PORTFOLIO OPTIMISATION USING ALTERNATIVE RISK MEASURES

Dissertation by

DOUGLAS AUSTEN LORIMER

Submitted in fulfilment of the requirements for the degree

Master of Philosophy in Finance and Investment

in the

Faculty of Economic and Management Sciences

at the

University of Pretoria

Supervisors:

Prof CH van Schalkwyk

Dr J Szczygielski

June 2023

UNIVERSITEIT VAN PRETORIA
UNIVERSITY OF PRETORIA
YUNIBESITHI YA PRETORIA

DECLARATION

1. I understand what plagiarism entails and am aware of the University's policy in this regard.

2. I declare that this assignment is my own, original work. That all sources used / quoted have been indicated and acknowledged by means of a complete reference system.

3. I did not copy and paste any information directly from an electronic source (e.g., a web page, electronic journal article or CD ROM) into this document.

4. I did not make use of another student's previous work and submitted it as my own.

5. This dissertation was not previously submitted for a degree at another university.

Signature

16 June 2023

Date

UNIVERSITEIT VAN PRETORIA
UNIVERSITY OF PRETORIA
YUNIBESITHI YA PRETORIA

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

LIST OF ABBREVIATIONS and acronyms

Abbreviation/ Acronym	Meaning
CVaR	Conditional value at risk
DSSD	Downside semi deviation
DSV	Downside semivariance
ETF	Exchange traded fund
EWMA	Exponentially-weighted moving average
GARCH	General Autoregressive Conditional Heteroscedasticity [model]
GICS	Global Industry Classification Standard
HMAE	Heteroskedasticity-adjusted mean absolute error
LPM	Lower Partial Moment
MAD	Mean absolute deviation
MPT	Modern Portfolio Theory
MSCI	Morgan Stanley Capital International
S&P	Standard and Poor's
SAD	Semi-absolute deviation
SD	Standard deviation
TER	Total Expense Ratio (of an ETF)
V	Variance
VaR	Value at risk
USD	United States Dollars
NASDAQ	National Association of Securities Dealers Automated Quotations
NYSE	New York Stock Exchange

LIST OF DEFINITIONS

Capital Market Line	The tangent line between the expected return of the risk-free asset and the efficient frontier
Conditional heteroskedasticity	The characteristic that changes in a risk measure are conditional upon past factors. In investment returns, this leads to clustering of volatility in time.
Conditional value at risk (CVaR)	The expected return considering only those returns below the VaR
Downside semivariance	The average of squared deviations from the mean, given such deviations are less than zero
Efficient diversification	The premise that a risk-averse investor will choose a portfolio with a lower risk for any given expected return, or equivalently, a portfolio with a higher expected return for any given level of risk
Heteroskedasticity	The characteristic that statistical dispersion (here, the risk measure) is non-constant with respect to an independent variable (here, time)
Heteroskedasticity-adjusted mean absolute error (HMAE)	Average of the size of the deviation between a risk measure and an estimate of that risk measure.
Inter-temporal dependence	Utility of consumption is a function of prior consumption
Lower partial moment (LPM(τ,k))	Class of downside risk measures taking into account only those returns lower than a specified target return, τ. The k^{th} LPM is denoted LPM(τ,k)
Lower partial standard deviation	Square root of downside semivariance
Mean-Absolute Deviation (MAD)	The average of the absolute value of deviations from the mean
Modern Portfolio Theory (MPT)	Mathematical framework for maximising expected return for a specified level of risk for a portfolio of assets
Non-linearity	Change in output is not proportionate to change in inputs

Optimal portfolio	The portfolio on the efficient frontier which connects the Capital Market Line to the risk-free asset
Rebalancing	The act of resetting portfolio weights back to the original portfolio weights
Rebalancing premium	The improvement in performance resulting from rebalancing
Risk measure	A value assigned to random variables interpreted as losses
RM-optimised portfolio	A portfolio optimised using the risk measure RM
Semi-absolute Deviation (SAD)	The average of the absolute value of deviations from the mean including only observations which are below the mean
Value at Risk (VaR)	The *n*th-percentile lowest return over a period

ABSTRACT

Markowitz' Modern Portfolio Theory (MPT) optimises the ratio of mean portfolio returns and portfolio risk in the form of the variance of returns, giving rise to criticism relating to, *inter alia*, minimising upside risk, the assumption of normally-distributed returns, and a failure to recognise heteroskedasticity. In addressing these criticisms, this research investigates the use of alternative risk measures to optimise risk and return in MPT investment strategies using non-parametric numerical methods to optimise portfolios comprising assets from the S&P 1200 and MSCI GICS world indices. It investigates, in particular, downside semivariance, downside semideviation, mean absolute deviation, semi-absolute deviation, value at risk and conditional value at risk. In addition, the study investigates optimisation using backward-looking and forward-looking risk measures through exponentially-weighted moving average forecasts of risk measures and return. In general, all the alternative risk measures investigated result in investment strategies with higher returns than traditional MPT variance-optimised strategies, with semi-absolute deviation-optimised strategies performing best of all. The introduction of risk and return forecasting does not materially impact on strategy performance.

