_{\mathbb{P}}[(\mathbf{R}-\boldsymbol{\mu})'(\mathbf{R}-\boldsymbol{\mu}) \otimes (\mathbf{R}-\boldsymbol{\mu}) \otimes (\mathbf{R}-\boldsymbol{\mu})],$$

where $\mathbb{E}_{\mathbb{P}}$ is the expectation operator under probability measure $\mathbb{P}$ and $\otimes$ is the Kronecker product (Boudt, Peterson & Croux, 2008, pp. 82-83). With this representation, it is straightforward to compute the third and fourth portfolio moments as:

$$m_3 = \mathbf{w}'\boldsymbol{M_3}(\mathbf{w} \otimes \mathbf{w}) \qquad m_4 = \mathbf{w}'\boldsymbol{M_4}(\mathbf{w} \otimes \mathbf{w} \otimes \mathbf{w})$$

Consequently, the portfolio skewness s_p and the portfolio excess kurtosis k_p are given by:

$$s_p = m_3/m_2^{3/2} \quad k_p = m_4/m_2^2 - 3$$

Note that the time subscripts are omitted in order to simplify the notation. In the following subsections, these definitions will be used to mathematically describe the portfolio construction techniques and their solution approaches.

4.2 Benchmark: Equal-Weight Portfolio

"Given the inherent simplicity and the relatively low cost of implementing the 1/N naive diversification rule, such a strategy should serve as a natural benchmark to assess the performance of more sophisticated asset-allocation rules" *(DeMiguel, Garlappi, & Uppal, 2007, p. 1921).*

In line with DeMiguel, Garlappi and Uppal, the equal-weight portfolio (also referred to as 1/N-portfolio) is used as a benchmark for the subsequent portfolio construction and optimization techniques. The authors claim that there are two main reasons for using this "naive" rule as a benchmark (pp. 1916 - 1920): First, the equal-weight portfolio is easily implemented because it does not rely on any estimation of the asset returns' moments. Therefore it is also straightforward to apply for a large number of assets. Second, the majority of investors continue to use such simple allocation rules despite the advanced theoretical models developed in the last decades. Moreover, the equal-weight portfolio seems to perform very well because it is not subject to estimation errors. DeMiguel, Garlappi and Uppal investigate 14 different models across seven empirical data sets and show that none of the sophisticated models is consistently better than the 1/N rule in terms of Sharpe ratio, certainty-equivalent return or turnover. This indicates that the out-of-sample gain from portfolio optimization is more than offset by estimation errors. However, these findings strongly depend on the amount of assets and the data history considered, with small samples of historical data leading to larger estimation errors. Formally, the equal-weight portfolio is defined as

$$\mathbf{w} = (w_1, \ldots, w_N)\ , \quad \text{where} \quad w_i = 1/N \text{ for } i = 1, \ldots, N.$$

Based on the equal-weight portfolio, two benchmarks will be used for the subsequent analyses: First, the equal-weight portfolio with frequent rebalancing, where the fraction $1/N$ of wealth is allocated to each of the N assets at each rebalancing date. Second, the equal-weight portfolio without rebalancing, where the equal weights are constructed only in the beginning of the sample period (simple buy-and-hold strategy). In the financial literature, the general consent is that frequent rebalancing increases returns and reduces risks. But frequent rebalancing requires frequent trading, which goes hand in hand with larger trading costs. By using two different benchmarks, rebalancing effects as well as the impact of trading costs on a retail investor's portfolio can be examined.

4.3 Mean-Variance: Maximum Sharpe Ratio Portfolio

4.3.1 Model Description

The mean-variance portfolio construction was pioneered by Harry Markowitz' seminal paper from the 1950s. He was awarded a Nobel prize in Economic Sciences in 1990 for his revolutionary work leading to the development of the Modern Portfolio Theory. In this kind of portfolio construction, the investor optimizes the tradeoff between the mean and the variance of portfolio returns. The attempt is to weight the assets in such a way that the expected return is maximized while the overall risk is being minimized. The underlying idea of this method is the comparison of all the portfolios that could potentially be built given a selection of assets (Chicarini & Kim, 2006, pp. 256 - 265). Theoretically, the portfolio risk as well as the portfolio's expected return can be computed for any portfolio using the underlying assets' empirical moments. Once this is completed, the investor can choose the portfolio with the lowest risk for a given level of expected return. In reality, this calculation would take forever since there is an infinite number of possible portfolio weight combinations. To overcome this hurdle, quadratic programming is employed when finding the optimal portfolio.

The mean-variance portfolio optimization is based on the first two moments of the assets' return distribution and optimizes the risk-return-tradeoff by defining the portfolio variance as risk measure. To see this, consider the mathematical notations and definitions outlined in 4.1 in combination with the following optimization problem that minimizes the portfolio variance with a target level of portfolio return denoted by $\overline{\mu}$:

$$\underset{\mathbf{w}\in\mathbb{R}^N}{\text{argmin}} \;\; \mathbf{w}'\boldsymbol{\Sigma}\mathbf{w} \;\; \text{subject to} \;\; \mu_P \equiv \mathbf{w}'\boldsymbol{\mu} = \overline{\mu}$$

In order to implement this model, the classic 'plug-in' approach can be applied, where the asset returns' means and their covariance matrix are replaced by the sample counterparts $\hat{\boldsymbol{\mu}}$ and $\hat{\boldsymbol{\Sigma}}$. This mean-variance optimization is a quadratic optimization problem and has a closed-form solution if the portfolio weights $\mathbf{w}$ are unconstrained. This allows for a simple implementation, which has led to the model's popularity. Throughout this thesis, however, the investor is subject to shortsale constraints, which makes this a quadratic optimization problem with inequality constraints. Such problems do not have a closed-form solution and can therefore only be solved numerically. Chincarini and Kim (pp. 289 - 295) provide an overview of the numerical methods applicable to deal with such an optimization problem. By solving above optimization problem, an investor can construct a portfolio of risky assets that will minimize the risk (defined by portfolio variance) for a given level of portfolio return. Because any target portfolio return can be chosen, there is a set of optimal portfolios. Repeating the optimization for all possible target levels provides the investor with the so-called efficient frontier. Likewise, the investor can construct a portfolio that maximizes the portfolio return for a given target level of risk. Because of the vast amount of possibilities in this Mean-Variance framework, this thesis focuses only on a very specific portfolio located on the efficient frontier: the tangency portfolio. The tangency portfolio achieves the highest possible Sharpe Ratio, which is the reason why it is also referred to as Maximum Sharpe Ratio portfolio. This portfolio has the most favourable reward-risk-ratio and is achieved by solving the following constrained optimization problem:

$$\underset{\mathbf{w}\in\mathbf{W}}{\operatorname{argmax}} \; \frac{\mathbf{w}'\boldsymbol{\mu} - r_f}{\mathbf{w}'\boldsymbol{\Sigma}\mathbf{w}} = \underset{\mathbf{w}\in\mathbf{W}}{\operatorname{argmax}} \; \frac{\mu_P - r_f}{\sigma_P^2}$$

where r_f denotes the risk-free rate, μ_P denotes the expected portfolio return and σ_P^2 denotes the portfolio variance. The risk-free rate r_f is set to zero for the Maximum Sharpe Ratio Portfolio in order to ensure the comparability between the various portfolio construction techniques. Again, due to the shortsale constraints, the solution to this optimization problem can only be obtained numerically.

4.3.2 Model Drawbacks

Despite its popularity, mean-variance optimization is subject to several disadvantages. First, the mean-variance optimization is limited only to the first two moments of the return distribution, meaning that higher moments such as skweness and kurtosis are entirely neglected. Therefore, low probability events are not appropriately taken into account when employing this portfolio optimization technique. Second, the portfolio variance is used as risk measure, which is a symmetric measure that penalizes gains and losses in the same way. In other words, there is no distinction between gains and losses because the fluctuations above and below the mean are considered and minimized. Third and most important from a practical perspective, the mean-variance framework is very sensitive to its input parameters. Chincarini and Kim (2006) mention that the mean-variance optimization acts as a chaotic investment decision system because small changes in input parameters imply enormous changes in the optimized portfolio weights (p. 257). Mean-variance optimization overuses statistically estimated information and magnifies the impact of estimation errors. Because sample counterparts are usually used to estimate the means and covariances (plug-in approach), this optimization technique completely ignores the possibility of estimation error of the asset returns' first two moments.

It is well known that sample mean estimates have large variability due to their sensitivity to outliers in historical returns. The mean estimates have a profound impact on the sensitivity of the model and are therefore seen as the Achilles heel of mean-variance optimizations (Martin, Clark & Green, 2010, p. 337). According to the authors, it moreover belongs to the conventional wisdom that historical estimates of the covariance matrix are more precise than historical sample mean estimates when the number of return observations is substantially larger than the number of assets in the portfolio. DeMiguel, Garlappi and Uppal (2007), however, claim that a shortsale-constrained portfolio can reduce estimation errors of a sample-based mean-variance portfolio since constraining the portfolio weights is equivalent to "shrinking" the expected return estimates towards their true average (p. 1925). Consequently, the robustness issues regarding mean estimations might be mitigated in the context of this thesis due to the shortsale constraint imposed on the retail investor's portfolio.

4.4 Minimum Variance Portfolio

Another special case of the mean-variance portfolio optimization is the minimum variance portfolio, which is also located on the efficient frontier – similar to the maximum Sharpe Ratio portfolio. The respective optimal weights can be obtained by solving:

$$\underset{\mathbf{w}\in\mathbf{W}}{\operatorname{argmin}} \ \mathbf{w}'\boldsymbol{\Sigma}\mathbf{w}$$

While the minimum variance portfolio optimization is also subject to the majority of the previously mentioned drawbacks of the mean-variance optimization, it has one crucial advantage: It belongs to the portfolios with moment restrictions and it only utilizes the estimate of the asset returns' covariance matrix while completely ignoring the estimations of the expected returns. To be more precise, DeMiguel, Garlappi and Uppal (p. 1924) show that expected returns still appear in the likelihood function needed to estimate the covariance matrix $\boldsymbol{\Sigma}$. However, they mention that under the assumption of normally distributed asset returns, the maximum likelihood estimator of the mean is always the sample mean for any estimator of the covariance matrix. This allows one to remove the dependence on expected returns for constructing the maximum likelihood estimator of $\boldsymbol{\Sigma}$. It is important to realize that the data at hand was thoroughly checked in chapter 3.2 and that a normality assumption is not valid in the present application. The dependence on mean estimates when estimating the covariance matrix in combination with non-normal return data could potentially be the reason if the minimum variance portfolio turns out to produce very sensitive results. However, similar to the mean-variance optimization, imposing a shortsale constraint on the minimum variance portfolio is equivalent to shrinking the elements of the covariance matrix, which leads to a mitigation of the estimation error. Jagannathan and Ma (2003, p. 1654) find that the sample covariance matrix with an imposed shortsale constraint performs almost as good as those specifically constructed to take into account estimation robustness (e.g. using factor models or shrinkage estimators). To conclude, it is expected that the optimal weights resulting from the constrained minimum variance optimization are less sensitive regarding the model inputs compared to the maximum Sharpe Ratio portfolio because the Achilles heel of the mean-variance optimization (i.e. the mean estimation) is not directly involved in the objective function.

4.5 Minimum Semivariance Portfolio

This portfolio construction technique attempts to solve the mean-variance optimization's drawback of using the portfolio variance as risk measure. By defining the lower partial second moment as risk measure, only the downside risk is taken into account when constructing the minimum semivariance portfolio. Despite the mean-variance optimization's popularity, it may be less known that Markowitz himself favoured the semivariance as measure of risk from the very beginning of his work (Estrada, 2008, pp. 57 - 58). The reason for this is that an investor usually worries about underperformance rather than overperformance. By minimizing the downside risk, this can be taken into account, making the semivariance a more plausible measure of risk compared to the variance. Back in the 1950s, variance had an edge over semivariance in terms of cost and convenience, which was the reason why Markowitz continued his work based on variance. The main issue with semivariance lies in the fact that the semicovariance matrix is endogeneous to the portfolio weights: a change in portfolio weights affects the periods in which the portfolio underperforms the benchmark, which in turn affects the elements of the semicovariance matrix. To see this, the respective formulas are investigated. Markowitz suggested estimating the semivariance of a portfolio with the expression

$$\Sigma_{PB}^2 = \sum_{i=1}^{N}\sum_{j=1}^{N} w_i w_j S_{ijB}, \quad \text{with} \quad S_{ijB} = 1/T * \sum_{t=1}^{K}(R_{i,t} - B)(R_{j,t} - B),$$

where B is any benchmark return chosen by the investor (also known as Minimum Acceptable Return, MAR). Throughout this thesis, the reference point B is chosen to be zero. Hence, only negative deviations from zero are taken into account. The periods 1 through K are those in which the portfolio underperforms the benchmark return B. Those periods crucially depend on the portfolio weights, which is the reason why Markowitz' definition of portfolio semivariance suffers from endogeneity. Nowadays, the issue of endogeneity can be solved using non-parametric methods as well as sophisticated algorithms that try to ensure convergence to the solution. Athayde (2005) proposes using non-parametric Kernel estimations of the returns while Liagkouras and Metaxiotis (2013) propose a multi-objective evolutionary algorithm to solve the constrained optimization problem. According to Estrada (p. 58), complicated black-box solutions are the reason why practitioners refrain from using the mean-semivariance framework. Instead of using 'obscure' numerical algorithms, Estrada proposes a heuristic approach to the estimation of the semicovariance matrix by defining it as

$$\mathbf{\Sigma}_{ijB} = E\left[Min(R_i - B, 0) * Min(R_j - B, 0)\right] = 1/T * \sum_{t=1}^{T}\left[Min(R_{i,t} - B, 0) * Min(R_{j,t} - B, 0)\right],$$

which produces a symmetric and exogenous semicovariance matrix. The endogeneity problem is avoided because the relevant knowledge is now whether the respective asset – and not the entire portfolio – has underperformed the benchmark B. As a consequence, the elements in the semicovariance matrix are invariant to the portfolio weights considered, leading to exogeneity. The simplicity of Estrada's approach lies in the fact that the portfolio semivariance is approximated by

$$\Sigma_{PB}^2 \approx \sum_{i=1}^{N}\sum_{j=1}^{N} w_i w_j \mathbf{\Sigma}_{ijB}.$$

Replacing the covariance matrix by the symmetric and exogenous semicovariance matrix $\mathbf{\Sigma}_{ijB}$ leads to the fact that the well-defined and well-known solutions of mean-variance problems can be applied to mean-semivariance problems. However, using above approximation and not relying on black-box numerical algorithms might hinder estimation accuracy. Estrada (p. 63) claims that there is an error due to the approximation, whose direction is predictable. Whenever there is an estimation error in the portfolio semivariance, Estrada shows that the approximation is always larger than its exact counterpart. A potential discrepancy always leads to a more conservative asset allocation by overestimating the downside risk of the portfolio. The extent of the overestimation is fairly small. Estrada tests over 1'100 portfolios and finds that the portfolio semivariances (exact and approximated) are highly correlated and very close in value. Moreover, he finds that this heuristic approach is particularly accurate in portfolio optimization contexts, where funds are allocated across various asset classes, which is exactly the scope of the present thesis. To conclude, the minimum semivariance portfolio is obtained by solving:

$$\underset{\mathbf{w} \in \mathbf{W}}{\operatorname{argmin}} \; \mathbf{w}'\mathbf{\Sigma}_{ijB}\mathbf{w}$$

4.6 Minimum Conditional Value-at-Risk Portfolio

4.6.1 Model Description

The minimum Conditional Value-at-Risk portfolio optimization uses a quantile-based risk measure and allows the investor to minimize the expected tail loss of his or her portfolio for a given confidence level. In order to understand the underlying risk measures, some mathematical explanations are required. Rockafellar and Uryasev (2000) introduced a very convenient approach to minimize the Conditional Value-at-Risk using linear programming, which overcomes mathematical as well as computational issues. The subsequent derivations are based on their contribution. Suppose that $f(\mathbf{w},\mathbf{R})$ is the portfolio loss function, which is essentially obtained by multiplying the portfolio returns with minus one. Assume that the random vector $\mathbf{R}$ has a probability density function of $p(\mathbf{R})$. Then, the cumulative distribution function of the loss associated with the weight vector $\mathbf{w}$ is defined as

$$P[f(\mathbf{w},\mathbf{R}) \leq \gamma] \equiv \Psi(\mathbf{w},\gamma) = \int_{f(\mathbf{w},\mathbf{R})\leq\gamma} p(\mathbf{R})d\mathbf{R},$$

which yields the probability of $f(\mathbf{w},\mathbf{R})$ not exceeding the threshold γ. The cumulative distribution function $\Psi(\mathbf{w},\gamma)$ is nondecreasing with respect to γ and continuous from the right, but not necessarily from the left due to the possibility of jumps. For the subsequent analyses, it is assumed that the cumulative distribution function is everywhere continuous with respect to γ. The definitions of the risk measures VaR and CVaR in case of discontinuities are provided in the appendix on p. 78. For a given confidence level $\alpha \in (0,1)$, the Value-at-Risk associated with a portfolio $\mathbf{w} \in \mathbf{W}$ is defined by

$$VaR_\alpha(\mathbf{w}) = \min\{\gamma \in \mathbb{R} : \Psi(\mathbf{w},\gamma) \geq \alpha\}.$$

The $VaR_\alpha(\mathbf{w})$ corresponds to the α-quantile of the portfolio loss distribution and has an intuitive interpretation: In (100*α)% of the cases, a negative portfolio return does not exceed this value (in absolute terms, because the VaR is reported as a positive number). In order to remain general and to ensure applicability to a portfolio of any size, the portfolio losses are considered in percentages rather than monetary values. The confidence level is set to $\alpha = 0.95$ throughout this thesis. It is important to realize that the Value-at-Risk is backward-looking because it is based on the portfolio return distribution over a given historical observation period. In fact, the Value-at-Risk suffers from many shortcomings.

First, the VaR is a purely statistical quantity without any economic principle because it does not take into account how large a potential exceedance of the VaR is. In other words, it simply tells the investor that a percentage of past portfolio returns exceeded the VaR, but the size of the exceedance is not reflected. Second, it is not a convex risk measure. Especially in a portfolio optimization context, convexity is a key property in order to efficiently find a global minimum because it eliminates the possibility of a local minimum being different from a global minimum (Rockafellar & Uryasev, 2000, p. 24). And third, the VaR fails to satisfy the sub-additivity property and therefore does not reflect diversification effects. Owing to this, the Value-at-Risk is not a coherent risk measure. Coherency can be thought of as a set of mathematical properties an investor expects to hold for a risk measure (Artzner et al., 1999). The necessary properties are *positive homogeneity*, *translation equivariance*, *monotonicity* and the one that the VaR lacks: *sub-additivity*. All definitions are provided and explained in the appendix on p. 77. Coherent risk measures are very appealing when it comes to portfolio optimization because they are convex by nature. The Conditional Value-at-Risk satisfies all the necessary properties and is therefore a coherent risk measure. Given a fixed $\mathbf{w} \in \mathbf{w}$ and a confidence level α, the CVaR is defined as:

$$CVaR_\alpha(\mathbf{w}) = \frac{1}{1-\alpha}\int_{f(\mathbf{w},\mathbf{R})\geq VaR_\alpha(\mathbf{w})} f(\mathbf{w},\mathbf{R})p(\mathbf{R})d\mathbf{R},$$

which is the expectation of the loss function associated with $\mathbf{w}$ conditional on that loss being equal to the $VaR_\alpha(\mathbf{w})$ or greater. In other words, $CVaR_\alpha$ can be interpreted as the equally weighted sum of quantiles of the portfolio loss distribution in the interval $[\alpha, 1]$. CVaR is considered more appealing than VaR because it takes the contributions from very rare losses into account. Moreover, it is a coherent risk measure and therefore also convex. However, it is obvious that above definition for the CVaR heavily depends on the mathematically troublesome VaR, which complicates the optimization of CVaR. Rockafellar's and Uryasev's main contribution is their proof that the CVaR optimization can be achieved by minimizing a function $F_\alpha(\mathbf{w},\gamma)$ with respect to $\mathbf{w}$ and γ:

$$\min_{\mathbf{w}\in\mathbf{W}} CVaR_\alpha(\mathbf{w}) \quad = \quad \min_{(\mathbf{w},\gamma)\in\mathbf{W}\times\mathbb{R}} F_\alpha(\mathbf{w},\gamma), \quad \text{where}$$

$$F_\alpha(\mathbf{w},\gamma) = \gamma + \frac{1}{1-\alpha}\int_{\mathbf{R}\in\mathbb{R}^N} max(f(\mathbf{w},\mathbf{R})-\gamma, 0)p(\mathbf{R})d\mathbf{R}.$$

They show that $F_\alpha(\mathbf{w},\gamma)$ is convex and continuously differentiable. Moreover, the CVaR can now be optimized without having to calculate the VaR on which its definition depends. In a last step, the continuous joint density function $p(\mathbf{R})$ is approximated with a number of discrete scenarios $\mathbf{R}_s$ for $s = 1,\dots,S$, which typically represent historical returns or simulated returns (e.g. via Monte Carlo or bootstrapping) resulting in the following approximation:

$$\widetilde{F}_\alpha(\mathbf{w},\gamma) = \gamma + \frac{1}{S(1-\alpha)}\sum_{s=1}^{S} max(f(\mathbf{w},\mathbf{R}_s)-\gamma,0).$$

In this thesis, the minimum Conditional Value-at-Risk portfolio is achieved by using historical returns. In section 4.8 a simulation-based Worst-Case CVaR optimization approach will be introduced. The minimization of $\widetilde{F}_\alpha$ over $\mathbf{W}\times\mathbb{R}$ can be further reduced to linear programming if the allowable set $\mathbf{W}$ is convex and the loss function $f(\mathbf{w},\gamma)$ is linear with respect to $\mathbf{w}$. By introducing auxiliary real variables z_s with linear constraints for $s = 1,\dots,S$, one can get rid of the nonlinear max-function and the optimization problem is:

$$\underset{\mathbf{w}\in\mathbf{W},\gamma\in\mathbb{R},z_1,\dots,z_S}{\text{argmin}} \quad \gamma + \frac{1}{S(1-\alpha)}\sum_{s=1}^{S} z_s, \quad \text{subject to} \quad z_s \geq 0 \quad \text{and} \quad z_s \geq f(\mathbf{w},\mathbf{R}_s)-\gamma.$$

Solving above optimization problem produces the portfolio allocation that exhibits the lowest Conditional Value-at-Risk possible. This linearized version of the optimization problem allows for highly efficient routines that enable the investor to find a solution for portfolios of virtually any size and complexity (Acerbi, 2007, p. 364). It is apparent that not just the weight vector $\mathbf{w}$ but also the quantile level γ must be optimized in order to achieve the minimum CVaR portfolio. Therefore, by optimizing CVaR, also the VaR is calculated as a side effect. This quantile-based portfolio optimization technique does not depend on any distributional assumption and therefore also works for non-normally distributed returns. In the special case where the underlying joint return distribution is multivariate normal, the results from minimizing CVaR coincide with the solutions obtained by minimizing the portfolio variance or the portfolio Value-at-Risk (Rockafellar & Uryasev, p. 29). This result is used in several studies to investigate the accuracy of the MinCVaR approach as will be discussed in the following subsection.