CHAPTER 1 - INTRODUCTION

1.1 INTRODUCTION

Modern Portfolio Theory (MPT), developed by Harry Markowitz, is a widely-lauded portfolio construction technique which allows the investor to trade off risk and reward (Markowitz, 1952). It led to Markowitz's award of the Nobel Prize for Economics in the 1950s and to Pensions & Investments magazine naming him the Man of the Century in 1999.

Despite these accolades, MPT has its detractors. MPT uses the variance of stock returns as its measure for risk (Markowitz, 1952) even though an investor who successfully minimises variance in his portfolio minimises not only downside variation, but also upside variation. This, along with other underlying assumptions and characteristics of MPT has garnered the attention of critics (Boasson et al., 2011; Byrne & Lee, 2004; Cardoso et al., 2019; Hunjra et al., 2020).

Other criticisms of MPT include the assumption of symmetrical, normally-distributed returns (Boasson et al., 2011; Stanković et al., 2020), as well as the assumption that a portfolio which optimises the ratio of risk and return will maintain that ratio into the future (Markowitz, 1952).

Markowitz's original model relied, *inter alia*, on risk being defined as variance of stock returns. Subsequently, he himself noted that variance is not an ideal risk measure for portfolio optimisation (Markowitz & Cowles, 1959).

However, despite these criticisms, the technique of optimising the risk-return ratio which underlies MPT can still be used to build portfolios with desirable characteristics. One means of applying this technique is to apply it to risk measures other than variance, leading to numerous studies having used alternative risk measures to optimise portfolios (Boasson et al., 2011; Byrne & Lee, 2004; Gaivoronski, 2005; Hunjra et al., 2020; Lohre, 2010; Sortino, 1994; Stanković et al., 2020). Accordingly, where MPT optimises the ratio between expected return and variance, it is possible to optimise the ratio between expected return and downside semivariance, value at risk or conditional value at risk. Different risk measures have different characteristics which may make their use more tractable; some require no assumptions as to distribution or parameterisation, while others are asymmetric, measuring downside risk but ignoring upside risk. The characteristics of different risk measures attract different investors with different needs and requirements (Rachev et al., 2008). By substituting variance with a risk measure which fulfils these needs, it may be possible to construct a portfolio which minimises the relevant investor's risk relative to a given return, and which may carry this optimised characteristic into the future.

Each alternative risk measure has different characteristics and advantages. Downside semivariance is similar to variance, but includes only those returns below the sample mean return, thereby avoiding the minimisation of upside variance (Boasson et al., 2011; Hunjra et al., 2020; Lohre, 2010). Mean absolute deviation measures both upside and downside deviation, but includes deviations from the sample mean linearly, rather than squared deviations. This makes mean absolute deviation less sensitive to outliers than variance (Byrne & Lee, 2004). Semi-absolute deviation includes only downside deviations, and similarly to mean absolute deviation, is less sensitive to outliers than downside semivariance. Value at risk allows the investor to specify risk in terms of the n^{th}-lowest percentile returns over a period. It may be estimated empirically and does not rely on normality of returns, but may be discontinuous, making portfolio optimisation difficult (Gaivoronski, 2005). Conditional value at risk is the expected return of returns lower than the value at risk, which shares some of the benefits of value at risk, but is also smooth and therefore more tractable to optimise.

In further criticism of MPT, Lukomnik and Hawley (2021) argue that MPT fails to capture ambiguous events, black swan events and systemic risk. They cite COVID-19 as an example, which led to the loss of 15% of payroll jobs in the USA in just two months in 2020, but which, for obvious reasons, was not captured by preceding returns data and therefore was not an investment input into MPT models. In corroboration, the 2020 market downturn is evident in both the benchmark portfolio and the MPT portfolios created in this study.

MPT also fails to capture investment influences which relate to corporate governance such as ESG (Lukomnik & Hawley, 2021). In combination, these shortfalls make MPT susceptible to market failures arising from policy, external influences and information asymmetry (Iyiola et al., 2012).

This leads to another criticism: That risk and return do not persist into the future, and out-of-sample performance of mean-variance optimised portfolios is often poor (Rigamonti & Lučivjanská, 2022). Markowitz's portfolio optimisation relies on the persistence of expected returns and variance if it is to build a portfolio which performs well in the future. However, risk is known to be heteroskedastic (Rachev et al., 2008) and volatility exhibits clustering (Dachraoui, 2018). In reality, the expected values of risk and return used in MPT do not incorporate the circumstances which presently exist, but which did not exist when historical data upon which the forecasts are based was generated. That is, investors must estimate future values from past data which says nothing about why the risk being modelled is actually arising (Iyiola et al., 2012). In contradiction of this criticism, Lohre *et al.* (2010) found only a small difference in performance between their in-sample and out-of-sample results.

This raises the question as to whether predictive models can be combined with MPT. There exist more sophisticated models for forecasting risk than to assume that the variance in the next period can be established by calculating the sample variance of the previous periods. If these more sophisticated models are used to forecast risk and expected returns, it may be that the MPT technique of optimising the risk-return ratio can be applied to the forecasts to construct an optimal portfolio which better maintains the desired risk-return ratio into the future.

One such predictive model which may be so applied is the exponentially-weighted moving average (EWMA), a special case of the General Autoregressive Conditional Heteroscedasticity (GARCH) models of variance, which are able to capture heteroskedasticity and volatility clustering (Danielsson, 2011).

This study will investigate whether applying these variations to MPT results in better performance, thereby providing additional evidence for the corroboration or contradiction of the portfolio optimisation techniques used in MPT.

1.2 PROBLEM STATEMENT

MPT has a number of shortfalls which this study aims to address. Firstly, there is an implicit assumption in MPT that variance must persist for at least some period into the future from the date at which the portfolio is constructed (Maeso & Martellini, 2020). If this implicit assumption is correct, then the desirable characteristics of the portfolio for which an investor has optimised today will persist such that his portfolio will retain those same desirable characteristics tomorrow. However, the variance of stock returns is known to be heteroskedastic (Rachev et al., 2008), violating MPT's implicit assumption of persistence of variance, and therefore weakening the premise that MPT will result in a better-performing portfolio in terms of risk and return compared to an arbitrary benchmark.

Secondly, the risk measure traditionally used in MPT is variance, which is symmetrical and therefore imparts a penalty to assets imparting upside deviations from the mean return as well as downside deviations. The impact thereof would be exacerbated in a bull market, where a strict application of MPT would reduce the weighting of assets with the largest upside returns relative to their peers.

Thirdly, since the variance of stock returns is heteroskedastic (Rachev et al., 2008) the investor must incorporate new information into his portfolio construction periodically for the portfolio to remain up-to-date with respect to changes in expected return and risk. MPT does not put forward a methodology for determining the importance and concomitant influence of new information as compared to old information in recalculating the portfolio at successive time steps, since variance, the risk measure proposed by Markowitz, weights each return observation in the sample period equally.

Finally, studies of MPT which optimise portfolios using different risk measures result in contradictory conclusions as to the efficacy thereof. More studies are required to build a body of evidence to demonstrate the merit or otherwise of applying MPT in this manner.

1.3 PURPOSE STATEMENT

The purpose of the study is to address the criticisms of MPT relating to minimising upside risk, the assumption of normally-distributed returns, and a failure to recognise heteroskedasticity through the use of alternative risk measures, non-parametric methods and the forecasting of risk measures.

1.4 OBJECTIVES OF THE STUDY

The study has three objectives; firstly, the study will investigate whether the use of alternative risk measures in portfolio optimisation leads to better portfolio performance. Further, the study will investigate whether the use of a forecasting model which takes into account the persistence and heteroskedasticity of the risk measures calculated on the returns results in better portfolio returns or risk-adjusted returns than the use of backwards-looking risk measures. Of the studies reviewed, exponentially-weighted moving averages have been applied only to variance forecasting in MPT.

The second purpose of the study will be to provide an empirical analysis of portfolios of sector exchange-traded funds (ETFs), represented by indices, optimised using a comprehensive variety of risk measures so as to facilitate comparisons with the majority of empirical studies of portfolio construction using alternative risk measures in other asset classes.

The third purpose of the study is to investigate the efficacy of recalculating the optimal portfolio over different periods, thereby testing whether the frequency of incorporating new information into the portfolio’s construction impacts upon its performance.

1.5 CONTRIBUTION

The study will provide a novel comparison of the performance of portfolios optimised using numerous unmodified risk measures as compared to the performance of portfolios optimised using an exponentially-weighted moving average risk forecasting model for those same risk measures. This will assess whether the application of an EWMA to the risk measures addresses the need to balance between the heteroskedasticity and short-term persistence of risk measures, thereby improving the performance of the portfolios they are used to construct and providing a method to improve the toolset of investors using MPT and MPT-like methods. Of the studies reviewed, only one applied an EWMA to a single risk measure, variance, used to optimise portfolios (Cardoso et al., 2019).

Further, by examining a wide variety of risk measures, the study will facilitate comparison to most existing empirical studies of portfolio optimisation using alternative risk measures based on different asset classes. This will provide a larger universe of risk measures for investors using MPT and MPT-like methods to choose from, through a side-by-side comparison of the relative merits and disadvantages of the use of each alternative risk measure. In particular, investors with a risk appetite which is not well characterised by portfolio variance, or who are required to measure their performance in terms of risk other than variance, may find, where classical MPT is precluded, that MPT-like methods associated with an alternative risk measure are useful in managing portfolios.