4.6.2 Model Discussion and Estimation Error Mitigation

Lim, Shanthikumar and Vahn (2011, pp. 163 - 168) claim that the CVaR is a coherent but fragile risk measure because a large number of observations is necessary to estimate this tail statistic. They exploit the aforementioned finding that the results from minimizing the portfolio CVaR, portfolio VaR and portfolio variance coincide under a multivariate normal distribution and conduct an extensive simulation study to assess the effect of estimation errors in the CVaR optimization. In order to isolate the errors stemming from the mean estimation, they consider the global minimum CVaR portfolio as defined in 4.6.1. Note that this is the same approach as in the minimum variance optimization explained in 4.4, where the estimation error was reduced by neglecting the mean of the mean-variance framework. They find that portfolios obtained by minimizing CVaR are unreliable due to estimation errors of the CVaR. More precisely, they argue that the sample estimator of $\widetilde{F}_\alpha(\mathbf{w}, \gamma)$ is biased in the direction of underestimation, resulting in a significant estimation error compared to the true CVaR. This underestimation is then magnified by the optimization procedure, leading to the issue that an investor can substantially underestimate the true portfolio CVaR and might therefore be exposed to more risk than suggested by the optimization. However, it is important to mention that they justify the assumption of a multivariate normal distribution by working with monthly return series. As shown in the appendix, daily returns usually lead to non-normal return distributions and because the CVaR optimization can be applied under any kind of distribution, it may still make sense to optimize with respect to CVaR. Moreover, Lim, Shanthikumar and Vahn's main results are based on the case where $S = 50$ (i.e. approximately 4 years of monthly data) are used to perform the CVaR estimations and they show that the drastic underperformance of MinCVaR compared to MinVar diminishes with increasing sample size (p. 170).

Another issue of optimizing with respect to CVaR relates to its axiomatic properties. According to Cont, Deguest and Scandolo (2008, pp 3-10), coherency (in fact, sub-additivity) and estimation robustness cannot coexist for the entire class of distribution-based risk measures. While the CVaR has the advantage of being a coherent risk measure, it lacks robustness with respect to small changes in the data set. The VaR, on the other hand, fails to satisfy the sub-additivity property and is therefore less sensitive to outliers in the data. Moreover, the researchers find that the same risk measure can exhibit very different sensitivities depending on the selected estimation procedure. Using a data set of 1000 observations they show that an additional outlier (a very large portfolio loss) changes the historical CVaR dramatically while the CVaR estimated from simulated returns (Laplace distribution) is less sensitive. According to Boudt, Peterson and Croux (2009, p. 79), the advantage of using the empirical distribution function over hypothetical distributions is that only information of the return series is used for the estimations and hence distributional mis-specifications can be avoided. They support the findings from Cont, Deguest and Scandolo by claiming that the disadvantage of using historical data is a larger variation from out-of-sample observations compared to a correctly specified parametric class of distribution functions. Moreover, they address the issue of estimating the portfolio moments in a robust way (pp. 88-90). Optimizing the CVaR requires the estimation of the first four moments of the portfolio return distribution. There are essentially two main approaches to obtain robust estimates of those moments. First, one can consider alternative, more robust estimators than the sample moments. Another approach – suggested by Boudt et al. – is to first clean the data in a robust way and to then take the sample moments of the cleaned data. They explicitly stress the utility of this approach in a portfolio construction context. Therefore, in an attempt to mitigate the impact of extremes on the obtained portfolio results, their data winsorization approach is applied throughout this thesis. This data cleaning method is designed in the way that only observations that belong to the $1-\alpha$ most extreme observations are adjusted:

As a first step, the data is sorted by extremeness. In order to assess the extremeness of the assets' return observations in period t denoted as r_t, they use the squared Mahalanobis distance defined as as $d_t^2 = (r_t - \boldsymbol{\mu})'\boldsymbol{\Sigma}^{-1}(r_t - \boldsymbol{\mu})$ for $t = 1, \ldots, T$. The parameters $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ are estimated in a corrected way in order to ensure their consistency as well as their robustness against the $1-\alpha$ most extreme returns. The result is an ordered sequence of estimated squared Mahalanobis distances $d^2_{(1)}, \ldots, d^2_{(T)}$, where $d^2_{(j)} \leq d^2_{(j+1)}$ (i.e. ordered in an increasing manner). Second, the assets' return observations are qualified as outliers when their estimated squared Mahalanobis distance is larger than the empirical α-quantile $d^2_{(\lfloor \alpha T \rfloor)}$ and exceeds a very extreme quantile of the chi-squared distribution function with N degrees of freedom, where $\lfloor . \rfloor$ is the floor function. In line with Boudt, Peterson and Croux (p. 90), the 99.9% quantile is chosen and is denoted as $\chi^2_{N,0.999}$. Third and last, the so-called multivariate winsorization is performed by cleaning only those returns that are qualified as outliers in the previous step by replacing those returns r_t with

$$r_t * \sqrt{max(d^2_{\lfloor \alpha T \rfloor}, \chi^2_{N,0.999})/d^2_t}.$$

It is important to realize that the added value of this data cleaning process lies in creating a more robust and stable estimation of the return distribution generating the majority of the data (ibid.). Moreover, the assets' returns of period t (r_t) are considered jointly when qualifying them as extreme observations such that only very extreme trading days are detected and reduced in magnitude. Obviously, this comes at the cost of less conservative risk measure estimates because some of the most extreme observations are scaled down. However, the advantage of this robust method is that it does not remove any data from the series but only decreases the magnitude of the very extreme events with the goal of improving the CVaR estimation.

In order to understand how this data winsorization approach affects the CVaR robustness as well as to investigate what role the number of observation plays in the CVaR estimations, a CVaR sensitivity analysis is performed. Similar to the analysis of Cont, Deguest and Scandolo (2008), an additional portfolio return outlier (a very negative portfolio return) is appended to three portfolio return series of different sizes. Consequently, the percentage change in the CVaR estimate is computed for different CVaR estimators with and without data winsorization. The results are reported in A.9. It is apparent that the data winsorization approach of Boudt et al. effectively binds the CVaR sensitivity, which would otherwise be unbounded due to its linear dependence on the outlier size (this stems from the sub-additivity property of CVaR). Moreover, the number of observations has a negative effect on the CVaR sensitivity, meaning that a larger sample size decreases the sensitivity of the CVaR estimate to an additional outlier in the sample.

Besides the discussed issues regarding the estimation process and potential errors therein, the CVaR optimization generally produces portfolio weights with a strong concentration in a few assets (Boudt, Carl & Peterson, 2013, p. 10). It is therefore very common that the optimization procedure 'kicks out' the majority of investible assets by assigning a weight of zero to them. This can lead to a large risk concentration in only a few assets, which can be avoided by employing the Minimum CVaR Concentration portfolio.

4.7 Minimum CVaR Concentration: MCC Portfolio

4.7.1 Model Description

The minimum CVaR concentration portfolio (henceforth abbreviated as MCC portfolio) attempts to create a well-diversified portfolio allocation with respect to the underlying assets' downside risk contributions. In short, the portfolio downside risk itself is being diversified. According to Pfaff (2013, pp. 192-193), such a portfolio construction approach is motivated by the empirical observation that the individual assets' risk contributions are a good predictor for the actual portfolio losses. Therefore, by diversifying the assets' risk contributions, portfolio losses might be limited in comparison to portfolio allocations with high risk concentrations. Following the definitions in Pfaff (ibid.), the risk contribution of an asset i is defined as

$$C_i M_{\mathbf{w}\in\mathbf{W}} = w_i \frac{\delta M_{\mathbf{w}\in\mathbf{W}}}{\delta w_i} \quad \text{for } i = 1, \ldots, N,$$

where $M_{\mathbf{w}\in\mathbf{W}}$ denotes a linear homogeneous risk measure and w_i is the portfolio weight of the ith asset, which belongs in the set of fully invested and shortsale-constrained portfolios $\mathbf{W}$. As such, $M_{\mathbf{w}\in\mathbf{W}}$ can be the portfolio standard deviation, Value-at-Risk or Conditional Value-at-Risk because all of these risk measures are linear homogeneous and by Euler's homogeneity theorem, the total portfolio risk equals the sum of the assets' risk contributions. As already discussed in 4.6.2 and outlined in the Appendix on p. 77, CVaR entails all properties of a coherent risk measure and is a convex function of the portfolio weights, which simplifies the optimization procedure enormously. Therefore, the portfolio CVaR is chosen as the risk measure entering above equation. In order to ease the interpretation of the CVaR risk contributions, they are standardized by the total CVaR. This produces the assets' percentage CVaR contributions, which are defined as:

$$\%C_i CVaR_\alpha(\mathbf{w}) = \frac{C_i CVaR_\alpha(\mathbf{w})}{CVaR_\alpha(\mathbf{w})} = \frac{w_i}{CVaR_\alpha(\mathbf{w})} \frac{\delta CVaR_\alpha(\mathbf{w})}{\delta w_i} \quad \text{for } i = 1, \ldots, N.$$

The CVaR risk contributions have a financial meaning and are directly linked to the portfolio downside risk concentration: the CVaR contribution of asset i corresponds to asset i's expected contribution to the portfolio return when the portfolio loss exceeds the $VaR_\alpha(\mathbf{w})$. In other words, a high (percentage) CVaR risk contribution of an asset implies a large portfolio downside risk concentration on that asset and vice versa.

According to Boudt, Carl and Peterson (2013, p. 40), for many practical applications, achieving a risk parity that requires all assets to contribute equally to portfolio risk is too restrictive to be solved. As an alternative, they propose to minimize the largest CVaR risk contribution in the portfolio. This bears the advantage that such portfolio optimizations can be easily combined with additional investor objectives and constraints (e.g. return targets, drawdown constraints etc.) while still generating portfolios that are similar to the risk parity portfolio. They define the portfolio CVaR concentration as the largest CVaR contribution among all assets in the portfolio:

$$C_\alpha(\mathbf{w}) = \max_i C_i CVaR_\alpha(\mathbf{w}) \quad \text{for } i = 1, \ldots, N.$$

To clarify the objective of MCC portfolio optimizations, the portfolio CVaR concentration can be rewritten in terms of the portfolio CVaR and the largest percentage CVaR contribution:

$$C_\alpha(\mathbf{w}) = CVaR_\alpha(\mathbf{w}) * max\{\%C_1 CVaR_\alpha(\mathbf{w}), \ldots, \%C_N CVaR_\alpha(\mathbf{w})\}.$$

Obviously, the first factor in above equation is minimized by the minimum CVaR portfolio as explained in 4.6.1. What is new in the MCC portfolio is the second term that is only minimized when the portfolio achieves an equal risk contribution property because the largest percentage contribution of all assets is minimized when they all exhibit the same risk contribution. To understand this, note that $max\{\%C_1 CVaR_\alpha(\mathbf{w}), \ldots, \%C_N CVaR_\alpha(\mathbf{w})\} \geq 1/N$. Minimizing the product of these two factors implies that the MCC portfolio achieves a balance between portfolio risk diversification and total risk minimization (p. 45). Consequently, the MCC portfolio produces a well-diversified (but not perfectly equal) allocation with respect to the assets' downside risk contributions when solving

$$\underset{\mathbf{w} \in \mathbf{W}}{\text{argmin}} \; C_\alpha(\mathbf{w}) = \underset{\mathbf{w} \in \mathbf{W}}{\text{argmin}} \max_i C_i CVaR_\alpha(\mathbf{w}) \quad \text{for } i = 1, \ldots, N.$$

The partial derivatives from $C_i CVaR_\alpha(\mathbf{w})$ enter the optimization procedure and their estimation is the key to the MCC portfolio. Theoretically, it is possible to estimate the partial derivatives using historical or simulated data. The issue, however, is that in such a portfolio optimization, the risk contributions need to be evaluated for a large number of possible portfolio weights such that fast and explicit estimators are required. The reason for this is that in many cases the CVaR concentration is not a convex function of the portfolio weights and $C_\alpha(\mathbf{w})$ may also not be differentiable. Therefore, a derivative-free global optimizer is necessary to find the MCC portfolio. In order to overcome this issue, the authors suggest to use the modified CVaR estimator (p. 43).

Boudt, Peterson and Croux (2008) were the first to provide all the formulas required for decomposing the modified VaR and modified CVaR into the risk contributions of the assets in the portfolio and to illustrate their practical use for portfolio construction. They claim that the definition of risk contribution is only useful in practice if a risk measure is chosen for which the derivative with respect to the portfolio weights can be easily computed (p. 82). Choosing the CVaR as a risk measure implies that the first four portfolio moments are involved in the estimation procedures, for which the definitions were already provided in 4.1. The computation of the estimated VaR's and CVaR's derivatives is very challenging because the estimator cannot be expressed as an explicit function of the portfolio weights (p. 84). As a simplification, normality could be assumed in order to obtain explicit formulas for VaR and CVaR. Obviously, this simplification comes at a cost of suboptimal estimates when the data does not follow a normal distribution. In order to take the stylized facts of financial return data into account, the modified VaR and CVaR will be introduced.

Generally speaking, the foundation of modified VaR and CVaR is to assume a Gaussian distribution and to correct its quantiles for the portfolio skewness and excess kurtosis using series expansions around the Gaussian distribution. In order to reduce the complexity, the subsequent definitions are based on the portfolio profit and loss distribution instead of the loss function (as in 4.6.1). The immediate consequence of this is that now the loss probability $\ell = 1 - \alpha$ enters the formulas and since the confidence level α is set to 0.95 throughout this thesis, ℓ equals 0.05. Moreover, the obtained VaR and CVaR values must be multiplied with (-1) in order to report the respective risk measure as a positive number. To begin the derivations of mVaR and mCVaR, the random variable $R_\mathbf{w}$ capturing the portfolio return is expressed in a location-scale representation:

$$R_\mathbf{w} = \mathbf{w}'\boldsymbol{\mu} + \sqrt{m_2} * u,$$

where u is a random variable with zero mean, unit variance and distribution function $G(.)$. Under this representation, the portfolio Value-at-Risk for the loss probability ℓ is given by

$$VaR_\ell(\mathbf{w}) = -\mathbf{w}'\boldsymbol{\mu} - \sqrt{m_2}G^{-1}(\ell),$$

where $G^{-1}(.)$ is the quantile function associated to $G(.)$. Generally, $G(.)$ is assumed to be Gaussian. This assumption is further improved by adjusting the Gaussian quantiles for higher moments in the data using the rth order Edgeworth expansion of $G(.)$ around the Gaussian distribution function $\Phi(.)$:

$$G_r(z) = \Phi(z) - \phi(z)\sum_{i=1}^{r} P_i(z),$$

where z denotes the critical value, $P_i(z)$ is a polynomial in z and $\phi(.)$ denotes the Gaussian probability density function. Correspondingly, the rth order Cornish-Fisher expansion of the quantile function $G^{-1}(.)$ around the Gaussian quantile function $\Phi^{-1}(.)$ equals

$$G_r^{-1}(\ell) = \Phi^{-1}(\ell) + \sum_{i=1}^{r} P_i^*(\Phi^{-1}(\ell)),$$

and only the polynomials relating to the second-order expansion are required (p. 84):

$$\begin{aligned}
P_1(z) &= P_1^*(z) = \frac{1}{6}(z^2-1)s_p \\
P_2(z) &= \frac{1}{24}(z^3-3z)k_p + \frac{1}{72}(z^5-10z^3+15z)s_p^2 \\
P_2^*(z) &= \frac{1}{24}(z^3-3z)k_p - \frac{1}{36}(2z^3-5z)s_p^2,
\end{aligned}$$

where s_p and k_p are the portfolio skewness and the portfolio excess kurtosis, respectively. Finally, the modified VaR is an explicitly defined estimator for VaR that is denoted as

$$\begin{aligned}
mVaR_\ell(\mathbf{w}) &= -\mathbf{w}\boldsymbol{\mu} - \sqrt{m_2}G_2^{-1}(\ell) \\
&= GVaR_\ell(\mathbf{w}) + \sqrt{m_2}\Big\{-\frac{1}{6}(z_\ell^2-1)s_p - \frac{1}{24}(z_\ell^3-3z_\ell)k_p + \frac{1}{36}(2z_\ell^3-5z_\ell)s_p^2\Big\},
\end{aligned}$$

where $z_\ell = \Phi^{-1}(\ell)$ and $GVaR_\ell(\mathbf{w}) = -\mathbf{w}'\boldsymbol{\mu} - \sqrt{m_2}\Phi^{-1}(\ell)$, which is the VaR for loss probability ℓ under a Gaussian distribution. It is worthwhile to note that mVaR coincides with GVaR when the skewness and excess kurtosis are zero (as is the case under a Gaussian distribution). This representation of mVaR allows for a numerical computation and enables its decomposition into the risk contributions of the different assets in the portfolio (remember that VaR is also a linear homogeneous risk measure, hence the risk decomposition can be performed for it as well). In order to make the MCC portfolio computable, the same approach is required for the CVaR. Similarly, Boudt et al (2008, pp. 85 - 86) consider an estimator of the portfolio CVaR that uses asymptotic expansions around the Gaussian distribution in order to take into account the third and fourth moment of the asset returns' distributions. The portfolio CVaR under the location-scale representation is given by

$$CVaR_{\ell}(\mathbf{w}) = -\mathbf{w}'\boldsymbol{\mu} - \sqrt{m_2}\,\mathbb{E}_G[z \mid z \leq G^{-1}(\ell)].$$

For a loss probability ℓ, they define the modified CVaR as the expected value of all returns below the ℓ Cornish-Fisher quantile, where the expectation is computed under the second order Edgeworth expansion of the true distribution function $G(.)$:

$$mCVaR_{\ell}(\mathbf{w}) = -\mathbf{w}'\boldsymbol{\mu} - \sqrt{m_2}\,\mathbb{E}_{G_2}[z \mid z \leq g_{\ell}], \text{ with } g_{\ell} = G_2^{-1}(\ell) \text{ and}$$

$$\mathbb{E}_{G_2}[z|z \leq g_{\ell}] = -\frac{1}{\ell}\left\{\phi(g_{\ell})+\frac{1}{24}[I^4-6I^2+3\phi(g_{\ell})]k_p+\frac{1}{6}[I^3-3I]s_p+\frac{1}{72}[I^6-15I^4+45I^2-15\phi(g_{\ell})]s_p^2\right\},$$

$$\text{where } I^q = \begin{cases} \sum_{i=1}^{q/2}\left(\frac{\prod_{j=1}^{q/2}2j}{\prod_{j=1}^{i}2j}\right)g_{\ell}^{2i}\phi(g_{\ell}) + \left(\prod_{j=1}^{q/2}2j\right)\phi(g_{\ell}), & \text{for } q \text{ even} \\ \sum_{i=0}^{q^*}\left(\frac{\prod_{j=0}^{q^*}(2j+1)}{\prod_{j=0}^{i}(2j+1)}\right)g_{\ell}^{2i+1}\phi(g_{\ell}) - \left(\prod_{j=0}^{q^*}(2j+1)\right)\Phi(g_{\ell}), & \text{for } q \text{ odd} \end{cases}$$

with $q^* = (q-1)/2$. The derivative of mCVaR with respect to the portfolio weights needs to be computed in order to have all the required pieces of solving for the MCC portfolio. In the light that the portfolio weights enter k_p, s_p and therefore also g_{ℓ}, G_2 and G_2^{-1}, this is an analytically very tedious exercise. The formulas for the derivatives of mVaR and mCVaR can be found in the appendix of Boudt, Peterson and Croux (2008). Despite their complexity, these formulas can be efficiently translated into an algorithm that computes both mCVaR and its derivative (similar for mVaR) within split seconds, even for portfolios with a very large number of assets. This workaround enables the implementation of the MCC portfolio and takes the non-normality of the return data into account.

4.7.2 Model Discussion and Implementation

It is crucial to realize that the modified VaR and CVaR are approximations of their true values. Boudt, Peterson and Croux (p. 87) find that these approximations are very reliable for moderate values of skewness and kurtosis and that they are certainly better than simply assuming a Gaussian distribution. However, they stress that modified CVaR is sensitive to extreme deviations from normality. Moreover, they find that in case of negative skewness, modified VaR and CVaR tend to be too pessimistic whereas the Gaussian counterparts are too optimistic. The relationship is inverted in the case of positive skewness, which indicates the necessity of thoroughly checking the return distributions before employing their methods. Another discussion topic is the loss probability ℓ. The researchers find that the approximation of mCVaR becomes less reliable for $\ell \to 0$ and drops to zero if ℓ is very close to zero. In order to avoid this mCVaR-related issue, the modified CVaR is replaced by the modified VaR if the mCVaR is lower than the mVaR by using:

$$mCVaR_{\ell}^{*}(\mathbf{w}) = -\mathbf{w}'\boldsymbol{\mu} - \sqrt{m_2} * min\{\mathbb{E}_{G_2}[z \mid z \leq g_{\ell}], g_{\ell}\}.$$

Finally, the MCC portfolio construction has to be implemented and solved using a derivative-free global optimizer. As shown in the previous subsection, the MCC portfolio optimization is heavily based on the estimation of the first four portfolio moments. Therefore – similarly to the explanations in 4.6.2 – the mCVaR is calculated using winsorized data in order to reduce the impact of outliers on the portfolio weights and to obtain robust moment estimations. The optimization problem is finally solved by using the Differential Evolution algorithm as explained in Ardia et al (2010), resulting in a portfolio that has a near equal CVaR contribution characteristic. In the appendix on p. 82, a CVaR contribution analysis is performed for the equal-weighted, MinCVaR and MCC portfolios in order to explicitly show the functionality of the MCC portfolio. The general insights are that an equal-weighted portfolio does not result in a well-diversified risk contribution allocation because of the non-linear dependence of the portfolio CVaR contributions on the weights. Moreover, the analysis provides indication for the analytical proof of Boudt, Carl and Peterson (2013, p. 46) that the MinCVaR approach results in rather concentrated portfolios regarding both portfolio weight and portfolio CVaR contribution. Lastly, the MCC portfolio construction technique achieves to diversify the risk contributions as well as to create a more balanced portfolio compared to the MinCVaR approach.

4.8 Worst-Case CVaR Approach with Multivariate Archimedean Copulas

4.8.1 Archimedean Copulas: Theory and Worst-Case Calibration

Optimizing an investor's portfolio with respect to CVaR implies a critical dependence on the asset returns' distributional assumption because the CVaR is defined as the expectation of the tail distribution above the VaR. Hence, in order to calculate the CVaR, it is necessary to make assumptions on the underlying distribution of the assets (see formulations in 4.6.1). In this section, the goal is to achieve potential improvements and estimation error mitigations of CVaR-based portfolio optimizations by using simulated instead of historical returns. The aggravating factor of such an endeavor, however, is that the returns' marginal distribution functions as well as their joint distribution function are unknown. Zhu and Fukushima (2009, p. 1156 - 1157) investigate robust portfolio optimization and worst-case techniques that relax the returns' distributional assumption. Instead of assuming the precise knowledge of the joint return distribution, they obtain the optimal portfolio solution by optimizing over a prescribed set of distributions that consists of mixtures of possible distribution scenarios. Because of the vast amount of potential distributions and corresponding CVaRs, they define the Worst-Case CVaR (WCVaR) as the CVaR when the worst-case probability distribution in this set of distributions occurs. They provide the analytical proof that WCVaR is a coherent risk measure, which implies that their approach of modelling the returns' uncertain distribution can be efficiently applied to a portfolio optimization context (due to convexity). Kakouris and Rustem (2014, p. 28) extend this robust framework of WCVaR by incorporating copulas that allow to model the tail dependence of the assets' co-movements. Similar to Zhu and Fukushima, they perform the portfolio optimization over a set of copulas, which allows the investor to exploit a variety of dependence structures. The subsequent derivations are based on their contribution with the goal of providing the reader with a clear understanding of how the Worst-Case Copula-CVaR approach works.

An N-dimensional copula is defined as a multivariate distribution function whose marginal distribution functions are uniformly distributed on the closed interval $[0, 1]$:

$$C(u_1, ..., u_N) = P[U_1 \le u_1, ..., U_N \le u_N],$$

where $U_i \sim$ Unif[0,1] and u_i are realizations of U_i for $i = 1,...,N$. Because one obtains an uniformly distributed random variable after transforming any random variable by its marginal distribution, the margins u_i can be replaced by $F_i(x_i)$, where x_i is a realization of a random variable and $F_i(.)$ denotes its marginal distribution function. According to Kakouris and Rustem (p. 29), Sklar's theorem is one of the most important results that links the concept of copulas to probability distributions. Sklar's theorem states that if F is an N-dimensional distribution function with marginal distributions $F_1,..,F_N$, then there exists an N-dimensional copula such that for all $x \in \mathbb{R}^N$:

$$F(x_1,...,x_N) = C(F_1(x_1),...,F_N(x_N)).$$

If the marginal distributions $F_1,...,F_N$ are continuous, then the N-dimensional copula $C(.)$ is unique. The theorem's corollary furthermore states that for any $\mathbf{u} = (u_1,...,u_N)$ with $\mathbf{u} \in [0,1]^N$, we have that

$$C(u_1,...,u_N) = F(F_1^{-1}(u_1),...,F_N^{-1}(u_N)),$$

where $F_i^{-1}(.)$ are the marginal distributions' quasi-inverses for $i = 1,...,N$. Now we have all the necessary parts to derive an important relation between probability density functions and copulas. To do so, let f be the multivariate probability density function of the N-dimensional probability distribution function F and suppose that $f_1,...,f_N$ are the univariate probability density functions of the marginal distributions $F_1,...,F_N$. If we assume that $F(.)$ and $C(.)$ are differentiable, it follows from Sklar's theorem that

$$\frac{\delta^N F(x_1,...,x_N)}{\delta x_1,...,\delta x_N} \equiv f(x_1,...,x_N) = \frac{\delta^N C(F_1(x_1),...,F_N(x_N))}{\delta x_1,...,\delta x_N} = c(u_1,...,u_N)\prod_{i=1}^{N} f_i(x_i).$$

Rearranging above expression shows that the copula probability density function can be represented as the ratio of the joint probability density function to what it would look like under independence (note that under the assumption of independence, the joint probability density function is obtained by the product of the random variables' marginal densities):

$$c(u_1,...,u_N) = \frac{f(x_1,...,x_N)}{\prod_{i=1}^{N} f_i(x_i)}.$$

Using this representation, a copula can be interpreted as the adjustment that we need to make in order to convert the independence probability density function into the true multivariate density function (Sabino da Silva, Ziegelmann & Tessari, 2017, p. 8). In short, copulas allow us to decompose the joint probability density function from the marginal densities. The marginal distributions can be any distribution of our choice (also non-normal) while the selected copula describes the monotonic relation between those marginal distributions. Therefore, the problem of finding the asset returns' joint distribution function is reduced to two steps: First, finding the assets' marginal distributions and second, finding the dependency between them.

The Worst-Case Copula-CVaR optimization approach is based on a very special family of copulas known as Archimedean. The advantage of Archimedean copulas is that they can be constructed using a specific function, which is called 'generator' and allows for a closed-form expression of the copula. Moreover, sampling from multivariate Archimedean copulas is based on the Laplace Transform and is relatively easy compared to other copula families (Melchiori, 2006, p. 2). Because there are many different kinds of copulas, it is very difficult to know which copula correctly describes the asset returns' dependency. Similar to above-mentioned distribution uncertainty setting, Kakouris and Rustem (2014) suggest to use a set of different copulas that are then combined to a mixture copula using different weights (p. 30). Their set consists of three Archimedean copulas called Clayton, Gumbel and Frank. These particular copulas are chosen due to two reasons: First, each of them describes another type of dependency. Clayton and Gumbel are asymmetric copulas that describe negative and positive dependencies, respectively. The Frank copula, on the other hand, is a symmetric copula that still captures the tail dependency better compared to a Gaussian copula. The second reason for working with these copulas is that each of them is easy to calibrate because they all have only one free parameter. Calibrating a copula essentially means to specify its parameter(s) in a way that appropriately describes the dependency in the data. In order to calibrate the three Archimedean copulas, Kendall's τ is selected as dependency measure because it measures the rank correlation and is independent from the random variables themselves (ibid). This makes it a more robust dependence measure compared to the well-known and common Pearson correlation. Moreover, each of the three Archimedean copulas' free parameter has a closed form relation with Kendall's τ, which simplifies the copula calibration enormously. The only issue is that – similar to the correlation matrix – there is a Kendall's τ for each of the asset pairs from which only one needs to be selected. By deliberately choosing the highest τ among all asset pairs, the most extreme behavior for the copulas is considered, which is the first out of two steps of implementing the worst-case approach. Once the copulas are calibrated in the worst possible way, they are combined in a mixture copula. A mixture copula is the convex combination of N-dimensional copulas, which is a copula itself (Hofert et al, 2018, pp. 127-128): Let $C_1,...,C_m$ be N-dimensional copulas and let $\boldsymbol{\pi} = (\pi_1,...,\pi_m)$ be a vector of weights such that $\pi_k \geq 0$ for all $k \in \{1,...,m\}$ with $\sum_{k=1}^{m} \pi_k = 1$ and $m \geq 2$. The mixture of $C_1,...,C_m$ with mixing vector $\boldsymbol{\pi}$ is the copula defined as

$$mix_{\boldsymbol{\pi}}(C_1,...,C_m)(\mathbf{u}) = \sum_{k=1}^{m} \pi_k C_k(\mathbf{u}), \quad \mathbf{u} \in [0,1]^N.$$

4.8.2 Worst-Case Copula-CVaR

This subsection provides the relevant derivations for understanding how the WCVaR and its implementation in portfolio optimization relates to the multivariate mixture copula extension. Note that many of the subsequent derivations are based on Rockafellar's and Uryasev's approach of minimizing CVaR explained in section 4.6.1. In line with Sabino da Silva, Ziegelmann and Tessari (2017, p. 9), let $h(\mathbf{w},\mathbf{u}) = h(\mathbf{w}, F(\mathbf{R}))$ be the loss function and $F(\mathbf{R}) = (F_1(R_1),...,F_N(R_N))^T$ be a set of marginal distributions for $\mathbf{R} \in \mathbb{R}^N$. Also, assume that $\mathbf{u} = (u_1,...,u_N)$ follows a continuous distribution with copula C(.) and $\mathbf{u} \in [0,1]^N$. In line with section 4.6.1, the portfolio loss function's cumulative distribution function is defined in a first step and then the definitions of portfolio VaR and CVaR are provided:

$$\begin{aligned} P[h(\mathbf{w},\mathbf{R}) \leq \gamma] &= \int_{h(\mathbf{w},\mathbf{R})\leq\gamma} f(\mathbf{R})d\mathbf{R} \;=\; \int_{h(\mathbf{w},\mathbf{R})\leq\gamma} c(F(\mathbf{R})) \prod_{i=1}^{N} f_i(R_i)d\mathbf{R} \\ &= \int_{h(\mathbf{w},\mathbf{F}^{-1}(\mathbf{u}))\leq\gamma} c(\mathbf{u})d\mathbf{u} \;=\; C(\mathbf{u} \mid \widetilde{h}(\mathbf{w},\mathbf{u}) \leq \gamma), \end{aligned}$$

where $f_i(R_i) = \frac{\delta F_i(R_i)}{\delta R_i}$, $\mathbf{F}^{-1}(\mathbf{u}) = (F_1^{-1}(u_1),...,F_N^{-1}(u_N))$ and $\widetilde{h}(\mathbf{w},\mathbf{u}) = h(\mathbf{w},\mathbf{F}^{-1}(\mathbf{u}))$. Correspondingly, the copula-based VaR and CVaR are defined as:

$$VaR_\alpha(\mathbf{w}) = \min\{\gamma \in \mathbb{R} : \; C(\mathbf{u} \mid \widetilde{h}(\mathbf{w},\mathbf{u}) \leq \gamma) \geq \alpha\}$$

$$CVaR_\alpha(\mathbf{w}) = \frac{1}{1-\alpha} \int_{\widetilde{h}(\mathbf{w},\mathbf{u})\geq VaR_\alpha(\mathbf{w})} \widetilde{h}(\mathbf{w},\mathbf{u})c(\mathbf{u})d\mathbf{u}.$$

As explained on p. 32 of this thesis, Rockafellar and Uryasev (2000) found that the CVaR can be calculated independently of VaR by minimizing a specific auxiliary function. Applied to the present copula-framework, the function to be minimized with respect to γ is:

$$H_\alpha(\mathbf{w},\gamma) = \gamma + \frac{1}{1-\alpha}\int_{\mathbf{u}\in[0,1]^N} max(\widetilde{h}(\mathbf{w},\mathbf{u}) - \gamma, 0)c(\mathbf{u})d\mathbf{u}.$$

In order to implement the uncertainty regarding what copula to choose, a set of copulas is considered in a mixture copula approach such that many possibilities are covered:

$$C(.) \in \mathscr{C}_{mix} \equiv \left\{ \sum_{i=1}^{d} \pi_i C_i(.) : \sum_{i=1}^{d} \pi_i = 1,\ \pi_i \geq 0,\ i = 1,\ldots,d \right\}.$$

Note that throughout this thesis, d is set to 3 because Gumbel, Clayton and Frank copulas are considered. In this mixture copula setting, the auxiliary function $H_\alpha(\mathbf{w},\gamma)$ changes to:

$$\begin{aligned} H_\alpha(\mathbf{w},\gamma,\boldsymbol{\pi}) &= \gamma + \frac{1}{1-\alpha}\int_{\mathbf{u}\in[0,1]^N} max(\widetilde{h}(\mathbf{w},\mathbf{u}) - \gamma, 0)\sum_{i=1}^{d}\pi_i c_i(\mathbf{u})d\mathbf{u} \\ &= \sum_{i=1}^{d} \pi_i H_\alpha^i(\mathbf{w},\gamma), \text{ where} \\ H_\alpha^i(\mathbf{w},\gamma) &= \gamma + \frac{1}{1-\alpha}\int_{\mathbf{u}\in[0,1]^N} max(\widetilde{h}(\mathbf{w},\mathbf{u}) - \gamma, 0)c_i(\mathbf{u})d\mathbf{u} \text{ for } i = 1,\ldots,d. \end{aligned}$$

Ultimately, the function $H_\alpha(\mathbf{w},\gamma,\boldsymbol{\pi})$ is used to calculate the Worst-Case Copula-CVaR:

$$WCVaR_\alpha(\mathbf{w}) = \min_{\gamma\in\mathbb{R}}\max_{\pi\in\Pi} H_\alpha(\mathbf{w},\gamma,\boldsymbol{\pi}) = \min_{\gamma\in\mathbb{R}}\max_{\pi\in\Pi}\sum_{i=1}^{d}\pi_i H_\alpha^i(\mathbf{w},\gamma),$$

where $\Pi = \left\{\pi \in \mathbb{R}^d \mid \sum_{i=1}^{d}\pi_i = 1,\ \pi_i \geq 0,\ \forall i = 1,\ldots,d\right\}$. Hence, the Worst-Case Copula-CVaR is given by the mixture copula that produces the largest CVaR. Similar to section 4.6.1, the integral in $H_\alpha^i(\mathbf{w},\gamma)$ can be approximated by generating a number of discrete scenarios $j = 1,\ldots,\mathscr{J}$ of copula i's uniform margins denoted as $\mathbf{u}_j^{[i]}$:

$$\widetilde{H}_\alpha^i(\mathbf{w},\gamma) = \gamma + \frac{1}{\mathscr{J}(1-\alpha)}\sum_{j=1}^{\mathscr{J}} max(\widetilde{h}(\mathbf{w},\mathbf{u}_j^{[i]}) - \gamma, 0) \text{ for } i = 1,\ldots,d.$$

Finally, the portfolio optimization problem that minimizes the Worst-Case Copula-CVaR over $\mathbf{w} \in \mathbf{W}$ using multivariate mixture copulas is defined as:

$$\min_{\mathbf{w}\in\mathbf{W}} WCVaR_\alpha(\mathbf{w}) = \min_{\mathbf{w}\in\mathbf{W}}\min_{\gamma\in\mathbb{R}}\max_{\pi\in\Pi} H_\alpha(\mathbf{w},\gamma,\boldsymbol{\pi})$$

4.8.3 Model Implementation

In this subsection, the implementation steps of the Worst-Case Copula-CVaR optimization are outlined and discussed. The selected implementation approach is in line with Sabino da Silva, Ziegelmann and Tessari (2017, pp. 11-12), who define the return dynamics as:

$$R_t = \mu_t + \epsilon_t, \text{ where } \epsilon_t = \sigma_t z_t \text{ and } z_t \overset{iid}{\sim} \text{skew-t.}$$

In their application, they investigate portfolios of 50 stocks from the S&P 500 Index and in order to capture the dynamics of the conditional mean (μ_t) and conditional volatility (σ_t), they estimate an AR(1)-GARCH(1,1) model with skew-t distributed innovations (z_t) for each of the univariate return series. Throughout this thesis, the retail investor's portfolio consists of various asset classes and another sample period is used to analyze the various portfolio construction techniques. Therefore, it does not seem to be appropriate to blindly follow the authors' approach of modelling the asset returns. The overarching goal is to simulate asset returns that capture the stylized facts of each of the univariate time series reasonably. Because this will be performed in a dynamic manner throughout the backtesting, it might not be optimal for the investor to employ the same model for each period and each asset. Therefore, the modelling approach is adjusted in two instances: First, the conditional means are not strictly modelled by an AR(1) model. Instead, all combinations of ARMA(p,q)-orders up to a maximal order (including pure AR(p) and MA(q) models) are estimated for each of the assets. Since daily returns are modelled, the conditional mean part is very close to zero and hence not as relevant as the conditional volatility part, where it is the goal to capture the stylized fact of volatility clustering. This can be done by estimating a GARCH(1,1) model, which serves as benchmark among volatility models. In order to improve the volatility modelling, each asset return series is tested for leverage effects. If there are significant leverage effects, a GJR-GARCH(1,1) model is selected for the volatility part. Otherwise, the standard GARCH(1,1) model is employed. The ARMA-(GJR)-GARCH models are estimated jointly and the most appropriate model is chosen according to the Bayesian Information Criterion (BIC). The model selection is based on BIC because it puts a larger weight on the penalty term and hence prefers smaller models compared to AIC. Once an appropriate model is selected, data is simulated for μ_t and σ_t. Consequently, the returns' dependency is modelled within the Worst-Case copula CVaR approach. The following steps summarize the entire procedure:

(i) Investigate the existence of significant leverage effects by performing negative and positive size bias tests for each of the assets' return series. If leverage effects are detected, use a GJR-GARCH(1,1) model in the next step. Otherwise, use a GARCH(1,1) model.

(ii) After analyzing ACF and PACF plots of the data, set the highest allowable AR and MA orders. Consequently, estimate all possible combinations of ARMA(p,q)-(GJR)-GARCH(1,1) models with skew-t distributed innovations for each of the assets and select the most appropriate model according to BIC for each asset.

(iii) Using the selected parametric model, simulate $\mathcal{J}$ conditional means ($\widetilde{\mu}_t$) and volatilities ($\widetilde{\sigma}_t$) from the selected models for each of the assets. Construct the standardized residuals (innovations) vectors according to $\hat{z}_{i,t} = \frac{\hat{\epsilon}_{i,t}}{\hat{\sigma}_{i,t}}$ for $i = 1, \ldots, N$ and $t = 1, \ldots, T$.

(iv) Convert the standardized residuals to pseudo-uniform observations: $u_{i,t} = \frac{T}{T+1} F_i(\hat{z}_{i,t})$, where F_i is the empirical distribution function and $\frac{T}{T+1}$ is an asymptotically negligible scaling factor, which ensures that the pseudo-uniform observations fall inside $(0,1)$.

(v) Combine the pseudo-uniform observations of all assets to the matrix $\mathbf{u}$ and calculate Kendall's τ correlation matrix of $\mathbf{u}$. Select the highest τ among all asset pairs and use it to calibrate the N-dimensional Gumbel, Frank and Clayton copulas: $\boldsymbol{\alpha}^* = (\alpha_G^*, \alpha_C^*, \alpha_F^*)$. Generate $\mathcal{J}$ scenarios from each of the copulas.

(vi) Create random copula weights that sum to 1. For each copula weight combination, generate $\mathcal{J}$ scenarios from the corresponding mixture copula (see Hofert et. al, 2018, p. 128) and calculate the CVaR by performing MinCVaR optimization with uniform margins as input. Find the worst-case copula weight combination $\boldsymbol{\pi}^* = (\pi_G^*, \pi_C^*, \pi_F^*)$ that yields the largest CVaR.

(vii) Construct the worst-case multivariate Gumbel-Clayton-Frank mixture copula: $C^{GCF}(\mathbf{u}) = \pi_G^* C^G(\alpha_G^*, \mathbf{u}) + \pi_C^* C^C(\alpha_C^*, \mathbf{u}) + \pi_F^* C^F(\alpha_F^*, \mathbf{u})$. Create the joint distribution function by combining C^{GCF} with the marginal skew-t distributions. Generate $\mathcal{J}$ scenarios for the innovations ($\widetilde{z}_t$). Lastly, determine the simulated returns $R_t^{\text{sim}} = \widetilde{\mu}_t + \widetilde{\sigma}_t \widetilde{z}_t$ for $t = 1, \ldots, \mathcal{J}$.

The simulated returns are used as inputs when optimizing the portfolio weights with CVaR-based objectives (MinWCVaR & WC-MCC optimizations). Throughout this thesis, the maximal orders for AR and MA are both set to 2. The random copula weights consist of 34 weight combinations and $\mathcal{J} = 1000$.

5 Empirical Results

5.1 Out-Of-Sample Backtesting

In order to assess how the various portfolio construction techniques would have fared throughout the past, an extensive backtesting procedure is performed. Inherent to the backtesting are the specifications on (i) what the retail investor does over time and (ii) how she does it. In chapter 4 it was extensively explained what needs to be done when implementing each of the portfolio construction techniques. Regarding the second aspect of backtesting – how the portfolio construction techniques are brought to practice – there is a vast amount of possible approaches. A relevant topic is the discussion about whether and how to rebalance the portfolio. As the financial markets are dynamic and change over time, the optimized portfolio weights are not optimal anymore as soon as the trading bell chimes. Therefore, it seems very plausible that rebalancing can improve the performance of long-term trading strategies. Dichtl, Drobetz and Wambach (2014) investigate periodic and interval rebalancing strategies. While periodic rebalancing demands a reallocation to the optimal portfolio weights at the end of each pre-defined period, interval rebalancing requires the implementation of a no-trade region around the optimal weights (p. 5). In accordance with the common financial wisdom, they find that rebalancing outperforms a simple buy-and-hold strategy. The choice of the particular rebalancing strategy, however, is only of minor economic importance. Assuming that the retail investor is interested in a periodic rebalancing strategy, it is important to realize that the definition of the rebalancing frequency comes with a trade-off: Obviously, the more often the portfolio is rebalanced, the more trading costs are incurred. In order to investigate potential performance improvements from periodic rebalancing, all the portfolio construction techniques are rebalanced yearly, quarterly, monthly and weekly. Daily rebalancing is intentionally excluded as it magnifies computational issues and, more importantly, seems impracticable for a retail investor.

Because the portfolio construction techniques involve estimations of portfolio moments and model parameters, it is necessary to specify how the retail investor performs the estimations over time. Throughout this thesis, an out-of-sample backtesting is carried out. This means that each period's estimations are exclusively based on historical data available up to the respective period. This way, the investor's uncertainty regarding future developments is reflected and biases relating to the backtesting procedure are mitigated. Similar to the retail investors' psychological biases in trading, the backtesting procedure itself is subject to many biases that might distort its results (Kemp, 2011, p. 166). Only a few examples are the 'look-back bias' and 'optimal period bias', where the investigated models are developed with past information in mind and the backtesting is carried out for a time period that produces very favorable results, respectively. Since the investigated portfolio construction techniques are purely data-driven, based on very general statistical and mathematical concepts and the required estimations are performed in an out-of-sample fashion, the look-back bias is mitigated to a great extent. The optimal period bias is avoided simply by using the maximum possible length of the data sample for the assets considered. As for the amount of historical data used throughout the estimations, there are two possibilities. Either the entire available history or a fixed amount of historical data can be used. Because the latter approach might base the portfolio optimizations on very positive market outcomes in periods of long bull markets, the entire available history is used for the estimations. This decision, in turn, opens the discussion about the representativeness of old historical data. As this thesis' sample period ranges from 01/1999 to 12/2019, 'only' 21 years of historical data is used for the very last optimizations in the backtesting procedure (see section 3.2), with the intention of performing the optimizations based on data that incorporates periods of market distress.

Concretely, the out-of-sample backtesting is implemented as follows: The sample's first year of daily data is used to construct the various techniques' initial portfolios. The next portfolio optimization is based on 1 year plus 1 week of daily data in the case of weekly rebalancing (equivalently for monthly, quarterly and yearly rebalancing). This procedure is continued until the very last portfolio optimization, which is based on the entire data sample. The result of this multiperiod optimization problem is a daily series of the respective technique's portfolio returns ranging from the beginning of year 2000 until the end of 2019.

5.2 Discussion of Results

The out-of-sample portfolio allocation performance and its associated risks are investigated by reporting a set of statistics for all portfolio construction techniques and rebalancing frequencies. Cumulative return and annualized return both provide an insight into the performance characteristic of each of the techniques. Volatility, semideviation, Value-at-Risk as well as Conditional Value-at-Risk (both at 95% confidence levels) and maximum drawdown deliver indications about the techniques' risk characteristics. In order to take the risk-reward tradeoff into account, Sharpe Ratio and Sortino Ratio are provided. Lastly, the portfolio turnover is reported in order to assess how much trading is required to implement the different portfolio construction techniques.

In light of the recent developments in the financial industry as well as the brokerage war among financial institutions (see section 2.1.1), no trading costs were assumed in the subsequent analyses. This way, the results remain on a general level and are relatable to a broad spectrum of retail investors with any trading account fee structure. While it is assumed that the reader understands all employed performance and risk statistics, it is worthwhile to provide an explicit definition for the portfolio turnover. In line with DeMiguel, Garlappi and Uppal (2009, p. 1929), let $w_{k,i,t}$ be asset i's portfolio weight at time t under portfolio optimization technique k, where $i = 1, \ldots, N$, $t = 1, \ldots, T$ and $k = 1, \ldots, K$. Moreover, let $w_{k,i,t+1}$ be the optimal portfolio weight after rebalancing at time $t+1$ and let w_{k,i,t^+} be the portfolio weight before rebalancing at $t+1$. Then, the portfolio turnover is defined as the average sum of the absolute value of the trades across all N available assets:

$$\text{Turnover}_k = \frac{1}{T} \sum_{t=1}^{T} \sum_{i=1}^{N} \left(|w_{k,i,t+1} - w_{k,i,t^+}| \right), \quad \text{for } k = 1, \ldots, K.$$

The portfolio turnover can be interpreted as the average percentage of wealth traded in each period. By construction, a higher rebalancing frequency (e.g. weekly or monthly) will lead to a larger portfolio turnover, indicating that more trading is required in order to implement the portfolio construction technique. All performance and risk statistics are based on simple portfolio returns except for Sharpe Ratio and Sortino Ratio, which are based on risk-adjusted simple portfolio returns using the interest rate on three-month Treasury bills as risk-free rate. In order to facilitate comparisons, some of the statistics are annualized (where it is considered appropriate). Table 3 reports the quantitative out-of-sample statistics of the various portfolio construction techniques from January 2000 to December 2019 for yearly, quarterly, monthly and daily rebalancing frequencies. Additionally, the quantitative results are supported qualitatively by wealth trajectory figures (see figures 1, 2 and 3).