The study will aim to provide some insight into the hyperparameters used (those parameters used to optimise the parameters themselves) and how they are arrived at, providing guidance to investors investigating the use of MPT or MPT-like optimisation methods. None of the literature reviewed discussed the process of parameterising the portfolios constructed explicitly, nor the relative merits of different hyperparameters.

The study will examine the efficacy of portfolio construction on global sector indices. Literature relating to the construction of portfolios using indices as a proxy for investable ETFs is limited. Of the literature reviewed, only one study made use of indices, using indices to mimic different classes of securities (Boasson et al., 2011). However, this related to only one of the risk measures proposed for this study. The study will be conducted on global sector index data used as a proxy for the ETFs which use them as a benchmark. The data sample selected for the study has the advantage of broad coverage, as the indices incorporate circa 70% of global market capitalisation (Standard and Poor's, 2022), giving the study applicability in many global markets.

In summary, the study addresses several of the criticisms of classical MPT, thereby improving the applicability of MPT-like methods in portfolio management.

1.6 DELIMITATIONS AND ASSUMPTIONS

The study will investigate only standard deviation, variance, downside semivariance mean absolute deviation, semi-absolute deviation, value at risk (5% and 10% levels) and conditional value at risk (5% and 10% levels) to establish whether they improve portfolio performance when compared to variance-optimised portfolios.

The relative merits of risk forecasting models other than EWMA (such as GARCH) will not be investigated in the study.

The study will not derive closed-form solutions for optimisations.

1.7 STRUCTURE OF THE STUDY

A literature review is included in chapter 2, introducing MPT, the main criticisms of MPT and the alternative risk measures used in the study. Past studies conducted applying MPT-like methods to MPT are discussed. Chapter 3 covers the research design and the data used in the study. The non-parametric numerical methods techniques for optimising each portfolio are discussed, as is the methodology used to forecast expected returns and risk measures into the future. Chapter 4 covers the results of the study and the analysis thereof. The chapter contains a comparison of all the strategies investigated, commentary on the benefits of using an EWMA for forecasting and discussion about the most successful of the strategies. Chapter 5 provides a conclusion, including possible implications and avenues for future research.

CHAPTER 2 - LITERATURE REVIEW

2.1 INTRODUCTION

This chapter will present the theoretical basis for the study, discuss the relevant existing research and scholarly debates, show how the study relates to them, and present the new insights it will contribute. Accordingly, portfolio management theory is discussed before addressing Modern Portfolio Theory (MPT), and the subsequent extension thereof through the inclusion of a risk-free asset.

The characteristics, advantages and disadvantages of the alternative risk measures investigated in the study are examined next, followed by a discussion of the issues and the best practice relating to forecasting these risk measures.

2.2 PORTFOLIO MANAGEMENT THEORY

Portfolio management is the process whereby an investor selects a portfolio of assets (typically securities) to invest into so as to meet his needs and preferences. It incorporates rebalancing, monitoring and evaluating the performance of the portfolio (Bodie et al., 2017). Portfolio management incorporates different frameworks, theories or principles, including that of MPT.

Portfolio Management Theory can be categorised into the traditional approach and the modern approach, each of which incorporates different theories. Under the traditional approach, Dow Jones theory is characterised by the idea that markets are efficient and cyclical, following trends in the long, medium and short term, which can be confirmed by adequate trade volumes (Schannep, 2012).

Hamilton (1922), an early proponent of another traditional approach, Formula Theory, built on Dow's articles in the Wall Street Journal, maintaining that the stock market moved in cycles of growth and decline which could be predicted using mathematical formulae.

Formula Theory also incorporates the proposal that mathematical models can be used to determine the intrinsic value of stocks, based on the characteristics of the underlying company and discounting of future predicted dividends (Malkiel, 2015).

Also under the traditional approach, Malkiel's (2015) Random Walk Theory argues that stock returns are predominantly random and unpredictable, advocating passive investment strategies over stock picking.

It is not, however, the traditional approach to portfolio management theory which is the subject of this study, but rather the modern approach.

2.3 MODERN PORTFOLIO THEORY

MPT falls under the modern approach to portfolio theory. It comprises a framework or set of principles which sets out how a rational investor can select a portfolio of assets in such a way as to trade off risk and return to arrive at an efficiently-diversified portfolio. The founding premise of MPT is efficient diversification, which provides that a risk-averse investor will choose a portfolio with a lower risk for any given expected return, or equivalently, a portfolio with a higher expected return for any given level of risk (Markowitz & Cowles, 1959; Sharpe, 1964).

MPT provides that the investor should not maximise discounted expected returns, since doing so would require placing the investor's full funds into a single stock, thereby ignoring diversification (Markowitz, 1952). Markowitz described a process by which an investor could identify a set of efficient portfolios, each of which, for each given expected return, would have a minimum risk; or equivalently, for each given level of risk, would have the greatest possible expected return. Markowitz used the variances and covariances of assets to package stocks with negative correlations into the portfolio and thereby, given a specified portfolio return, minimise the variance of the portfolio. Markowitz used the variance of stock returns as the measure of risk to identify these optimal portfolios because of variance's relative computational tractability, but noted that he believed semivariance would have been a better risk measure (Markowitz & Cowles, 1959). Markowitz concluded that mean-variance efficiency might be a reasonable working hypothesis for investors seeking yield but trying to avoid risk.

2.4 INCLUSION OF A RISK-FREE ASSET

Given a risky portfolio and a risk-free asset, the Capital Allocation Line (CAL) is the set of all possible linear combinations on the standard deviation-return plane (Bodie et al., 2017). The Capital Market Line (CML), PφZ in Figure 0.1, is a special case of the CAL, with the risky portfolio having been specified as that which makes the CAL tangent to the efficient frontier (AφB), joining the risk-free asset (P) and the efficient frontier, where the efficient frontier is that set of portfolios described by Markowitz (1959), each of which, for each given expected return, would have a minimum risk; or equivalently, for each given level of risk, would have the greatest possible expected return.

Figure 2.1

The Efficient Frontier and the Capital Market Line

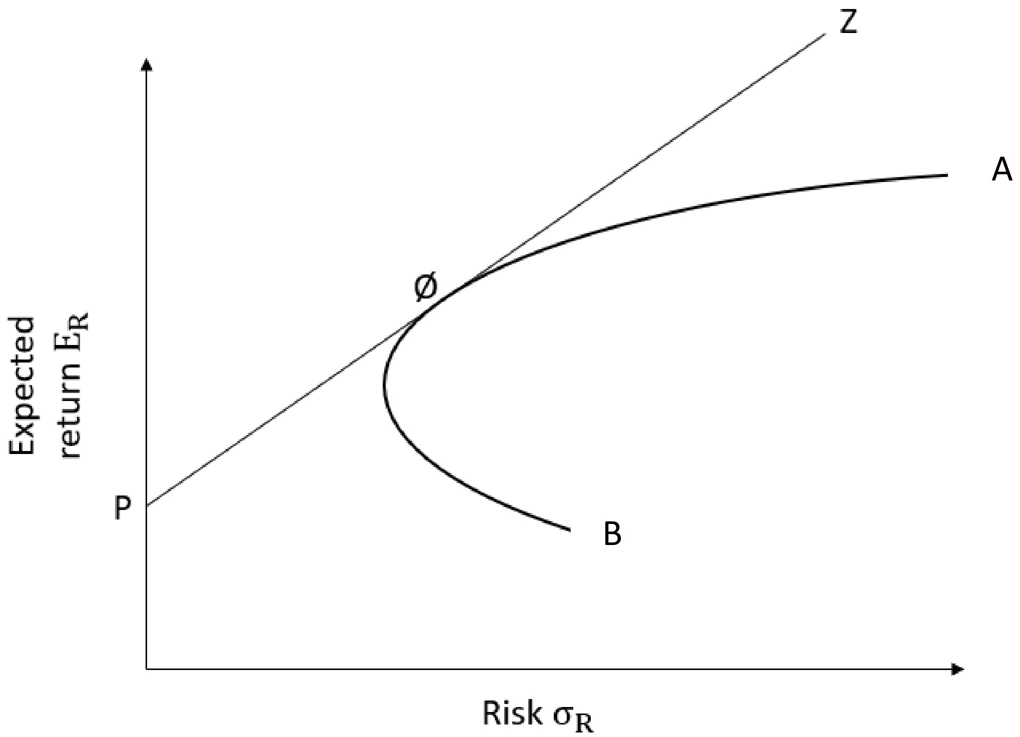

Note. Graph of the efficient frontier and the capital market line. AφB is the efficient frontier. PφZ is the Capital Market Line.

Source: Sharpe, 1964.

Analogously, the risky portfolio is that portfolio on the efficient frontier at which the CML is tangential to the efficient frontier (at φ). Consequently, given an efficient risk-expected return frontier of possible portfolios (AφB) and a risk-free asset (P), any combination of risk σ_R and expected return E_R on the CML, PφZ, is attainable if the investor allocates a portion of his assets to the risk-free asset and a portion to the efficient portfolio at φ; or alternatively, borrows at the same risk-free rate and invests the proceeds thereof into the efficient portfolio at φ. Consequently, other than where the investor invests all his assets into the portfolio at φ with no borrowing, the investor is able to achieve a higher expected return for the same level of risk, or a lower level of risk for the same expected return than by investing into any portfolio on the efficient frontier (Sharpe, 1964). In consequence, the portfolio φ is the optimal efficient portfolio of all portfolios which may be constructed from the risky assets being considered.