	1/N	MaxSR	MinVar	MinSemiVar	MinCVaR	MCC	MinWCVaR	WC-MCC
Yearly Rebalancing:								
Cumul. Return (%)	206.66	113.48	176.80	155.24	177.52	177.86	163.47	207.84
Ann. Return (%)	5.81	3.90	5.27	4.84	5.28	5.29	5.00	5.83
Ann. Volatility (%)	12.26	7.27	5.03	5.10	5.11	7.31	5.49	9.83
Ann. Semideviation (%)	8.77	5.16	3.40	3.44	3.47	5.14	3.69	6.91
Ann. Sharpe Ratio	0.33	0.31	**0.72**	0.63	0.71	0.49	0.61	0.41
Ann. Sortino Ratio	0.53	0.48	1.06	0.94	1.05	0.73	0.92	0.64
VaR$_{0.95}$	0.011	0.007	0.005	0.005	0.005	0.007	0.005	0.009
CVaR$_{0.95}$	0.019	0.011	0.007	0.007	0.007	0.011	0.008	0.015
Max. Drawdown (%)	50.00	30.01	17.82	11.43	17.87	34.27	10.69	45.48
Turnover (%)	0.039	0.069	0.021	0.009	0.024	0.032	0.071	0.087
Quarterly Rebalancing:								
Cumul. Return (%)	201.25	164.13	177.55	154.57	176.57	178.27	157.51	196.52
Ann. Return (%)	5.72	5.02	5.28	4.82	5.26	5.29	4.88	5.63
Ann. Volatility (%)	12.38	6.82	5.00	5.11	5.06	7.25	5.21	9.02
Ann. Semideviation (%)	8.85	4.78	3.38	3.44	3.43	5.10	3.52	6.37
Ann. Sharpe Ratio	0.31	0.49	**0.72**	0.63	0.71	0.49	0.62	0.43
Ann. Sortino Ratio	0.51	0.73	**1.07**	0.94	1.05	0.73	0.93	0.66
VaR$_{0.95}$	0.011	0.007	0.005	0.005	0.005	0.007	0.005	0.008
CVaR$_{0.95}$	0.019	0.010	0.007	0.007	0.007	0.011	0.007	0.013
Max. Drawdown (%)	50.10	28.08	17.49	11.34	17.31	33.83	10.99	39.74
Turnover (%)	0.074	0.090	0.035	0.014	0.040	0.059	0.179	0.237
Monthly Rebalancing:								
Cumul. Return (%)	195.42	165.49	175.27	154.49	176.40	174.72	166.14	196.26
Ann. Return (%)	5.61	5.05	5.24	4.82	5.26	5.23	5.06	5.63
Ann. Volatility (%)	12.55	6.73	5.00	5.11	5.07	7.31	5.28	9.02
Ann. Semideviation (%)	8.97	4.72	3.38	3.45	3.44	5.15	3.57	6.36
Ann. Sharpe Ratio	0.30	0.50	**0.72**	0.62	0.71	0.48	0.65	0.43
Ann. Sortino Ratio	0.50	0.74	1.06	0.94	1.05	0.72	0.97	0.66
VaR$_{0.95}$	0.011	0.007	0.005	0.005	0.005	0.007	0.005	0.008
CVaR$_{0.95}$	0.019	0.010	0.007	0.007	0.007	0.011	0.007	0.013
Max. Drawdown (%)	**50.90**	27.78	17.72	11.37	17.85	34.57	10.61	40.28
Turnover (%)	0.121	0.154	0.055	0.022	0.069	0.101	0.456	0.690
Weekly Rebalancing:								
Cumul. Return (%)	203.70	163.85	178.28	153.99	179.06	178.44	145.57	**215.57**
Ann. Return (%)	5.76	5.01	5.29	4.81	5.31	5.30	4.63	**5.96**
Ann. Volatility (%)	**12.68**	6.67	5.01	5.11	5.10	7.40	5.17	8.94
Ann. Semideviation (%)	**9.03**	4.68	3.40	3.45	3.46	5.20	3.49	6.29
Ann. Sharpe Ratio	0.31	0.50	**0.72**	0.62	0.71	0.48	0.58	0.47
Ann. Sortino Ratio	0.51	0.74	**1.07**	0.93	1.06	0.72	0.87	0.72
VaR$_{0.95}$	**0.011**	0.007	0.005	0.005	0.005	0.007	0.005	0.008
CVaR$_{0.95}$	**0.019**	0.010	0.007	0.007	0.007	0.011	0.007	0.013
Max. Drawdown (%)	50.64	27.81	17.88	11.26	18.09	34.32	10.70	37.75
Turnover (%)	0.231	0.312	0.107	0.036	0.135	0.196	1.433	**2.667**

Table 3: Performance Overview – All Strategies and Rebalancing Frequencies

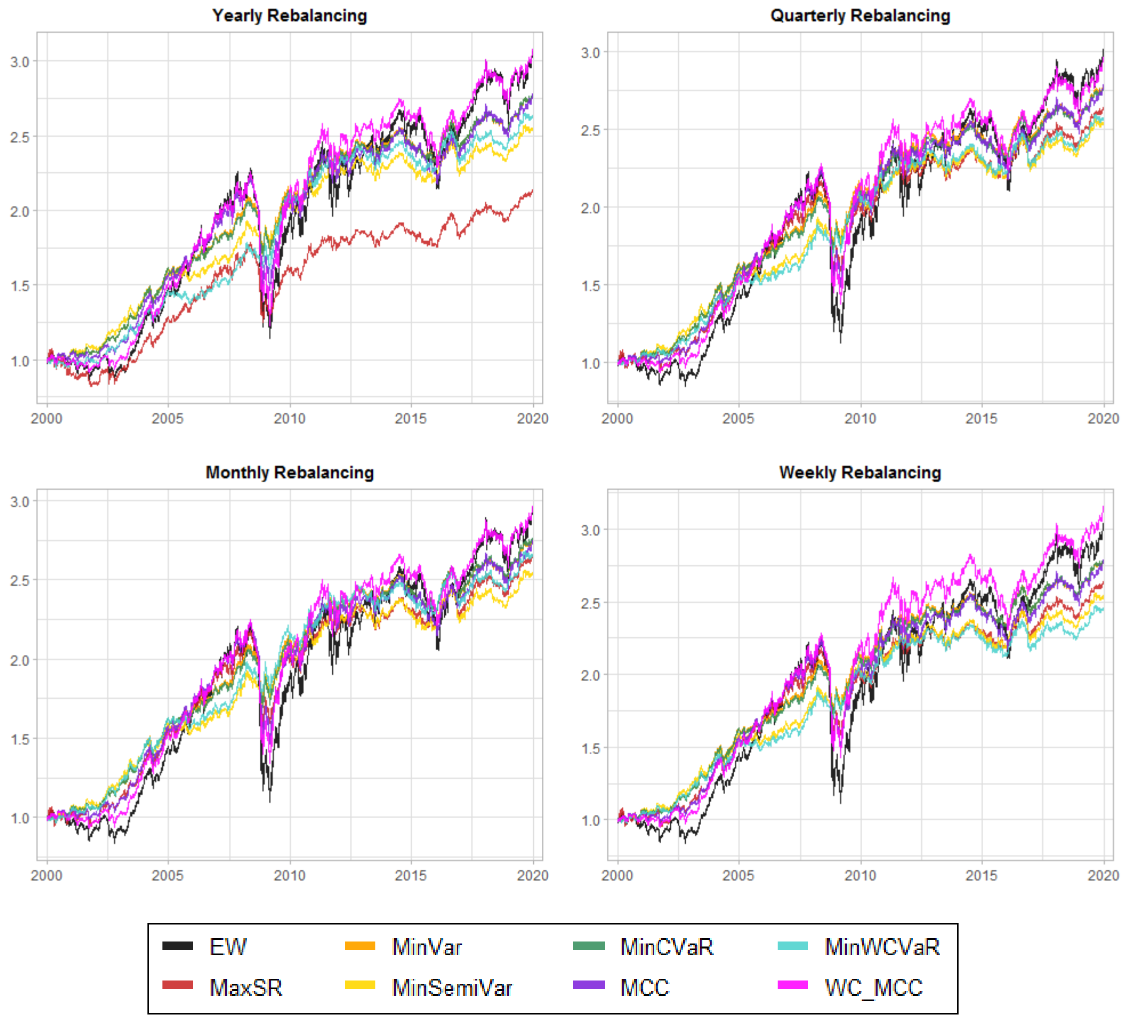

Figure 1: Wealth Trajectory – All Strategies and Rebalancing Frequencies

As a general insight, almost all the very extreme risk and return characteristics (bold in table 3) can be found for weekly rebalancing. It seems that complexity pays off since the sophisticated Worst-Case MCC portfolio [WC-MCC] achieves the largest cumulative as well as annualized returns with weekly rebalancing among all portfolio construction techniques and rebalancing frequencies. As for the pure risk characteristics, the simplistic equal-weight portfolio [1/N or EW] seems to be the worst choice because it exhibits the largest annualized volatility, semideviation, maximum drawdown as well as the largest Value-at-Risk and Conditional Value-at-Risk for all rebalancing frequencies. Again, the largest values are found for weekly rebalancings. These findings indicate that all of the optimization-based portfolio construction techniques are able to reduce the risk characteristics of the portfolio compared to the simple approach of allocating the wealth equally among all investible assets – no matter the rebalancing frequency.

On a risk-adjusted basis, the maximum Sharpe Ratio portfolio [MaxSR] fails to achieve its objective as the largest annualized Sharpe (and Sortino) Ratio is achieved by the minimum variance portfolio [MinVar] for all the rebalancing frequencies. Interestingly, the MinVar and minimum CVaR portfolios [MinCVaR] produce very similar portfolio weights, which leads to the almost identical risk-return characteristics. According to the theory (see 4.6.2), this happens only if the investigated returns are normally distributed. As was thoroughly checked in section 3, the return data of all the assets does not follow a normal distribution. However, among all the investible assets, the World Bond Index produces returns that are closest to a normal distribution. As the portfolio allocation figures A.12 - A.15 point out, both the MinVar and MinCVaR portfolios result in a large concentration of this respective index, which might explain why they do not produce perfectly equal but similar performances. Despite the similarity to the MinCVaR portfolio, the MinVar portfolio clearly achieves its objective and produces the lowest annualized volatilities among all techniques for each of the rebalancing frequencies. The minimum semivariance portfolio [MinSemiVar], however, does not achieve the lowest semideviations since those are achieved by the MinVar portfolio. Nevertheless, the MinSemiVar approach is able to effectively reduce the downward movements of the portfolio as can be seen in figure 1 and as indicated by the maximum drawdowns for each rebalancing frequency. The differences in annualized portfolio semideviation between MinVar and MinSemiVar are vanishingly small and might be the result of the heuristic approach, which is based on an approximation. It could be interesting to measure the extent of the potential approximation errors within the heuristic MinSemiVar approach by evaluating various existing approaches to minimizing the portfolio semivariance. However, this is left open for further research.

Among all portfolio construction techniques, the Minimum CVaR Concentration portfolio [MCC] achieves a balance between risk and return and is located in the middle field. Regarding portfolio turnover, however, the MCC portfolio requires the most trading for its implementation (among all techniques that are based on historical data). As already explained, the portfolio turnover increases in rebalancing frequency as more trades are required to sustain the optimal asset allocation over time. The Worst-Case approaches that are based on simulated data require the most trading, which is discussed further below.

Since it might be difficult to spot differences in figure 1, each portfolio construction technique is displayed and discussed separately. All techniques' performances are compared to an equally-weighted buy-and-hold strategy, which is comparable to being an absolutely passive investor (i.e. 1/N portfolio is constructed at the beginning of the period and held until the sample period end without any changes). The goal is to show what each portfolio construction technique achieves as well as to obtain indications of whether it pays off to rebalance the portfolio more frequently over time.

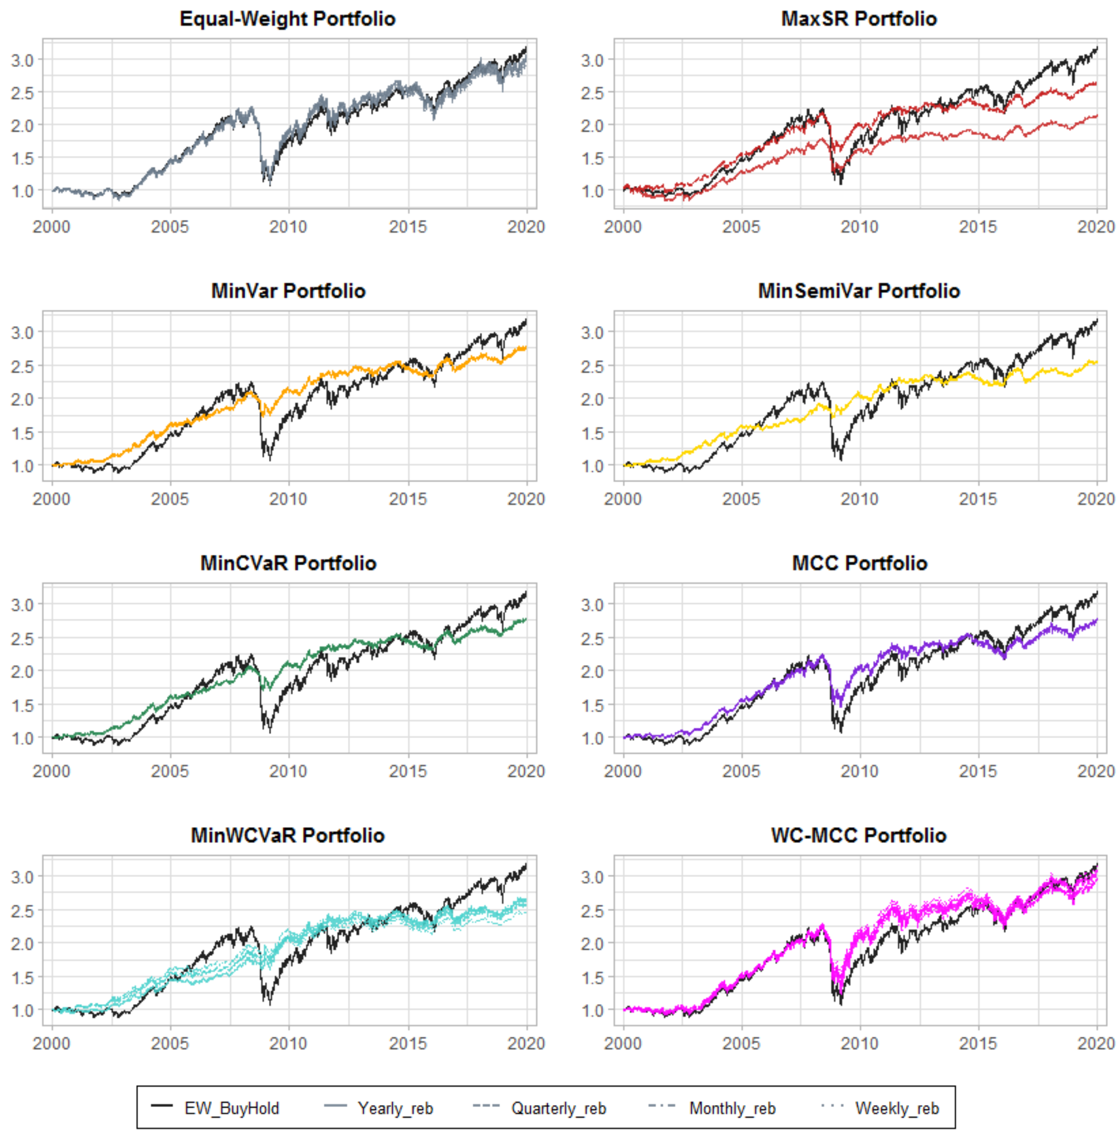

Figure 2: The Effect of Rebalancing for all Portfolio Construction Techniques

Rebalancing the equally-weighted portfolio seems not to pay off substantially in the present context. Up until 2009, all rebalancing frequencies are close to the buy-and-hold strategy. After the financial crisis, however, rebalancing the portfolio on a periodical basis increases the cumulative return evidently. This bull market effect can be explained by the rebound of all assets after the crisis (see figure A.1): An equally-weighted portfolio allows the investor to participate in the rebounds of all asset classes, resulting in an attractive return potential. However, rebalancing the equal-weight portfolio does not improve its risk characteristics.

The maximum Sharpe Ratio portfolio profits the most from periodic rebalancing. Yearly rebalancing (bottom line) is clearly dominated by the more frequent quarterly, monthly and weekly rebalancings (top line, as they are all very close to each other). The MaxSR portfolio strikes a balance between risk and return: the risk (volatility, VaR, CVaR and maximal drawdown) is effectively reduced compared to the equal-weight portfolios for all rebalancing frequencies. As can be deduced from table 3, the Sharpe Ratio increases for the MaxSR portfolio if rebalancing takes place more often. As discussed previously, the MaxSR portfolio fails to achieve the largest Sharpe Ratio among all strategies.

The figures A.12 - A.15 show that the MinVar, MinSemiVar and MinCVaR techniques produce very concentrated portfolios that are similar to each other. Apparently, the retail investor can achieve a very steady and safe growth of capital over time when employing those strategies. The large concentration in one asset, however, can be very problematic since diversification effects are almost nonexistent. As for the effect of rebalancing, MinVar, MinSemiVar and MinCVaR portfolios seem not to benefit substantially from rebalancing the portfolio more often – regarding both return and risk characteristics.

Section 4.7.1 explained that the MCC portfolio strikes a balance between portfolio risk diversification and total risk minimization. Boudt, Carl and Peterson (2013, p. 56) discovered that the MCC portfolio inherits the good risk-adjusted return properties of a MinCVaR portfolio as well as the positive return potential of an equally-weighted portfolio. Figure 2 reflects this finding: The MCC portfolio achieves a similar upward movement as the equally-weighted buy-and-hold portfolio while additionally limiting the risk. This can be seen during the financial crisis (also see table 3). The rebalancing frequency, however, seems not to play a significant role for its implementation. Conversely, the rebalancing frequency matters for the worst-case counterparts of MinCVaR and MCC as examined next.

Table 3 indicates that the copula-based worst-case approaches of MinCVaR and MCC portfolios require the most trading in order to be implemented. This finding can also be obtained from the portfolio allocation plots (A.12 - A.15) that show fluctuating allocations for each rebalancing date. The open question is now whether it is worthwhile to take on the trading effort and to implement such strategies. The multivariate worst-case mixture copula extensions of MinCVaR and MCC portfolio optimizations do not use historical data for the optimizations. Rather, they use the historical data to estimate models and to construct distributions from which worst-case data is simulated. Figure 3 depicts the wealth trajectories of both MinCVaR and MCC portfolios together with their simulation-based extensions in order to obtain clear insights into the effect of rebalancing for those techniques.

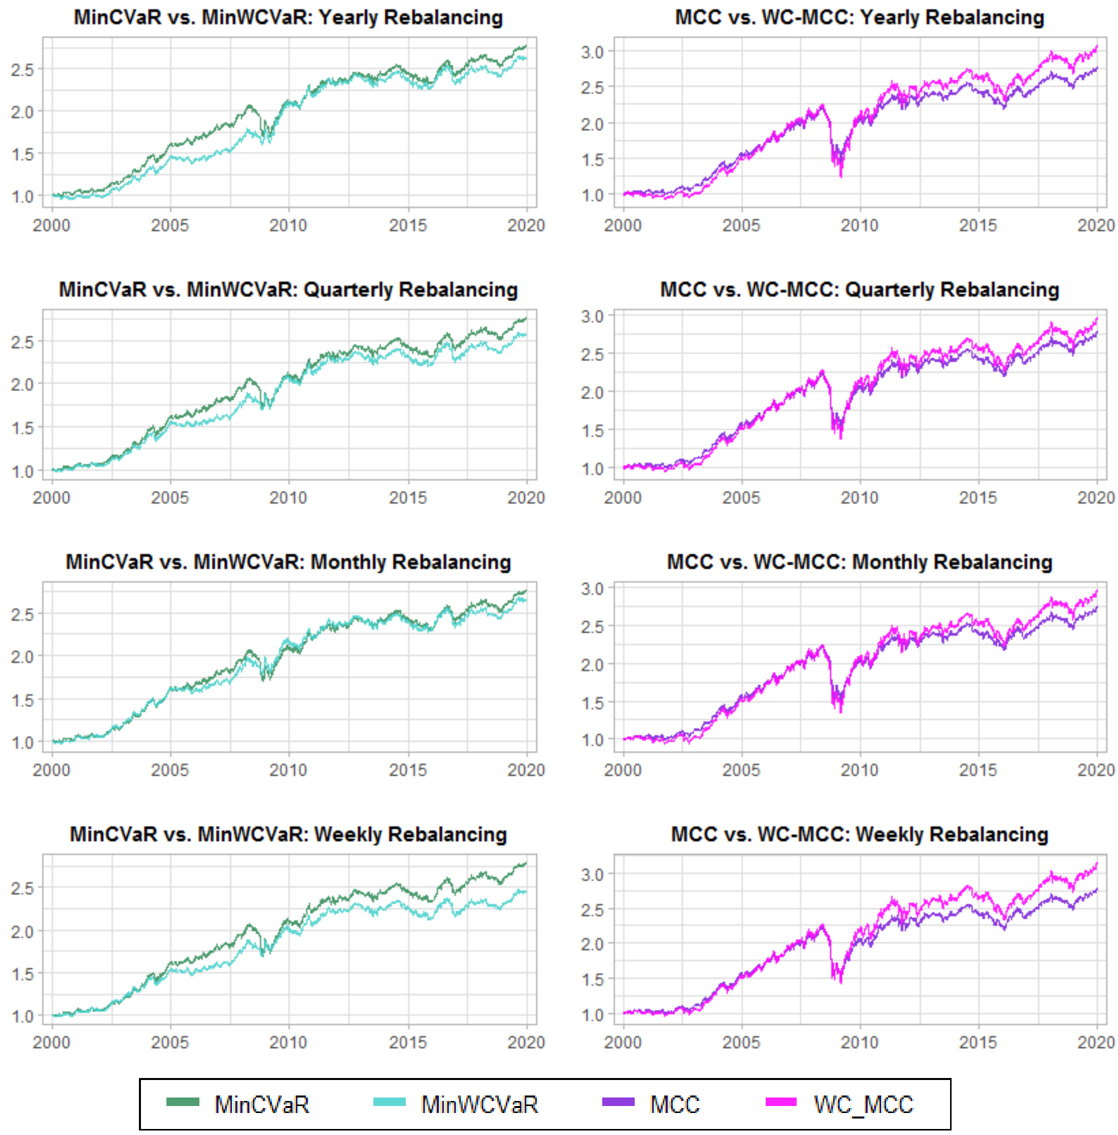

Figure 3: The Effect of Simulating Data for MinCVaR and MCC Portfolio Optimizations

The Worst-Case MinCVaR [MinWCVaR] approach produces a very conservative asset allocation that protects the investor from the worst possible scenario. It can be deduced from table 3 that the MinWCVaR portfolio achieves the lowest CVaRs and maximum drawdowns among all strategies, making it a reasonable choice for highly risk-averse investors. Obviously, this safety component comes at the cost of smaller cumulative returns as can be seen for all rebalancing frequencies in figure 3. Moreover, figures A.12 - A.15 indicate that the MinWCVaR portfolio produces very concentrated asset allocations that lead to the same diversification problem as for MinVar, MinSemiVar and MinCVaR. The drawback of MinCVaR and MinWCVaR portfolios is that the concentrated portfolio weights translate one-to-one to highly concentrated CVaR contributions (Boudt, Carl and Peterson, 2013, p. 46). This means that the entire CVaR risk is driven almost exclusively by only one asset, which is not what a highly risk-averse investor might be interested in. A CVaR contribution analysis is carried out on p. 82 and it outlines that the full investment constraint is responsible for the CVaR risk concentration of MinCVaR and MinWCVaR portfolios. It is important to realize that this thesis uses various asset classes to construct the portfolio, which leads the conservative optimization techniques to produce concentrated portfolios. In an investment universe that incorporates numerous defensive investments (e.g. bond portfolio), this portfolio concentration drawback might disappear.