2.5 CRITICISMS OF MODERN PORTFOLIO THEORY

Since its inception in 1952, academics and investors have raised many criticisms about MPT. These relate to the quantitative assessment of risk and return and its relationship with future risk and return as well as MPT’s underlying assumptions about investor behaviour.

2.5.1 MPT assumes probability distributions are constant through time

One of the criticisms of Markowitz’s mean-variance portfolio optimisation is the requirement for the simplifying assumption that the probability distributions of asset returns is static in time (Markowitz, 1952), which is not the case. Engle and Patton (2001) discuss several characteristics of volatility of returns, which included that volatility exhibits persistence, that it is mean-reverting and that it is influenced by exogenous variables. These characteristics imply that volatility is not constant over time, and that the probability distribution of returns is therefore also not constant over time. Engle (2001) goes on to note that heteroskedasticity is often an issue in financial time series, stating that “even a cursory look at financial data suggests that some time periods are riskier than others”. French et al. (1987) similarly observed that the volatility of stock returns is not constant when they examined the relationship between expected market premium and volatility. Reducing reliance on this simplifying assumption requires variance (or other risk measure) forecasts which vary from time-to-time, the estimation of which is crucial to portfolio optimisation, as is purported by Gosier (2005) who concluded that the use of shorter sampling intervals improves risk forecasting, not least because it permits the use of more timely information and reduces staleness. While *observed* expected returns and covariance for some period of the past for the constituents of a hypothetical portfolio are tentative inputs for their *expected* return and covariance, better methods than these must be available (Markowitz, 1952).

2.5.2 MPT assumes risk and return persist into the future

Another criticism of MPT is that risk and return do not persist into the future, and out-of-sample-performance of mean-variance optimised portfolios is often poor, as noted by Rigamonti and Lučivjanská (2022) who explore tractable methods to estimate the downside semivariance correlation matrix. Investors must estimate future values from past data which says nothing about why the risk being modelled is actually arising (Iyiola et al., 2012).

Lohre et al. (2010) used the proportion of periods in which the optimised portfolio showed lower risk than the benchmark as a measure to indicate the efficacy of mean-risk measure optimised portfolios. They reported findings on mean-variance-optimised portfolios which contradicted the criticism that out-of-sample-performance of mean-variance optimised portfolios is often poor – they found that the number of periods in which the volatility of their Eurostoxx variance-optimised portfolio was lower than their benchmark moved from 94.5% when optimised over the full study period (that is, in-sample optimisation), to 93.15% when optimised out of sample.

2.5.3 MPT minimises upside variation

A further criticism of MPT is that it uses the variance of stock returns as its measure for risk (Markowitz, 1952) despite the fact that an investor who successfully minimises variance in his portfolio minimises not only downside variation, but also upside variation. This, along with other underlying assumptions and characteristics of MPT has garnered the attention of critics (Boasson et al., 2011; Byrne & Lee, 2004; Cardoso et al., 2019; Hunjra et al., 2020).

Other risk measures exist which may substitute for variance as the risk measure to be optimised against (Hunjra et al., 2020). For example, where MPT optimises the ratio between expected return and variance, it is possible to optimise the ratio between expected return and downside semivariance, value at risk or conditional value at risk. Different risk measures have different characteristics which may make their use more tractable; some require no assumptions as to distribution or parameterisation, while others are asymmetric, measuring downside risk but ignoring upside risk. The characteristics of different risk measures attract different investors with different needs and requirements (Rachev et al., 2008). By substituting variance with a risk measure which fulfils these needs, it may be possible to construct a portfolio which minimises the relevant investor's risk relative to a given return, and which may carry this optimised characteristic into the future.

2.5.4 MPT assumes normally-distributed returns

The MPT assumption of symmetrical, normally-distributed returns is incorrect (Boasson et al., 2011; Stanković et al., 2020). The market exhibits variance of returns which is greater and more frequent than the normal distribution would allow (Omisore et al., 2012). Richardson and Smith (1993) explore multiple tests for multivariate normality of returns where stock returns are highly correlated, noting that violation of the normality assumption in empirical studies can lead to incorrect inferences. They point out that while univariate normality does not imply multivariate normality, multivariate normality implies univariate normality, before showing that a sample of stock returns from the Dow Jones is not multivariate normal for the majority of the periods examined.

Additionally, the distribution of returns evidenced in the market is not symmetrical, while normally-distributed market returns would be. Richardson and Smith (1993) found skewness in 45 of 130 tests on return distributions investigated. Skewness of returns makes variance an inappropriate risk measure for optimisation of portfolios, since variance penalises upside dispersion relative to the mean return just as it does downside dispersion. While this would not be problematic if returns were symetrically distributed, since upside dispersion and downside dispersion would be equal, upside dispersion and downside dispersion are not equal if returns are skewed (Boasson et al., 2011).

2.5.5 MPT fails to capture systemic events

Lukomnik and Hawley (2021) argue that modern portfolio theory fails to capture ambiguous events, black swan events and systemic risk. They cite COVID-19 as an example, which led to the loss of 15% of payroll jobs in the USA in just two months in 2020, but which, for obvious reasons, was not captured by preceding returns data and therefore was not an investment input into MPT models.

Lukomnik and Hawley (2021) look at the impact the portfolio selection undertaken by large institutional investors has on systemic risk and conclude that MPT also does not capture investment influences which relate to corporate governance such as ESG. They posit that institutional investors such as pension funds have a responsibility to promote the health and welfare of their clients, not only to maximise their monetary wealth. They note an increasing trend in the world's largest institutional investors to seek to influence corporate governance positively. They argue that portfolio risk cannot be diversified away by selection of portfolio constituents through the use of modern portfolio theory, but rather that, in an investment market dominated by large institutions, such selection drives beta and systemic risk, leading to a feedback loop between portfolio risk management and systemic risk (Hawley & Lukomnik, 2018).

In combination, these shortfalls make MPT susceptible to market failures arising from policy, external influences and information asymmetry (Iyiola et al., 2012). The study does not propose a solution to the criticism that MPT fails to capture systemic events, since the alternative risk measures investigated in the study are, similarly to variance, based on historical data and are therefore not able to take ambiguous or black swan events into account. Similarly, the use of EWMAs to better forecast risk and return makes use of historical data and is therefore unable to address this criticism.

2.5.6 MPT ignores transaction costs

Another criticism of MPT is that its early formulations ignore transaction costs (Omisore et al., 2012) because of the additional computational complexity. Traditionally, portfolios were optimised without considering costs, since it was the function of the trading desk to control and manage these, leading to portfolios with large trading costs, with potentially severe impacts on returns (Kolm et al., 2014).

However, some empirical studies have shown a marked difference in the performance before and after taking account of transaction costs. Rigamonti and Lučivjanská (2022) note that their strategy results in high turnover and comensurately high fees, which they include as a fee of 10bp per trade after optimisation without fees, while Nguyen et al. (2018) similarly note that the weights of some of the portfolios which they generate change materially over time, thereby incurring considerable transaction costs.

More modern empirical research, assisted by more tractable computing capabilities, allow for the inclusion of fees. Kolm et al. (2014) note that it is common to amend the MPT framework with various constraints, including the costs of trading, with an example provided by Nguyen et al. (2018), who take the approach of incorporating transaction costs into the optimisation function rather than simply including the effect of fees after optimisation.

2.6 RISK MEASURES

Some of the criticisms of MPT may be addressed through optimisation of risk and return using different risk measures to variance. In particular, the criticisms relating to minimisation of upside risk and MPT's assumption of normally distributed returns can be addressed through the use of non-parametric risk measures, and asymmetric risk measures (Boasson et al., 2011; Byrne & Lee, 2004; Hunjra et al., 2020; Liu et al., 2019). This section discusses these different risk measures and the literature relating to their use in MPT.

Risk measures assign values to random variables, typically interpreted as losses. This aggregation results in a loss of information, but also allows investors' preferences to be ordered relative to each other, thereby enabling tractable optimisation in portfolio theory (Rachev et al., 2008).

Although variance is widely used in portfolio optimisation (Byrne & Lee, 2004), there is nonetheless a widely-held view that it is not an ideal risk measure, not least because it also results in the minimisation of upside risk (Boasson et al., 2011; Byrne & Lee, 2004; Hunjra et al., 2020; Liu et al., 2019). Affirming this, studies of portfolio optimisation have been carried out using a variety of risk measures. Examples include studies of portfolio optimisation using downside semivariance (Boasson et al., 2011; Byrne & Lee, 2004; Hunjra et al., 2020; Lohre, 2010; Sortino, 1994), value at risk and Conditional value at risk (Gaivoronski, 2005; Hunjra et al., 2020; Lohre, 2010; Stanković et al., 2020), Mean-Absolute Deviation (Hunjra et al., 2020; Stanković et al., 2020) and Semi-absolute Deviation (Stanković et al., 2020).

Different risk measures have different desirable characteristics, including (Rachev et al., 2008) asymmetry of the risk probability distribution and the concomitant ability to take into account downside risk, aggregated risk, being the ability to take into account different risks, transaction costs, computational complexity, investor risk aversion, inter-temporal dependence, whereby current investor utility is a function of the investor's prior utility, non-linearity, whereby change in the risk measure output need not be a linear function of its inputs, and correlation and diversification. The complexity of comparing different risk measures is alluded to by Ramos et al. (2023) who derived linear optimisation models for several different risk measures in order to do so.

The ten risk measures used to optimise portfolios in this study are discussed in the remainder of this section.