The Worst-Case MCC portfolio [WC-MCC] produces superior returns compared to the MCC portfolio for all rebalancing frequencies. In fact, the simulation-based WC-MCC portfolio performs slightly worse in the beginning of the sample period and then starts to pick up after the financial crisis. One needs to understand that the modified CVaR estimator takes into account the skewness and kurtosis of the financial return data (see 4.7.1). By simulating worst-case data, this is directly incorporated in the mCVaR estimate, which is ultimately used in the portfolio optimization. Since WC-MCC achieves clear advantages over MCC after the global financial crisis, this might indicate that the WC-MCC portfolio requires a sufficiently long data history which involves market downturns in order to outperform the classical MCC portfolio. Unfortunately, these indications remain descriptive. An extensive simulation study is required in order to investigate the effect of the number of simulations, the random seed, the random copula weights as well as the confidence level on such multivariate worst-case mixture copula approaches. Nevertheless, the WC-MCC portfolio remains a very interesting and promising portfolio construction technique. It might be suited for retail investors who look for a long-term investment strategy that combines the upside potential of an equally-weighted portfolio with the downside limitation of the MinCVaR portfolio construction technique.

6 Conclusion

In light of recent developments in the financial industry and technology in general, retail investors are in an unprecedented position to invest on their own. The goal of this thesis was to implement and investigate various risk-based portfolio construction techniques that make use of mathematics and statistics rather than rules of thumb for investing. Over the course of the past 20 years, a retail investor would have been able to triple the invested money. However, to what extent the money would have actually grown and how many sleepless nights the investor would have had, crucially depends on the employed portfolio construction technique.

Using extensive out-of-sample backtestings for yearly, quarterly, monthly and weekly rebalancing, I find that the naive equal weight portfolio achieves significant returns while being subject to major losses during market downturns. In order to improve the portfolio's risk characteristics, the retail investor is advised to employ more complex approaches to portfolio construction. I find that all purely risk-based portfolio construction techniques which minimize a risk measure (variance, semivariance and CVaR) effectively reduce the portfolio risk. However, they all produce concentrated portfolios that are similar to each other. This finding is explained by the fact that the retail investor's portfolio consists of only six index ETFs that all cover different asset classes such that the optimizations tilt the portfolio towards the safest constituent. Additionally, I find that more frequent rebalancing does not substantially improve both risk and return characteristics for those portfolio construction techniques such that the investor can save the trading cost and effort by rebalancing the portfolio on a yearly basis. More recent developments in portfolio construction and optimization propose the use of techniques that diversify the investor's portfolio with respect to the assets' risk contributions. The minimum CVaR concentration [MCC] portfolio uses the modified CVaR and employs derivative-free global optimizers to produce an asset allocation that is well-diversified with respect to the assets' CVaR contributions. It resolves the issue of portfolio concentration and inherits the upside potential of an equally-weighted portfolio while limiting the downside risk. This makes it a highly attractive investment strategy that excels during bull markets (Boudt, Carl & Peterson, 2013, p. 60).

Besides performing the portfolio optimizations based on historical data, the effect of using simulated data is investigated for CVaR-based portfolio construction techniques. More precisely, a novel worst-case multivariate mixture copula approach is implemented and examined for MinCVaR and MCC portfolios. As for the MinCVaR portfolio, the worst-case mixture copula approach leads to an even more conservative asset allocation since the optimization technique attempts to protect the investor from the worst possible case – as suggested by its name. It effectively manages to reduce the portfolio risk measured by CVaR and maximum drawdown. However, these advantages come at a cost of lower returns and the same portfolio concentration issue exists as for the Minimum Variance, Minimum Semivariance and Minimum CVaR portfolios. Interestingly, the worst-case approach leads to highly promising results for the Worst-Case MCC portfolio: It produces superior returns compared to the classical MCC portfolio, which is purely based on historical data. This improvement becomes evident and gets stronger after the global financial crisis when the optimization procedure starts to take data from excessive market downturns into account. A conclusive statement, however, is not possible since the results of the WC-MCC portfolio are based on simulations that themselves depend on the number of simulations, the random seed, the amount of random copula weights as well as the CVaR confidence level. Therefore, an extensive simulation study is required to provide indications for why the WC-MCC portfolio dominates the classical MCC portfolio optimization technique in the present context.

In conclusion, this thesis shows that today's retail investors have access to all the pieces required to implement sophisticated portfolio construction techniques on their home computers, which might encourage some of them to take charge of their own financial destiny.

References

Acerbi, C. (2002). Spectral Measures of Risk: A Coherent Representation of Subjective Risk Aversion. *Journal of Banking and Finance. 26(7),* 1505-1518. Retrieved January 30, 2020 from `https://www.sciencedirect.com/`

Acerbi, C. (2007). Coherent measures of risk in everyday market practice. *Quantitative Finance. 7(4),* 359-364. doi: 10.1080/14697680701461590

Antoniewicz, R., & Heinrichs, J. (2014). Understanding Exchange-Traded Funds: How ETFs Work. *ICI Research Perspective. 20(5),* 1-39. doi: 10.2139/ssrn.2523540

Ardia, D., Boudt, K., Carl, P., Mullen, K. M. & Peterson, B. G. (2010). Differential Evolution with DEoptim: An Application to Non-Convex Portfolio Optimization. *The R Journal. 3(1), 27-34.* Retrieved January 30, 2020 from `https://ssrn.com/abstract=1584905`

Artzner, P., Delbaen, F., Eber, J. M. & Heath, D. (1999). Coherent Measures of Risk. *Mathematical Finance. 9(3),* 203-228. doi: 10.1111/1467-9965.00068

Athayde, G. (2005). The mean-downside risk portfolio frontier: a non-parametric approach. In S. Satchell & Scowcroft A. (Editors), *Advances in Portfolio Construction and Implementation* (pp. 290-309). Oxford: Elsevier.

Ben-David, I., Franzoni, F., & Moussawi, R. (2017). Exchange Traded Funds (ETFs). *Annual Review of Financial Economics. 9*, 169-189. doi: 10.1146/annurev-financial-110716-032538

Bhattacharya, U., Loos, B., Meyer, S., & Hackethal, A. (2017). Abusing ETFs. *Review of Finance. 21(3),* 1217-1250. doi: 10.1093/rof/rfw041

BlackRock. (2017). *A Primer on ETF Primary Trading and the Role of Authorized Participants.* Retrieved January 30, 2020 from `https://www.blackrock.com/corporate/library?keyword=Authorized+Participant`

Blocher, J., & Whaley, R. E. (2016). *Exchange-traded Funds and Securities Lending.* Retrieved January 30, 2020 from `https://ssrn.com/abstract=2474904`

Boudt, K., Carl, P. & Peterson, B. G. (2013). Asset Allocation with Conditional Value-at-Risk Budgets. *The Journal of Risk. 15(3), 39-68.* doi: 10.21314/JOR.2013.258

Boudt, K., Peterson, B. G. & Croux, C. (2008). Estimation and decomposition of downside risk for portfolios with non-normal returns. *The Journal of Risk. 11(2), 79-103.* doi: 10.21314/JOR.2008.188

Chincarini, L. B. & Kim, D. (2006). *Quantitative Equity Portfolio Management. An Active Approach to Portfolio Construction and Management*. New York: McGraw-Hill Companies, Inc.

Cont, R. (2005). *Volatility Clustering in Financial Markets: Empirical Facts and Agent-Based Models*. Retrieved January 30, 2020 from https://ssrn.com/abstract=1411462

Cont, R., Deguest, R. & Scandolo, G. (2008). *Robustness and Sensitivity Analysis of Risk Measurement Procedures*. Retrieved January 30, 2020 from https://ssrn.com/abstract=1086698

Dalbar. (2018). *Quantitative Analysis of Investor Behavior. 2018 QAIB Report.* Retrieved January 30, 2020 from https://content.swanglobalinvestments.com/hubfs/ThirdPartyDocuments/Dalbar2018QAIBReport-QuantitativeAnalysisofInvestorBehavior-AdvisorEdition.pdf

DeMiguel, V., Garlappi, L. & Uppal, R. (2007). Optimal Versus Naive Diversification: How Inefficient is the 1/N Portfolio Strategy? *The Review of Financial Studies. 22(5),* 1915-1953. doi: 10.1093/rfs/hhm075

Dichtl, H., Drobetz, W. & Wambach, M. (2014). *Where is the Value Added of Rebalancing? A Systematic Comparison of Alternative Rebalancing Strategies*. Retrieved January 30, 2020 from https://ssrn.com/abstract=2139915

Estrada, J. (2008). Mean-Semivariance Optimization: A Heuristic Approach. *Journal of Applied Finance. 18(1),* 57-72. Retrieved January 30, 2020 from https://ssrn.com/abstract=2698700

Evans, A. D. (2009). A Requiem for the Retail Investor? *Virginia Law Review, 95(4),* 1105-1130. Retrieved January 30, 2020 from https://repository.law.umich.edu/articles/118

Ferri, R. A. (2008). *The ETF Book. All You Need To Know About Exchange-Traded Funds*. New Jersey: John Wiley and Sons Inc.

Fitzgerald, M. (2019). The end of commissions for trading is near as TD Ameritrade cuts to zero, matching Schwab. *CNBC*. Retrieved January 30, 2020 from https://www.cnbc.com/

Fry, D. (2008). *Create Your Own ETF Hedge Fund. A Do-It-Yourself ETF Strategy for Private Wealth Management*. New Jersey: John Wiley and Sons Inc.

Hirshleifer, D. A. (2001). Investor psychology and asset pricing. *Journal of Finance, 56(4),* 1533–1597. doi: 10.2139/ssrn.265132

Hofert, M., Kojadinovic, I., Mächler, M. & Yan, J. (2018). *Elements of Copula Modeling with R*. Cham: Springer.

IOSCO. (2019). *The Application of Behavioural Insights to Retail Investor Protection. Final Report.* Retrieved January 30, 2020 from https://www.iosco.org/library/pubdocs/pdf/IOSCOPD626.pdf

Jagannathan, R. & Ma, T. (2003). Risk Reduction in Large Portfolios: Why Imposing the Wrong Constraints Helps. *Journal of Finance. 58(4),* 1651-1684. doi: 10.1111/1540-6261.00580

Kakouris, I. & Rustem, B. (2014). Robust portfolio optimization with copulas. *Journal of Operational Research. 235,* 28-37. doi: 10.1016/j.ejor.2013.12.022

Kemp, M. H. (2011). *Extreme Events. Robust Portfolio Construction in the Presence of Fat Tails*. West Sussex: John Wiley and Sons Ltd.

Liagkouras, K. & Metaxiotis, K. (2013). The Constrained Mean-Semivariance Portfolio Optimization Problem with the Support of a Novel Multiobjective Evolutionary Algorithm. *Journal of Software Engineering and Applications. 6, 22-29.* doi: 10.4236/jsea.2013.67B005

Lim, A. E., Shanthikumar, J. G. & Vahn G. (2011). Conditional value-at-risk in portfolio optimization: Coherent but fragile. *Operations Research Letters. 39(3),* 163-171. doi: 10.1016/j.orl.2011.03.004

Marshall, B. R., Nguyen, N. H., & Visaltanachoti, N. (2017). *Do Liquidity Proxies Measure Liquidity Accurately in ETFs?* Retrieved January 30, 2020 from `https://ssrn.com/abstract=2701608`

Martin, R. D., Clark, A. & Green, C. G. (2010). Robust Portfolio Construction. In J. B. Guerard, Jr. (Editor), *Handbook of Portfolio Construction - Contemporary Applications of Markowitz Techniques* (pp. 337-380). New York: Springer.

Mauzy, S. (2015). The Overlooked Advantages of the Individual Investor. *Wyatt Investment Research*. Retrieved January 30, 2020 from `https://www.wyattresearch.com/`

Melchiori, M. R. (2006). *Tools for sampling Multivariate Archimedean Copulas*. Retrieved January 30, 2020 from `https://ssrn.com/abstract=1124682`

Meziani, A. S. (2016). *Exchange-Traded Funds. Investment Practices and Tactical Approaches*. London: Macmillan Publishers Ltd.

Montier, J. (2007). *Behavioural Investing. A practitioner's guide to applying behavioural finance*. West Sussex: John Wiley and Sons Ltd.

Neves, M. E., Fernandes, C. M., & Martins, P. C. (2019). Are ETFs good vehicles for diversification? New evidence for critical investment periods. *Borsa Istanbul Review. 19(2),* 149-157. doi: 10.1016/j.bir.2019.01.002

Pagano, M., Sánchez Serrano, A., & Zechner, J. (2019). Can ETFs contribute to systemic risk? *European Systemic Risk Board, Reports of the Advisory Scientific Committee, No 9*. Retrieved January 30, 2020 from `https://www.esrb.europa.eu/`

Pan, K., & Zeng, Y. (2019). *ETF Arbitrage Under Liquidity Mismatch.* Retrieved January 30, 2020 from `https://ssrn.com/abstract=2895478`

Petajisto, A. (2017). Inefficiencies in the Pricing of Exchange-Traded Funds. *Financial Analysts Journal. 73(1),* 24-54. doi: 10.2469/faj.v73.n1.7

Pfaff, B. (2013). *Financial Risk Modelling and Portfolio Optimization with R*. West Sussex: John Wiley and Sons Ltd.

Rockafellar, R. T. & Uryasev, S. (2000). Optimization of Conditional Value-at-Risk. *The Journal of Risk. 2,* 21-41. Retrieved January 30, 2020 from `https://www.ise.ufl.edu/`

Sabino da Silva, F. A., Ziegelmann, F. A. & Tessari, C. (2017). *Robust Portfolio Optimization with Multivariate Copulas: A Worst-Case CVaR Approach*. Retrieved January 30, 2020 from `https://ssrn.com/abstract=3076283`

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

Zhu, S. & Fukushima, M. (2009). Worst-Case Conditional Value-at-Risk with Application to Robust Portfolio Management. *Operations Research, 57(5),* 1155-1168. doi: 10.1287/opre.1080.0684

Appendix

1.) The Indexes' Price Evolutions

The following plots provide an overview of the price evolution of all selected indexes over the entire sample period. The underlying price data are daily adjusted closing prices of the respective indexes. During the financial crisis around 2008 - 2009, sharp drops in prices can be found for all indexes except for the World Bond Index. General upward trends during the sample period can be identified for all index price series, the one of the Commodity Index being the weakest overall. The different scales of the y-axes show that asset prices are not scale free, which is one of the reasons why researchers usually work with financial return data and not with raw price data.

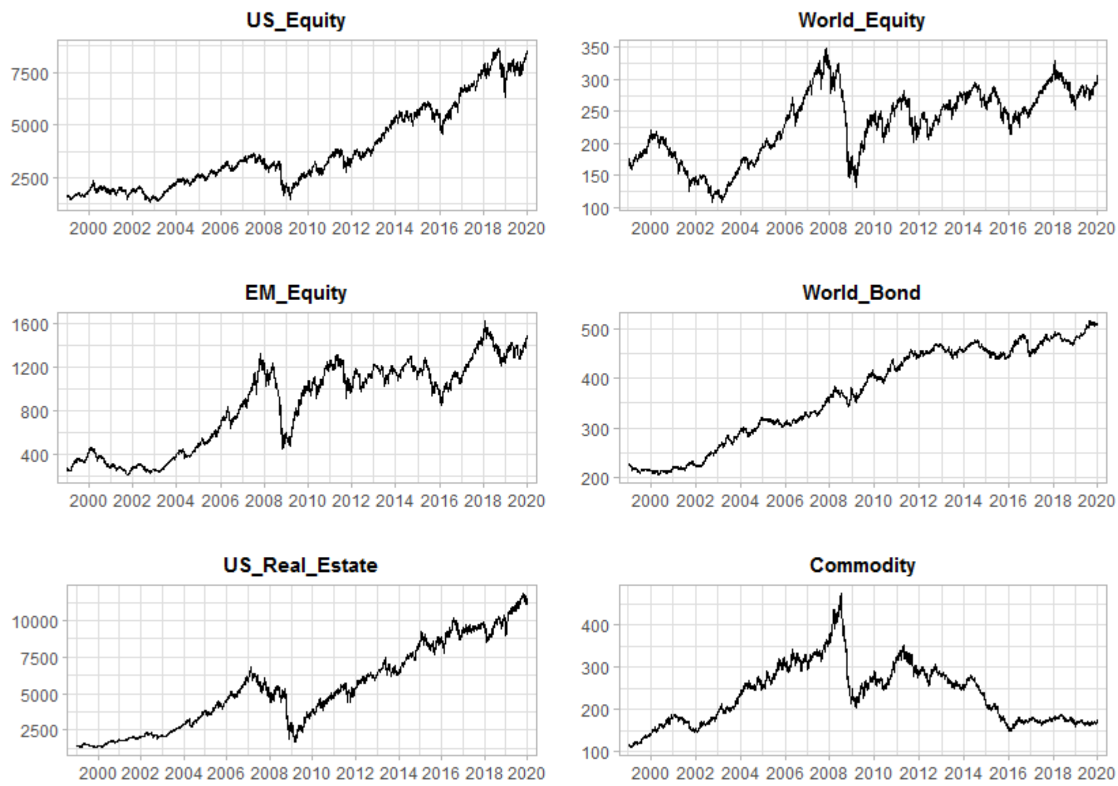

Figure A.1: Price Evolution of the Indexes

2.) The Indexes' Simple Return Evolutions

The subsequent plots provide a glimpse into the stylized fact that daily financial return series fluctuate around their means, which are very close to zero. The most extreme returns (positive and negative) seem to have occurred during the global financial crisis 2008 - 2009 for all indexes. The daily return series of the US Real Estate Index exhibits the most extreme behavior because this asset class was directly affected during the global financial crisis (US housing bubble burst). The Commodity Index' return series already gives the indication of a high volatility, which will be further investigated in the next plots. Overall, the World Bond Index produces the most moderate returns over time.

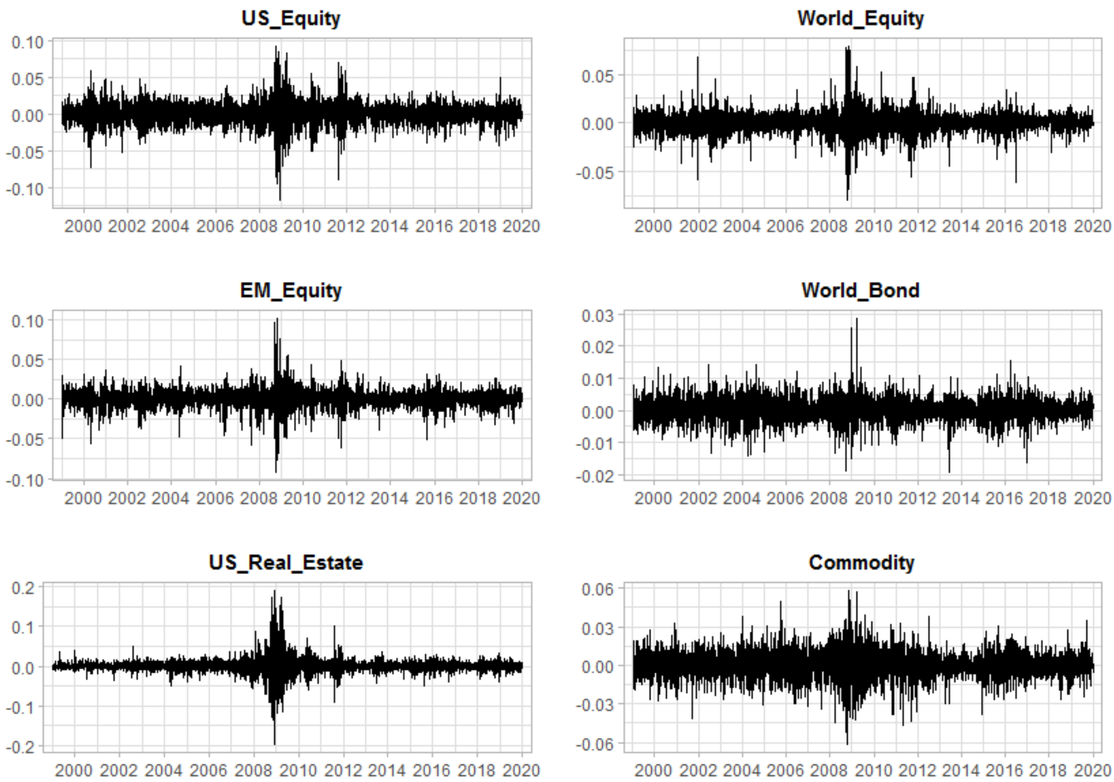

Figure A.2: Daily Simple Return Series of the Indexes

3.) The Indexes' Squared Simple Return Evolutions

The following plots provide an overview of the indexes' squared simple returns, which are used as a proxy for volatility. The volatilities seem to behave similarly across the indexes. It is apparent that the volatility is not constant over time and occurs in so-called volatility clusters. In the period from 2000 until 2007 a relatively low fluctuation in daily returns can be observed for almost all indexes, the Commodity Index being the exception. Afterwards, the daily volatilities of all indexes have peaked around the global financial crisis 2008 - 2009. Again, the US Real Estate Index stands out with its extreme volatility during said crisis. In the subsequent years, the volatility declined to pre-crisis levels until the end of 2011, where another volatility peak occurred for almost all indexes. Overall, the Commodity Index seems to be frequently volatile. The World Bond Index exhibits the lowest volatilities among all indexes.

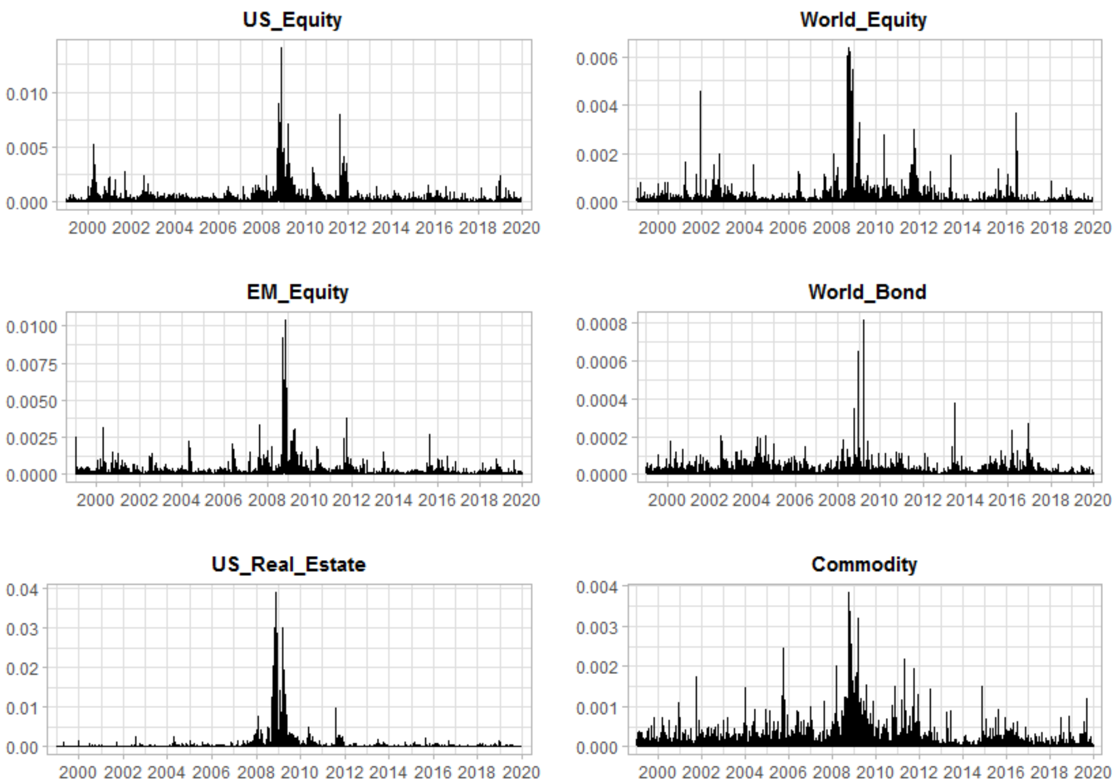

Figure A.3: Daily Squared Simple Return Series of the Indexes

4.) Normality Tests of the Indexes' Univariate Simple Return Series

In order to investigate the potential normality of all the index return series, a qualitative as well as quantitative approach is performed. First, QQ-plots are investigated for which the empirical quantiles are plotted against those of a normal distribution with estimated mean and variance. This provides a first impression of the respective return distributions in comparison with a normal distribution. Moreover, those qualitative results are double-checked quantitatively with Jarque-Bera normality tests. Both approaches provide clear indications of non-normality for all the index return series.

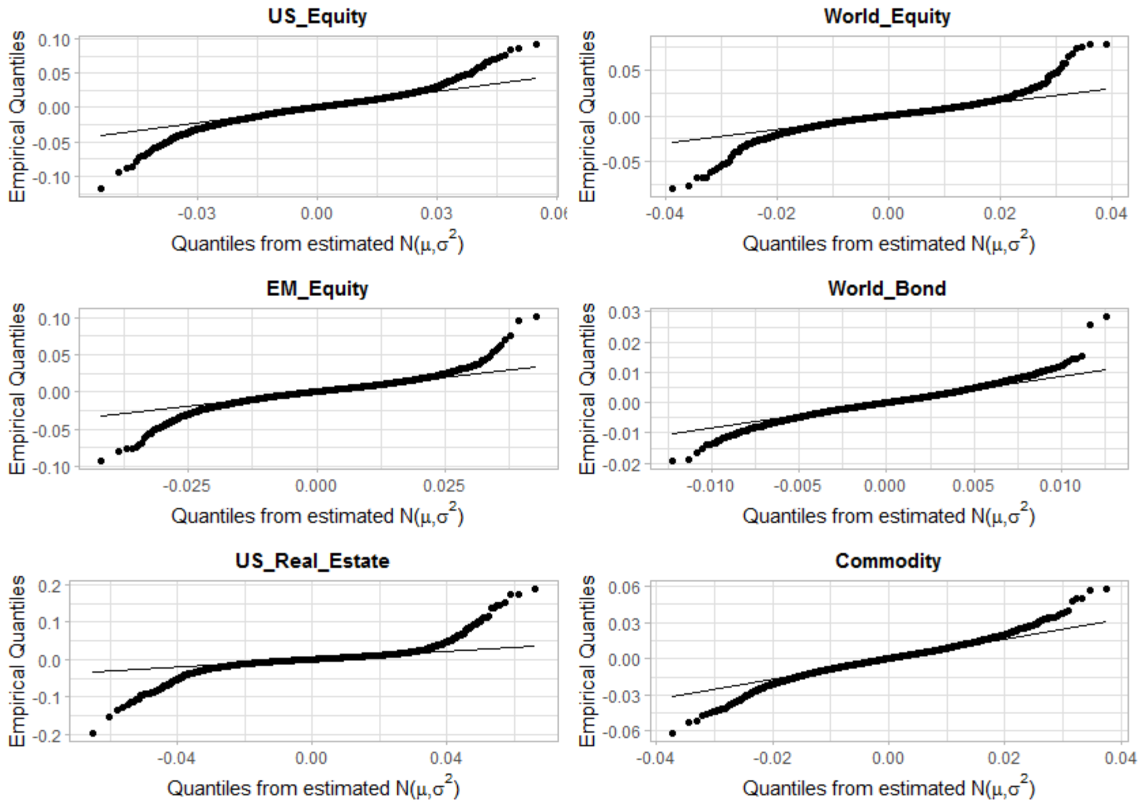

Figure A.4: QQ-Plots – Empirical Quantiles vs. Estimated Normal

As is often the case with financial return data, the QQ-plots of all indexes get a very good resolution in the center of the distribution (where the probability density is high). In the tails, however, the resolution is poor. It is obvious that the non-normality of all index return series stems from the upper and lower tails of the respective empirical return distributions. In reality, extreme positive and negative returns seem to occur more often than extreme realizations from a normal distribution with estimated parameters. This is a materialization of the stylized fact that financial return series exhibit leptokurtosis (fat tails).

According to Tsay (2010, pp. 9 - 10), skewness measures the symmetry of a random variable X while kurtosis measures the tail behavior of X. The third and fourth central (normalized) moments of the random variable X are defined as:

$$S(x) = E\left[\frac{(X-\mu_x)^3}{\sigma_x^3}\right] \quad \text{and} \quad K(x) = E\left[\frac{(X-\mu_x)^4}{\sigma_x^4}\right],$$

where μ_x denotes the mean of X and σ_x is the standard deviation of X. See chapter 4.1 for an equivalent definition in a portfolio context. Under the null hypothesis of normality, the sample skewness ($\hat{S}(x)$) and sample excess kurtosis ($\hat{K}(x)-3$) are asymptotically normal with zero mean and variances $6/T$ and $24/T$, respectively. Exploiting this feature, Jarque and Bera (1987) proposed the following test statistic to test the normality of asset returns:

$$JB_i = \frac{\hat{S}^2(R_i)}{6/T} + \frac{[\hat{K}(R_i)-3]^2}{24/T} \quad \text{for} \quad i = 1, \ldots, N,$$

where R_i denotes the return series of asset i. The test statistic JB_i is asymptotically a chi-squared random variable with two degrees of freedom. The null hypothesis of normality is rejected if the p-value of the JB_i-statistic is less than the selected significance level. The Jarque-Bera tests confirm the suspicions from the QQ-plots with p-values very close to zero. Table A.1 displays the test results and indicates a statistically significant non-normality for all index return series.

	US Eq	World exUS Eq	EM Eq	World Bond	US RE	Commodity
JB	4’951.3	13’331.0	13’435.5	2’125.9	111’392.0	1’527.1
p-value	0.00	0.00	0.00	0.00	0.00	0.00

Table A.1: Jarque-Bera Tests

5.) Autocorrelations of the Indexes' Simple Return Series

In a time series setting, the concept of correlation is generalized to autocorrelation if the linear dependence between a weakly stationary return series $R_{i,t}$ and its past values $R_{i,t-k}$ is of interest (pp. 30 - 32). The correlation coefficient between $R_{i,t}$ and $R_{i,t-k}$ is called the lag-k autocorrelation of $R_{i,t}$ and is commonly denominated by $\rho_{i,k}$. Under the assumption of weak stationarity, $\rho_{i,k}$ is only a function of the lag k:

$$\rho_{i,k} = \frac{Cov(R_{i,t}, R_{i,t-k})}{\sqrt{Var(R_{i,t})Var(R_{i,t-k})}} = \frac{Cov(R_{i,t}, R_{i,t-k})}{Var(R_{i,t})} = \frac{\gamma_{i,k}}{\gamma_{i,0}} \quad \text{for} \quad i = 1, \ldots, N,$$

where the property $Var(R_{i,t}) = Var(R_{i,t-k})$ of a weakly stationary time series is used. In the finance literature, it is common to assume weak stationarity for asset return series (p. 30). The return series $\{R_{i,t}\}$ is weakly stationary if

(i) $E[R_{i,t}^2] < \infty, \ \forall t \in \mathbb{Z}$ (ii) $E[R_{i,t}] = \mu_i, \ \forall t \in \mathbb{Z}$ (iii) $Cov(R_{i,t}, R_{i,t-k}) = \gamma_{i,k}, \ \forall t, k \in \mathbb{Z}$

for $i = 1, \ldots N$. Weak stationarity is based only on the first two moments of the stochastic process $\{R_{i,t}\}$, which is the reason why it is also referred to as second-order stationarity. For a given sample of returns $\{R_{i,t}\}_{t=1}^{T}$ with sample mean $\mu_i = \sum_{t=1}^{T} \frac{R_{i,t}}{T}$, the lag-$k$ sample autocorrelation of $R_{i,t}$ is defined as

$$\hat{\rho}_{i,k} = \frac{\sum_{t=k+1}^{T}(R_{i,t} - \mu_i)(R_{i,t-k} - \mu_i)}{\sum_{t=1}^{T}(R_{i,t} - \mu_i)^2}, \quad \text{for} \quad 0 \leq k < T-1 \quad \text{and} \quad i = 1, \ldots, N.$$

It is essential to realize that $\hat{\rho}_{i,k}$ is a biased estimator of $\rho_{i,k}$ in the case of finite samples. The bias equals $1/T$, which can be substantial when the sample size T is low. In most financial applications, however, the sample size T is relatively large such that the bias is negligible.

The autocorrelation function (ACF) plots below provide insights into the autocorrelation of the respective return series. For each lag up to $k = 20$, the autocorrelation is estimated. The horizontal blue lines correspond to a 95% confidence interval. The ACF plots hint at a slight autocorrelation of first order for all index return series except for the Commodity Index. There is almost no serial correlation of the simple returns for all indexes, which is in line with the corresponding stylized fact of financial return data. As will be shown in the next plots, all return series exhibit a dependence despite the lack of autocorrelation.

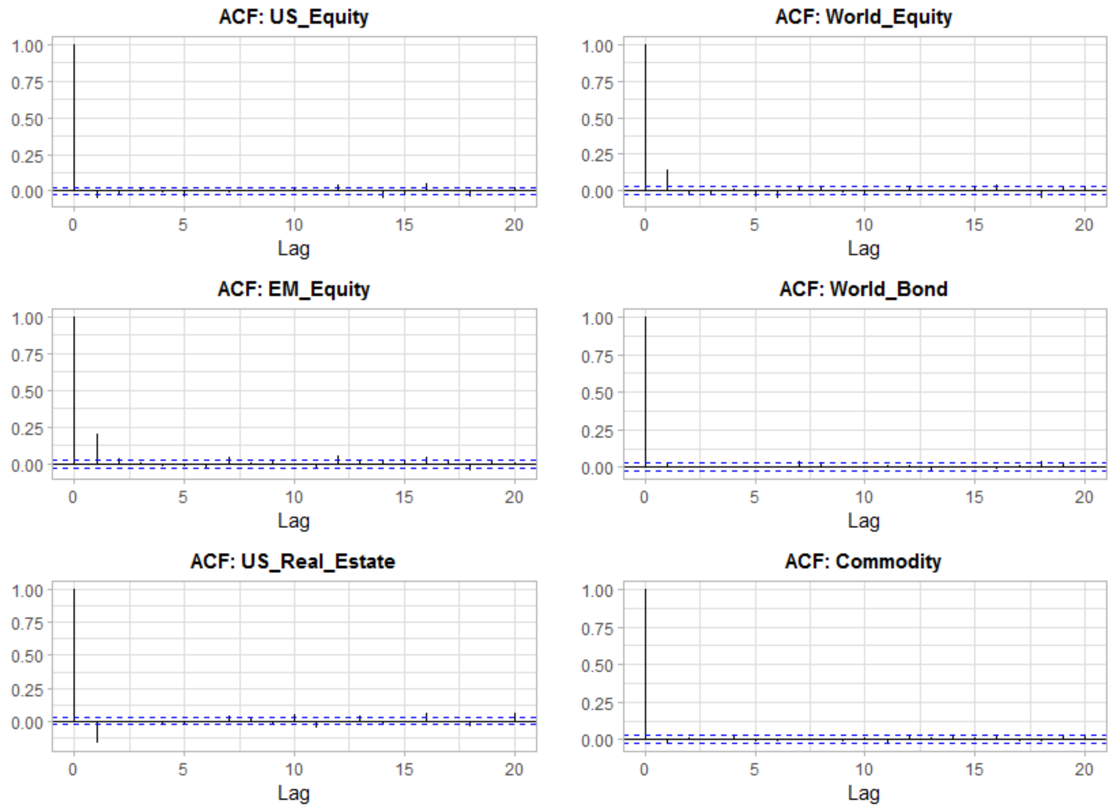

Figure A.5: ACF Plots of the Indexes' Simple Returns

6.) Autocorrelations of the Indexes' Squared Simple Return Series

In contrast to the previously depicted ACF plots of simple returns, figure A.6 shows that the squared returns of all indexes are clearly autocorrelated. The only exception is the squared return series of the World Bond Index, which in fact does exhibit significant autocorrelations for several lags but in a much smaller extent compared to the other indexes. A similar result can be obtained when absolute returns are considered. Clearly, the autocorrelations of the indexes' squared returns are significantly different from zero at a 95% confidence level for several lags and they taper off only slowly in a hyperbolic fashion, indicating a strong time dependence in volatility. This is an important feature in financial market modeling, which has led to the emergence of so-called conditional heteroskedasticity models (ARCH & GARCH Models).

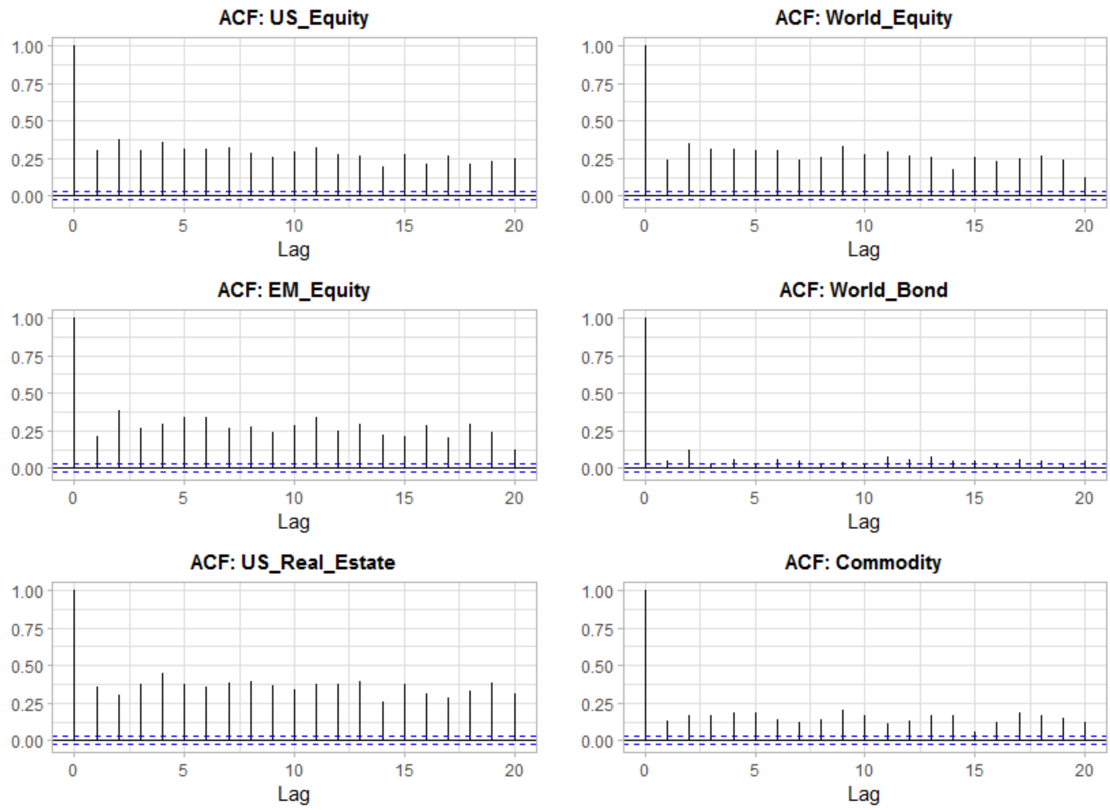

Figure A.6: ACF Plots of the Indexes' Squared Returns

7.) Cross-Correlations of the Selected Indexes' Simple Return Series

Because the indexes' return series are analyzed in a portfolio context, cross-correlations between the assets must be considered as well. In the general case of N assets, there are $N*(N-1)/2$ unique asset pairs to investigate. In order to avoid redundancy, only a selection of asset pairs is disclosed here. The plots of the cross-correlations are fairly similar for all possible asset pairs and the general insight is similar to the univariate case: the cross-correlations of simple returns are less pronounced and not significant compared to those of squared returns. Hence, also in the multivariate case there is no significant autocorrelation between the indexes' simple returns (except for 1 lag). The ACF plots of the squared returns, however, show that the return series are dependent despite the lack of significant autocorrelation in simple returns.

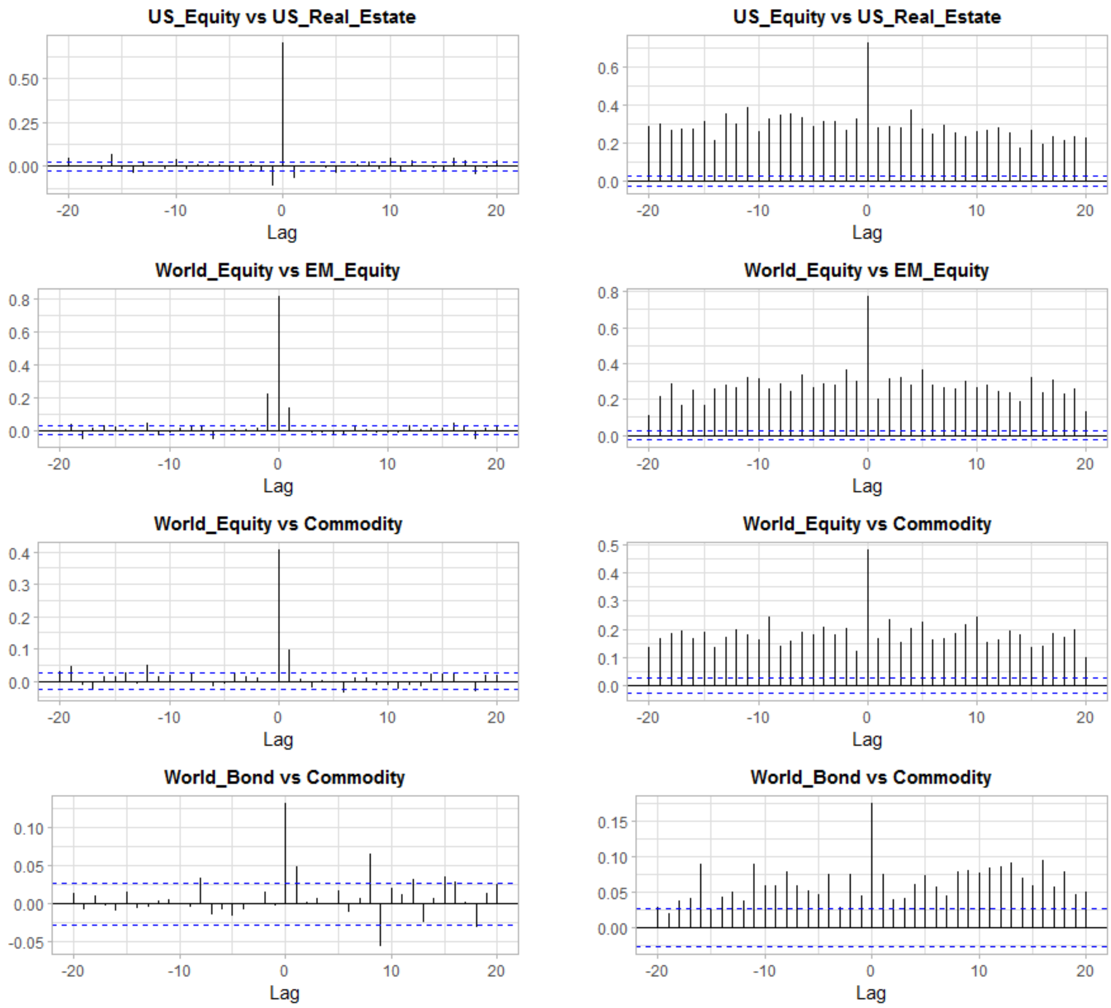

Figure A.7: CCF Plots of Selected Index Return Series (Simple & Squared)

Finally, the time-varying correlations between the asset pairs are investigated. Figure A.8 depicts correlations based upon a moving window of 250 observations. Again, there is a total of $N*(N-1)/2$ plots, from which only a selection is displayed (same asset pairs as above). It is crucial to realize that the correlations between assets are dynamic and change over time. The plots below portray the dynamic correlations for selected asset pairs. The correlations seem to have increased prior to the financial crisis in 2008 for almost all asset pairs considered. The exception is the asset pair of World Bond & Commodity, which shows decreasing correlations prior to the crisis. Nevertheless, the correlation of this asset pair has increased subsequently after the crisis in concert with the other asset pairs. The plots furthermore provide indication of diversification effects across asset classes: The rolling correlations of asset pairs involving different asset classes tend to exhibit lower correlations than pairs of the same asset class. More specifically, the asset pair World Equity & EM Equity exhibits the largest correlation among all the asset pairs in figure A.8. This could be the result of aforementioned diversification effects. The exception to this observation, however, is the asset pair US Equity & US Real Estate, which exhibits very high correlations during and after the financial crisis even though two different asset classes are involved. Intuitively, the reason for the high correlation between US Equity and US Real Estate can be explained by the strong exposure of both asset classes to the same geography. This stresses the importance of geographic diversification in portfolio construction.

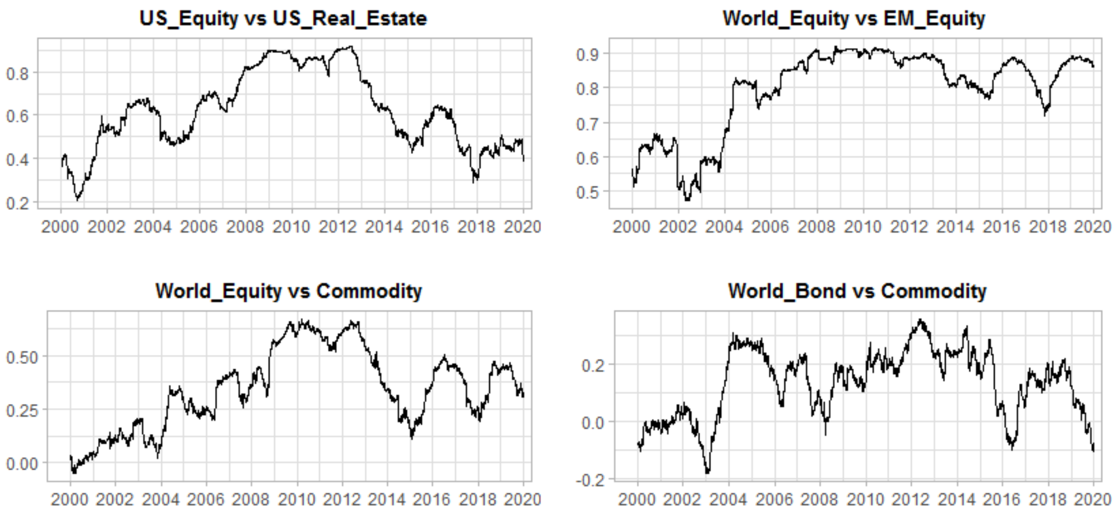

Figure A.8: Selected Rolling Correlations (T = 250)

In order to provide an insight into all possible asset pairs, Table A.2 shows the rolling correlation ranges for each asset pair as well as the most recent correlation based on the last 250 observations (in parentheses; as of December 31, 2019). The lower and upper triangles of the correlation matrix show that there are 15 unique asset pairs when 6 assets are considered. The largest interval is detected for the asset pair World Bond & World Equity: the correlation ranges between −0.39 and 0.56 in the sample period. The smallest correlation interval is found in the asset pair US Real Estate & EM Equity with correlations between 0.04 and 0.52. The lowest correlations can be found in all asset pairs where the World Bond Index is involved, indicating its strong diversification benefits. Currently, the lowest correlation can be detected for the asset pair World Bond & US Equity (−0.39, based on past 250 observations). The highest correlation, on the other hand, is currently between Emerging Markets Equity & World Equity with 0.86. This supports the previous argument that exposure to the same asset class can lead to high correlations for such an asset pair.

	US Equity	World Equity	EM Equity	World Bond	US Real Estate	Commodity
US Equity	1	[0.16; 0.69] (0.55)	[0.01; 0.60] (0.48)	[−0.44; 0.28] (−0.39)	[0.20; 0.92] (0.41)	[−0.11; 0.58] (0.37)
World Equity	[0.16; 0.69] (0.55)	1	[0.47; 0.92] (0.86)	[−0.39; 0.56] (−0.12)	[0.07; 0.62] (0.19)	[−0.05; 0.67] (0.33)
EM Equity	[0.01; 0.60] (0.48)	[0.47; 0.92] (0.86)	1	[−0.27; 0.44] (−0.15)	[0.04; 0.52] (0.08)	[−0.07; 0.61] (0.30)
World Bond	[−0.44; 0.28] (−0.39)	[−0.39; 0.56] (−0.12)	[−0.27; 0.44] (−0.15)	1	[−0.42; 0.31] (0.10)	[−0.18; 0.36] (−0.08)
US Real Estate	[0.20; 0.92] (0.41)	[0.07; 0.62] (0.19)	[0.04; 0.52] (0.08)	[−0.42; 0.31] (0.10)	1	[−0.26; 0.51] (0.07)
Commodity	[−0.11; 0.58] (0.37)	[−0.05; 0.67] (0.33)	[−0.07; 0.61] (0.30)	[−0.18; 0.36] (−0.08)	[−0.26; 0.51] (0.07)	1

Table A.2: Rolling Correlation Ranges (T = 250)

8.) Axiomatic Definition of Risk Measures

Artzner et al. (1999) introduced and justified a set of four desirable properties for risk measures. They advocated the use of coherent risk measures that fulfill all of their framework's mathematical axioms. In order to understand the tradeoff between coherency and estimation robustness outlined on p. 79, a basic knowledge of these axioms is required. Let $(\Omega, \mathscr{F}, \mathbb{P})$ be a given probability space and L^0 be the space of all random variables. Let X be a random variable that represents the profit and loss of a portfolio over a specified horizon with $X \in L \subset L^0(\Omega, \mathscr{F}, \mathbb{P})$, where negative values of X correspond to losses. A risk measure on L is the real-valued map $\rho : L \longrightarrow \mathbb{R}$ assigning a number to each profit and loss $X \in L$ representing its degree of riskiness (Cont, Deguest & Scandolo, 2008, p. 7). In order to obtain a coherent risk measure, the following properties must be fulfilled:

1. *Positive Homogeneity*: $\rho(\lambda X) = \lambda \rho(X)$ for any $\lambda \geq 0$. Increasing the size of a portfolio by a factor λ must scale its risk by the exact same factor.
2. *Translation Equivariance*: $\rho(X + c) = \rho(X) - c$ for any $c \in \mathbb{R}$. Adding risk-free cash to a portfolio must decrease its risk by this amount.
3. *Monotonicity*: $\rho(X) \leq \rho(Y)$ for $X \geq Y$. If a risky position X provides larger payoff than position Y, then the risk of X must be less or equal the risk of Y.
4. *Sub-Additivity*: $\rho(X + Y) \leq \rho(X) + \rho(Y)$. Merging two positions cannot increase the total portfolio risk due to potential diversification effects.

In a portfolio optimization context, it is important to know that axioms 1 and 4 lead to convexity, which simplifies the optimization procedure immensely. Therefore, an elementary consequence of coherency is the convexity of risk surfaces (Acerbi, 2007, pp. 359 - 360). The Value-at-Risk lacks the Sub-Additivity property, which is the reason why it is not considered a coherent risk measure and is therefore not convex. The lack of Sub-Additivity implies that using VaR as risk measure could result in a lower total portfolio risk when the risky positions are split into single positions. Conversely, it can happen that the hedging benefit turns out to be negative, which is nonsensical from a risk-theoretical perspective. The VaR satisfies the Sub-Additivity property and achieves coherency only in the special case where the assets' returns are jointly normally distributed. The Conditional Value-at-Risk, on the other hand, always satisfies above axioms and is therefore a coherent risk measure, which is the reason why it became very popular in risk management in recent years. However, Cont, Deguest and Scandolo (2008) prove that CVaR's coherency comes at a cost of estimation robustness issues as will be discussed on p. 79 in the appendix.

9.) Quantile-Based Portfolio Optimization with Discontinuity: VaR and CVaR

In case of any uncertainties regarding the subsequent notations, reconsider the section 4.1. How to optimize with respect to CVaR is extensively explained in 4.6.1, where the crucial assumption of the cumulative distribution function of the loss $\Psi(\mathbf{w},\gamma)$ being everywhere continuous with respect to γ is imposed. If this is not the case, the definitions of VaR and CVaR must be adjusted. Again, let $f(\mathbf{w},\mathbf{R})$ denote the portfolio loss function, where $f(\mathbf{w},\mathbf{R}) = -[w_1R_1 + ... + w_NR_N] = -\mathbf{w}^T\mathbf{R}$. For any confidence level $\alpha \in (0,1)$ and any portfolio $\mathbf{w} \in \mathbf{W}$, the Value-at-Risk at level α associated with a portfolio $\mathbf{w} \in \mathbf{W}$ is defined as

$$VaR_\alpha(\mathbf{w}) = inf\{\gamma \in \mathbb{R} : \Psi(\mathbf{w},\gamma) \geq \alpha\} = F_L^{\leftarrow}(\alpha),$$

where the function $F_L^{\leftarrow}(\alpha)$ is the generalized inverse of the losses' cumulative distribution function $\Psi(\mathbf{w},\gamma) = \mathbb{P}[f(\mathbf{w},\mathbf{R}) \leq \gamma]$, allowing us to obtain the α-quantile (i.e. the VaR) of the loss distribution in case of discontinuities. The generalized inverse coincides with the usual inverse function when the cumulative distribution function is continuous. Therefore, it provides us with a broader applicability. Moreover, it is rich enough to also handle empirical distributions, which are highly discontinuous step-functions. To move on, the Conditional Value-at-Risk is defined as

$$\begin{aligned} CVaR_\alpha(\mathbf{w}) = \frac{1}{1-\alpha}\Bigg[&\mathbb{E}_{\mathbb{P}}[f(\mathbf{w},\mathbf{R})\ \mathbb{1}_{\{f(\mathbf{w},\mathbf{R})\geq VaR_\alpha(\mathbf{w})\}}] \\ &- VaR_\alpha(\mathbf{w})\Big(\mathbb{P}[f(\mathbf{w},\mathbf{R}) \geq VaR_\alpha(\mathbf{w})] - (1-\alpha)\Big)\Bigg]. \end{aligned}$$

The second part is responsible for CVaR's coherency under discontinuity. Note that this correction disappears under continuity because $\mathbb{P}[f(\mathbf{w},\mathbf{R}) \geq VaR_\alpha(w)] = 1-\alpha$. Equivalently, the Conditional Value-at-Risk can be defined in a more convenient way, which makes obvious that it can be interpreted as equally weighted average of quantiles in $[\alpha, 1]$:

$$CVaR_\alpha(\mathbf{w}) \ = \ \frac{1}{1-\alpha}\int_\alpha^1 F_L^{\leftarrow}(p)dp \ = \ \frac{1}{1-\alpha}\int_\alpha^1 VaR_q(\mathbf{w})\ dq.$$

Acerbi (2002) uses this representation and shows that the way how the tails' confidence level slices are weighted can be modified in order to model investor-specific risk aversions, opening an entire spectrum of coherent risk measures (so called spectral risk measures). This topic is very interesting when it comes to modifying the portfolio optimization procedure for different kinds of investors. However, it is not further investigated throughout this thesis and is left open for further research.

10.) Sensitivity Analysis of CVaR Estimation

"Estimation or mis-specification errors in the portfolio loss distribution can have a considerable impact on risk measures, and it is important to examine the sensitivity of risk measures to these errors" *(Cont, Deguest & Scandolo, 2008, p. 3).*

In a portfolio context, the application of risk measures is a two step procedure. It is first necessary to estimate the portfolio profit and loss distribution either by using an empirical distribution from historical or simulated data or a parametric form whose parameters are estimated from available data (e.g. using maximum likelihood). As a second step, the risk measure is applied to this estimated profit and loss distribution. Cont, Deguest and Scandolo (ibid.) investigate the sensitivity of quantile-based risk measures while looking at both the choice of estimation method and risk measure jointly in what they call a 'risk measurement procedure'. They define robustness of a risk measure estimator in the sense of sensitivity with respect to an additional (extreme) data point: An estimator of the theoretical risk measure is considered robust if small variations in the loss distribution result in small variations in the estimator. Their striking results indicate that coherence (in fact, sub-additivity) and robustness cannot coexist for the entire class of distribution-based risk measures. As explained in 4.6.2, they use a data set of 1000 observations and find that an additional outlier changes the historical CVaR dramatically while the CVaR estimated from simulated portfolio returns is less sensitive. In order to investigate how the number of observations relates to their findings as well as to examine whether the data winsorization approach of Boudt et al. can mitigate the robustness issues associated with coherent risk measures, qualitative checks of the various CVaR estimation methods are performed.

In an effort to avoid errors when estimating the portfolio profit and loss distribution, the subsequent CVaR sensitivity analysis is based on a historical portfolio return series constructed from an equally-weighted portfolio with monthly rebalancing (no estimations involved). The effect of sample size is investigated by generating the equally-weighted portfolio return series for three different time periods. First, one year of daily return data is selected. Second, half of the available return data is chosen. And third, the entire data set is utilized. An additional outlier is appended to each of the portfolio return series and the corresponding CVaR percentage change is computed for historical and modified (see 4.7.1) CVaRs with and without data winsorization. Figure A.9 depicts the corresponding results.

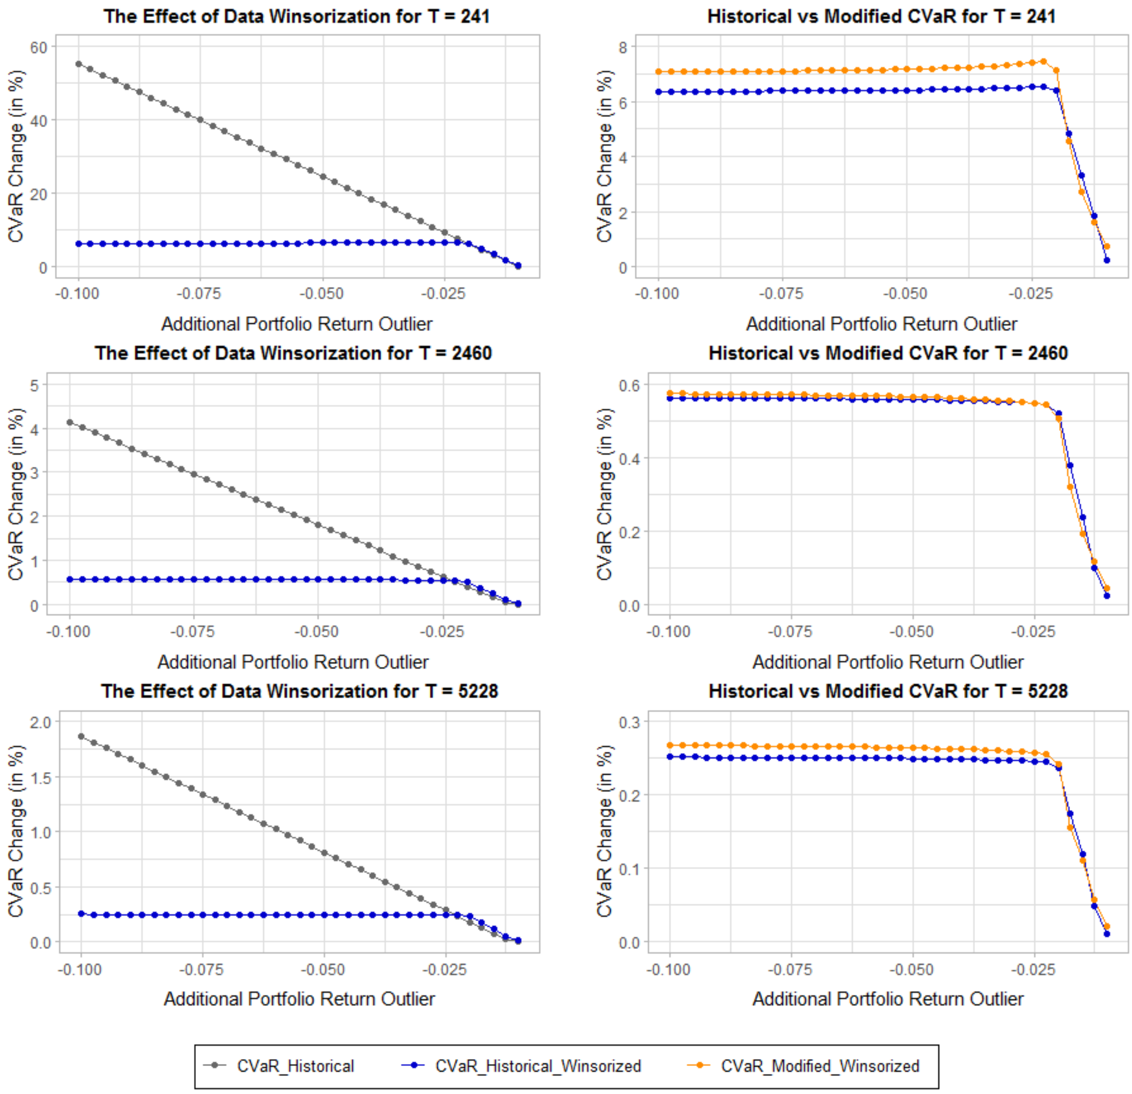

Figure A.9: CVaR Sensitivity Analysis for Different Sample Sizes and Estimators

The sensitivity analysis is based on CVaR estimates at a confidence level of 95% and the first year of return data is excluded because it serves as initial sample throughout the backtesting procedure (see section 5.1). All plots show the sensitivity of the respective risk estimator to the addition of a new data point in the sample. The plots on the left show the effect of data winsorization on the sensitivity of historical CVaR estimates for different sample sizes. The plots on the right side compare the sensitivities of historical and modified CVaR estimates using winsorized data samples of different sizes.

In line with Cont, Deguest and Scandolo (p. 20), the sensitivity of historical CVaR (grey) is unbounded and evolves in a linear manner depending on the outlier size. The degree of sensitivity clearly depends on the sample size. This makes intuitively sense because a larger sample goes hand in hand with the existence of more extreme data points within it, such that an additional outlier has a decreasing effect on the CVaR percentage change.

The data winsorization approach proposed by Boudt et al. (see section 4.6.2) breaks the linearity in historical CVaR sensitivity (blue) and makes it bounded. For very small additional portfolio losses (-0.01 to -0.02), however, data winsorization has no effect and the CVaR sensitivity increases linearly in outlier size. For larger outliers (in an absolute sense), the winsorization kicks in and binds the CVaR sensitivity. This observation makes clear how the data winsorization approach works: For small portfolio losses, the additional portfolio return is not qualified as outlier and is therefore not included in the multivariate winsorization. In this case, the additional observation has the same effect on CVaR sensitivity as for the non-winsorized CVaR estimation. More extreme outliers, on the other hand, are identified as such and are reduced in magnitude by the data winsorization, which makes the CVaR sensitivity bounded. Interestingly, the bounded CVaR sensitivity decreases in magnitude when a larger sample is utilized for the estimation. It is important to realize that the 'threshold', where the winsorization approach begins to identify an additional data point as outlier, crucially depends on the data set at hand: For a data set with very extreme and frequent portfolio losses, it requires a more extreme portfolio loss to be detected by the winsorization and to thereby break the linearity in CVaR sensitivity.

For the modified CVaR (orange), the same negative effect of sample size on CVaR sensitivity can be detected. The right-hand plots of figure A.9 compare the CVaR sensitivities of winsorized data samples for historical and modified CVaRs. In all cases, both the modified and historical CVaR estimations exhibit the same linearity in CVaR sensitivity when the data is not winsorized (for very small additional outliers that are not identified as such). For more extreme outliers – which are identified by the winsorization approach and therefore reduced in magnitude – it can be concluded that the modified CVaR behaves very similar as the historical CVaR, even though it exhibits a slightly larger sensitivity compared to the historical CVaR estimator. However, this difference seems to shrink for large sample sizes.

11.) CVaR Contribution Analysis and MCC Portfolio Implementation

To demonstrate how the MCC portfolio construction works, what it achieves and what pitfalls exist, a CVaR risk contribution analysis is conducted. Figure A.10 depicts the portfolio weights (left) and the CVaR risk contributions (right) of the equal-weight, MinCVaR and MCC portfolio construction techniques based on yearly rebalancing. CVaR contributions are computed using the modified CVaR and its derivative. Moreover, the data is winsorized in order to make the CVaR estimates more robust to outliers (see 4.6.2 and 4.7.1). As in the backtesting procedure, the first year of daily data is used to construct the initial portfolios.

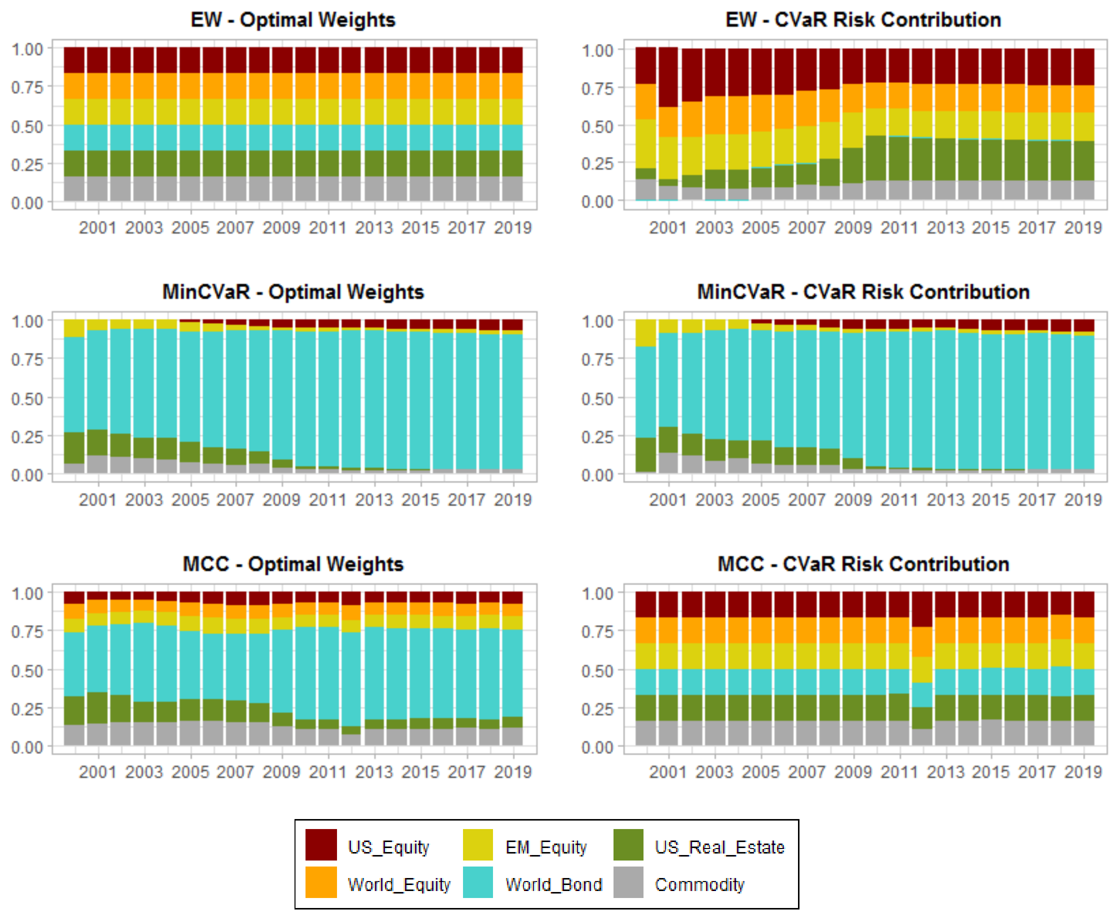

Figure A.10: CVaR Risk Contribution Analysis

It is apparent that the equal-weight portfolio achieves equal weights for each rebalancing date. However, this does not translate to an equal risk contribution of all assets in the portfolio. For example, the US Equity Index contributes more to the portfolio CVaR than its weight in the portfolio would suggest. On the other hand, the World Bond Index does almost not contribute to the portfolio CVaR at all. In fact, at the beginning of the sample period (years 2000, 2001, 2003 and 2004) it contributes slightly negatively to the portfolio CVaR, indicating its role as CVaR reducer in the portfolio. According to Boudt, Carl and Peterson (2013, p. 54), it is the non-linear dependence of portfolio CVaR contributions on the weights that results in relatively poor ex ante risk diversification of the equal-weight portfolio.

The MinCVaR portfolio optimization identifies that the portfolio CVaR can be reduced when associating a relatively high proportion of the investor's wealth to the World Bond Index. This comes at the cost of kicking the World Equity Index out of the portfolio. Moreover, the weight in the US Real Estate Index tends towards zero the more data points are used in the optimization procedure (starting in the time period around the global financial crisis 2008, where this asset class experienced extreme ups and downs), resulting in a rather concentrated portfolio. Interestingly, the assets' CVaR risk contributions are identical to the portfolio weights resulting from minimizing the portfolio CVaR. This observation is perfectly in line with Boudt, Carl and Peterson (p. 46), who provide the analytical proof for the fact that the full investment constraint on the portfolio weights leads to this feature of the MinCVaR portfolio optimization. Therefore, the strong concentration in portfolio weights translates one-to-one to a high concentration in CVaR risk contributions when using MinCVaR under a full investment constraint.

The weights resulting from the MCC portfolio optimization are more balanced across all asset classes compared to the MinCVaR approach. Moreover, it achieves its objective and produces a well-diversified CVaR risk contribution allocation. However, it is important to realize that the CVaR risk contributions of the MCC portfolio approach are not perfectly equal, which can be clearly seen for year 2012. The reason for this is that the maximal CVaR contribution among all assets is minimized, which only serves as a reasonable approximation to achieve MCC's objective (see 4.7.2). However, the difference in CVaR risk contributions among the assets is vanishingly small and cannot be seen with the bare eye for the majority of the optimizations. In year 2012, however, the US Equity Index seems to be overweighted in the sense that it contributes more to the CVaR than it should. Crucial to understanding the MCC portfolio optimization is that the calculations are performed by applying a derivative-free global optimizer. In this thesis, the Differential Evolution algorithm is employed, which is a stochastic optimizer that typically finds only near-optimal solutions and depends on the random seed. This has to be considered when implementing the MCC portfolio optimization. To stress this pitfall, figure A.11 outlines the outcome (including neighbour solutions) of an arbitrary MCC portfolio optimization that fails to achieve an equal risk contribution allocation. In most cases, these issues can be mitigated to a great extent by adjusting the convergence tolerance as well as the tuning parameters of the algorithm that solves the MCC optimization. In a backtesting setting, however, a rather conservative convergence tolerance leads to more iterations in each rebalancing period, such that a trade-off between computation time and objective precision arises.

Figures A.10 and A.11 are only for illustrative purposes. The goal is to show the characteristics and potential pitfalls of the MCC portfolio implementation. Throughout this thesis, the results of both the MCC and Worst-Case MCC portfolios are based on a very conservative convergence tolerance with a maximum of 200 iterations for each rebalancing period, which leads to an equal risk contribution of all assets up to the sixth digit after the comma.

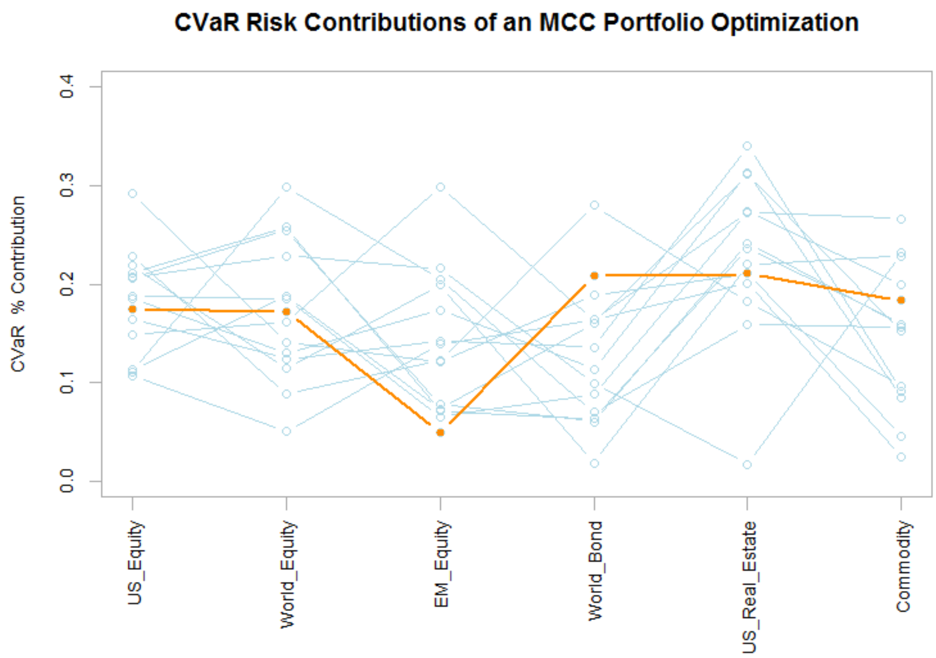

Figure A.11: Near-Optimal Results of an Arbitrary MCC Portfolio Optimization

12.) Backtesting Results – Portfolio Weights

The subsequent figures (A.12 - A.15) show the resulting portfolio weights for each portfolio construction technique for yearly, quarterly, monthly and weekly rebalancing. It is central to understand that the portfolio weights are the core of each portfolio construction technique. Once the optimal portfolio weights are obtained, the respective portfolio returns as well as the corresponding wealth trajectories can be investigated. As a general observation, it can be noted that all optimization-based techniques experience some sort of change in portfolio weights around the year 2008 as very extreme outliers enter the optimizations.

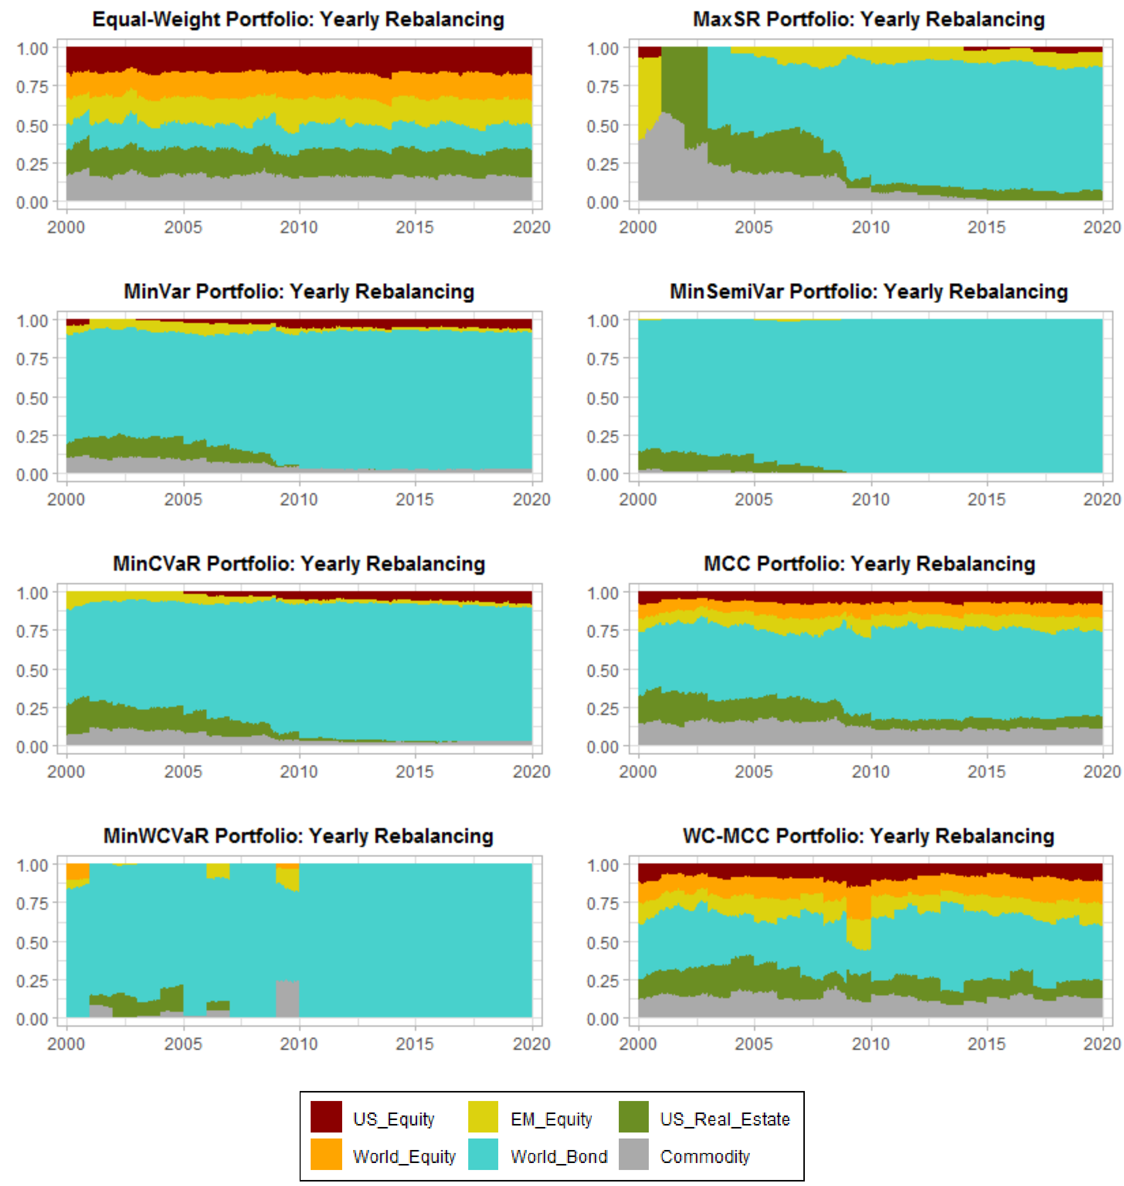

Figure A.12: Backtesting Results – Portfolio Weights for Yearly Rebalancing

In comparison to the yearly rebalancing, performing the portfolio rebalancings each quarter leads to more stable asset allocations. This can be clearly seen for the equal-weight portfolio. The MaxSR portfolio still produces rather wild weights, which might be attributed to the sensitivity of mean-variance models with respect to their input parameters. MinVar, MinSemiVar and MinCVaR produce very concentrated portfolios that do not change substantially if they are implemented more often over time. MinCVaR, MCC and their copula-based worst-case extensions produce similar portfolio weights, even though the results stemming from optimizations based on simulated data fluctuate more than those based on historical data.

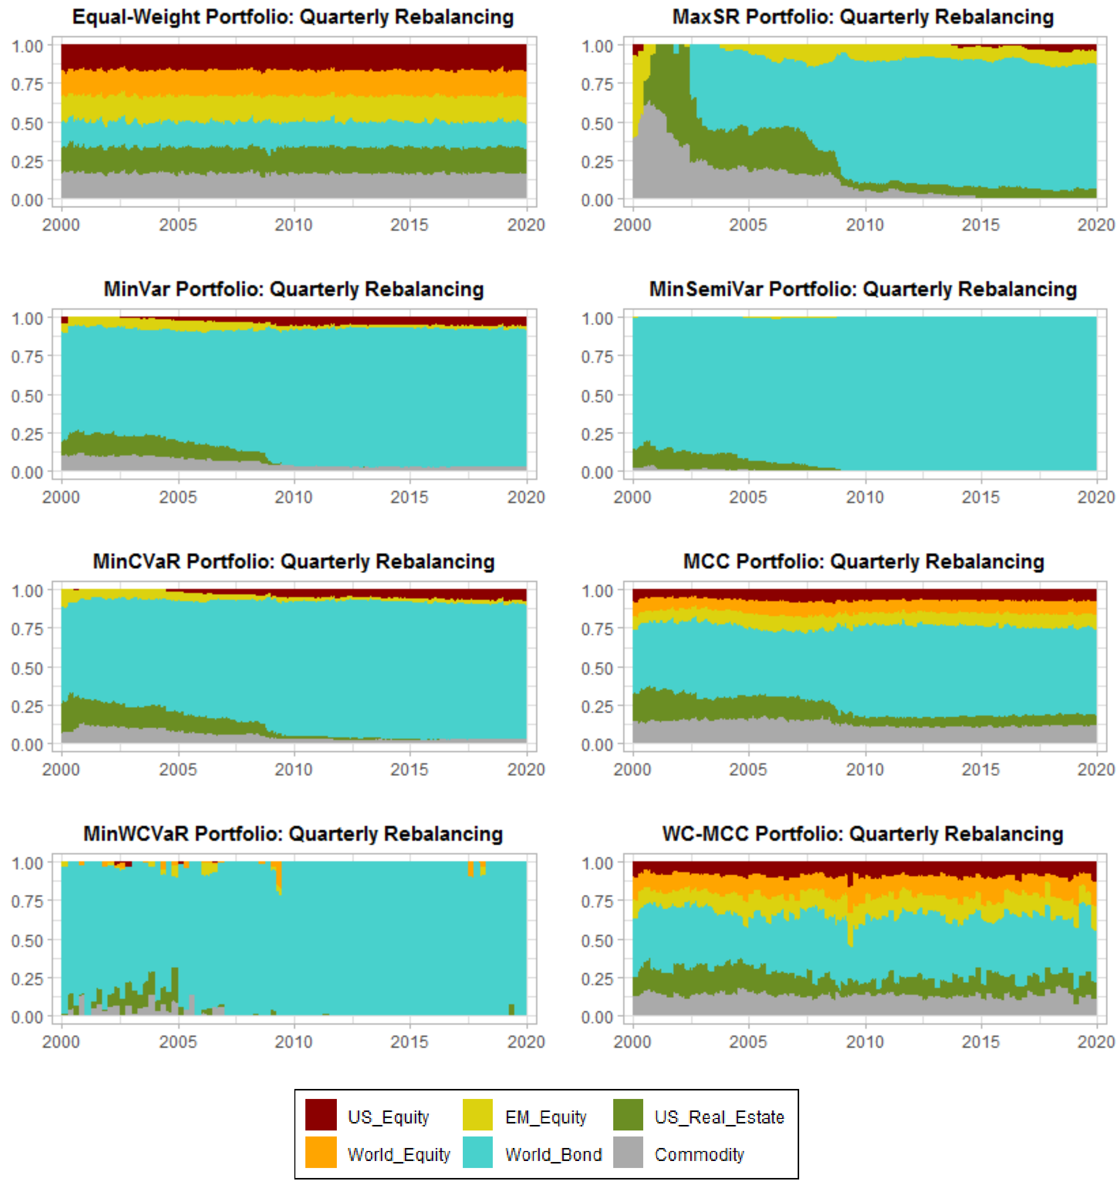

Figure A.13: Backtesting Results – Portfolio Weights for Quarterly Rebalancing

The stability of the portfolio weights over time increases even more for monthly rebalancing. Interestingly, MinVar and MinCVaR produce very similar portfolio weights. As explained in the theory part in section 4.6.2, this happens only if the assets' returns are normally distributed. However, section 3 showed that the return data in the present context does not follow a normal distribution. Nevertheless, among all the investible assets, the World Bond Index is the only asset that produces returns that are the closest to a normal distribution. The large concentration of both optimizations in this respective index might explain why they do not produce perfectly equal but similar portfolio weights.

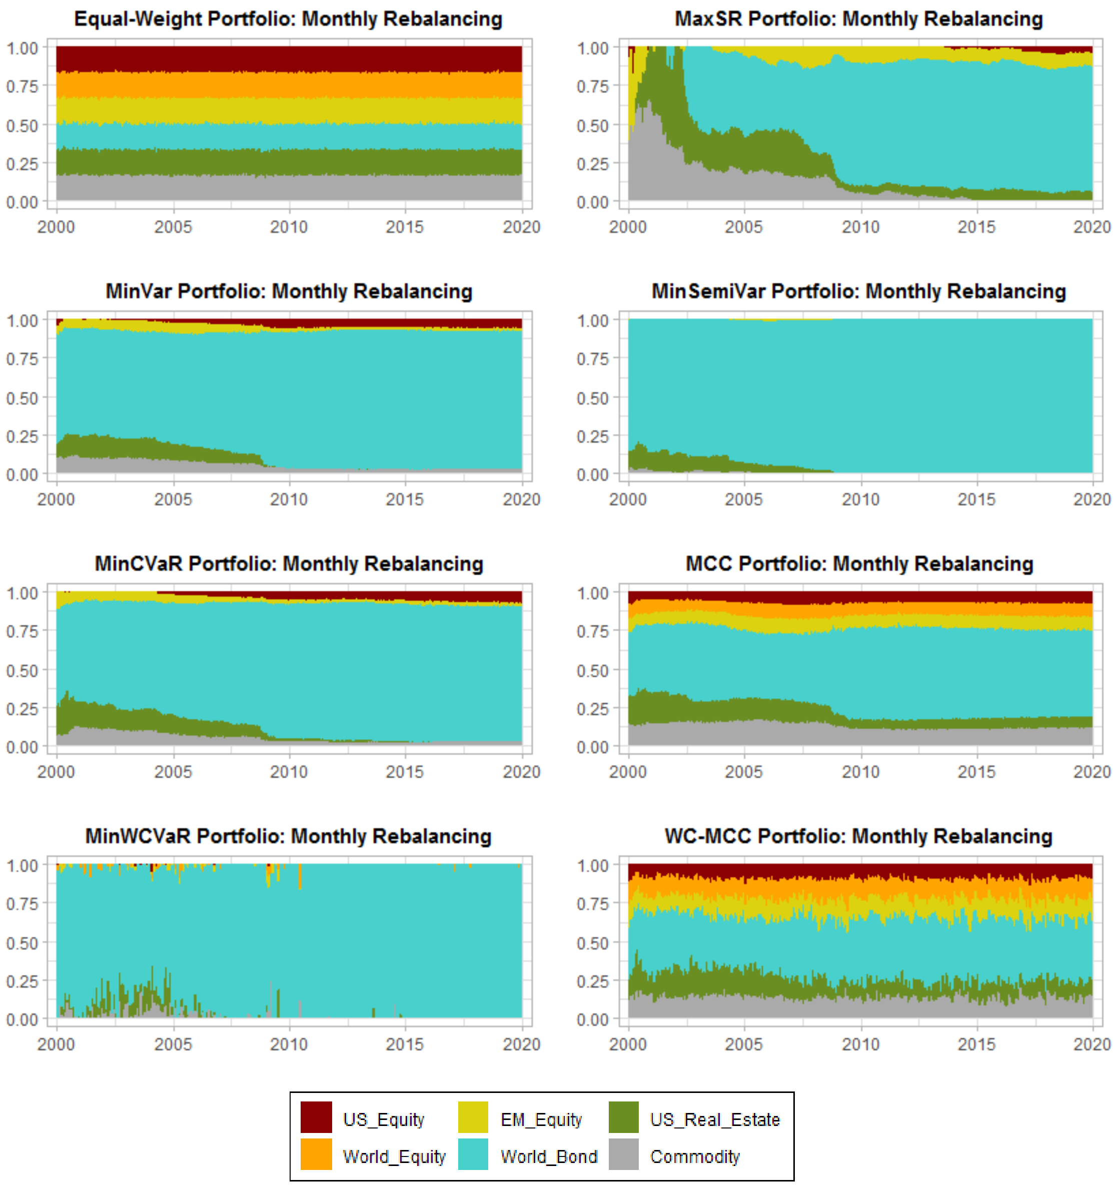

Figure A.14: Backtesting Results – Portfolio Weights for Monthly Rebalancing

Apparently, the results from simulation-based optimizations of MinCVaR and MCC become more stable and get closer to the results of the optimizations with historical data for weekly rebalancing. The MinWCVaR portfolio produces a highly conservative portfolio as to protect the investor from the worst-case possible. The WC-MCC portfolio, on the other hand, produces a well-balanced asset allocation similar to the equal-weight portfolio but additionally diversifies the portfolio with respect to risk contribution. Both worst-case approaches require a lot of trading as is reflected in the portfolio turnover (see Table 3).

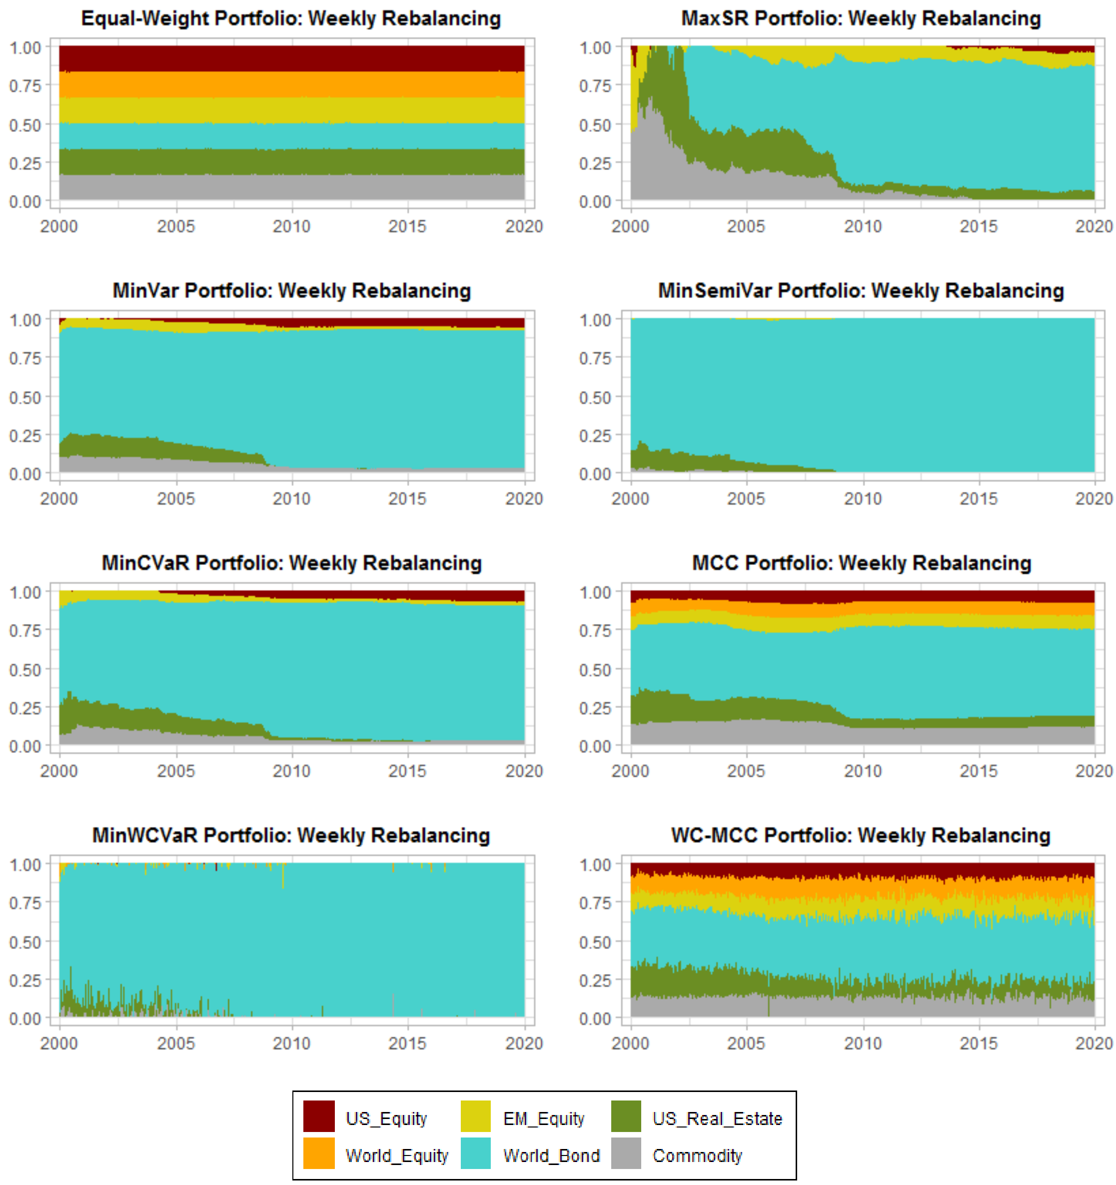

Figure A.15: Backtesting Results – Portfolio Weights for Weekly Rebalancing

Declaration of Authorship

I hereby declare

- that I have written this thesis without any help from others and without the use of documents and aids other than those stated above;

- that I have mentioned all the sources used and that I have cited them correctly according to established academic citation rules;

- that I have acquired any immaterial rights to materials I may have used such as images or graphs, or that I have produced such materials myself;

- that the topic or parts of it are not already the object of any work or examination of another course unless this has been explicitly agreed on with the faculty member in advance and is referred to in the thesis;

- that I will not pass on copies of this work to third parties or publish them without the University's written consent if a direct connection can be established with the University of St.Gallen or its faculty members;

- that I am aware that the University will prosecute any infringement of this declaration of authorship and, in particular, the employment of a ghostwriter, and that any such infringement may result in disciplinary and criminal consequences which may result in my expulsion from the University or my being stripped of my degree.

By submitting this academic term paper, I confirm through my conclusive action that I am submitting the Declaration of Authorship, that I have read and understood it, and that it is true.

St. Gallen, February 2020

Ferid Dupljak