2.6.1 Portfolio variance or standard deviation

Portfolio variance measures the likelihood of deviations in the portfolio return away from the expected return (Bodie et al., 2017) and is the risk measure utilised by Markowitz in his seminal paper *Portfolio Selection* in 1952. Standard deviation (and therefore variance) remains the most popular risk measure, despite being a measure of uncertainty rather than of risk (Rachev et al., 2008). Variance has certain advantages over other risk measures, including cost, convenience and familiarity (Markowitz & Cowles, 1959) where cost is a proxy for computational complexity. It is non-linear (that is, $Var(aX) \neq aVar(X)$ for random distribution X) and can take into account correlation and diversification and can be used intertemporally (Rachev et al., 2008). However, it is symmetrical and consequently fails to differentiate between upside and downside risk (Boasson et al., 2011).

2.6.2 Lower partial moments and lower partial standard deviation

Lower partial standard deviation (or the analogous downside semivariance) is part of the class of lower partial moment (LPM) risk measures. An LPM takes into account only those returns which perform more poorly than a specified target (Lohre, 2010). Accordingly, optimising a portfolio with respect to an LPM risk measure addresses the criticisms or MPT relating to minimisation of upside risk and MPT's assumption of normally-distributed returns.

In particular, the second lower partial moment with a target return equal to the expected portfolio return is the downside semivariance of returns (Lohre, 2010). Accordingly, it may be calculated similarly to portfolio standard deviation or variance, as a measure of deviation from the mean, but observations where the return exceeds the mean return are excluded. Other target returns used include the risk-free rate and a zero return (Lohre, 2010).

Lower partial standard deviation is not as widely used in portfolio management as volatility (Gosier, 2005). However, as with variance, it is non-linear; can take into account correlation and diversification; and can be used intertemporally. Unlike variance, it is asymmetrical and able to differentiate between upside and downside risk. Using lower partial standard deviation has the advantage of ignoring the risk of upside performance (captured in the portfolio variance) and quantifying only downside performance risk (Bodie et al., 2017). This risk measure is particularly appropriate if the portfolio manager believes returns are skewed, and that the assumption of normality of returns is therefore unrealistic (Bodie et al., 2017), since the risk measure will be asymmetric about the mean return and minimising downside risk will not necessarily minimise upside risk as well. Downside semi deviation captures risk commensurately with semi deviation when returns are symmetrical and the benchmark is the mean return, while improving thereon when returns are asymmetric or the benchmark is any return other than the mean (Estrade, 2006).

One of the challenges associated with constructing mean-semivariance portfolios (hereafter, semivariance-optimised portfolios) is that the estimates for the expected return and the deviations from the mean for each of the portfolio constituents are more affected by estimation errors than when constructing mean-variance portfolios (hereafter variance-optimised portfolios) (Rigamonti & Lučivjanská, 2022) since the incidence of downside deviations, used for the calculation of the semivariance, will be lower than that of both downside and upside deviations, used for the calculation of variance. This challenge may be overcome by applying principal components analysis using minimum average partial on the downside correlation matrix used to estimate the semivariance of the portfolio. Rigamonti & Lučivjanská (2022) demonstrate the efficacy of this method, applying it to a 15-year data sample of stocks from the Dow Jones industrial average, the NASDAQ 100 and the S&P 100. Their out-of-sample tests show their adjusted semivariance-optimised portfolios outperforming unadjusted mean-semivariance portfolios and variance-optimised portfolios before fees. However, their investment strategy trades significantly more often than the unadjusted strategies, with the result that after fees, the performance of the various strategies is broadly aligned.

Lohr et al. (2010) conducted a study of the efficacy of portfolio construction using various downside measures, including downside semivariance, on a sample from the Dow Jones Euro Stoxx 50 between 1992 and 2009. They found that portfolio optimisation using downside semideviation and downside semivariance resulted in portfolios with lower downside semideviation and downside semivariance than their benchmark portfolio, but not lower than a portfolio optimised using volatility, while the total returns over the out-of-sample period were very similar for these different portfolios, with both outperforming the benchmark over the period.

Stanković et al. (2020) concluded that a variance-optimised portfolio outperformed a semivariance-optimised portfolio on the Belgrade Stock Exchange over a 250-day study, out of sample, with the variance-optimised portfolio achieving a return of 3.3%, as compared to 2.19% achieved by the semivariance-optimised portfolio. However, the Sharpe ratio of the semivariance-optimised portfolio was roughly 65% of that of the variance-optimised portfolio.

In contrast, Rocha (2016) observed that variance-optimised portfolios resulted in lower returns than semivariance-optimised portfolios in most cases, while Lam et al. (2010) observed similar returns between variance-optimised portfolios and semivariance-optimised portfolios based on stocks on the Kuala Lumpur stock exchange, but considerably lower risk on the semivariance-optimised portfolio.

Hunjra et al. (2020) conducted a study on the Pakistan, Bombay and Dhaka stock exchanges for the period 2003 to 2015, wherein they compared the performance of, *inter alia*, variance-optimised portfolios as compared to semivariance-optimised portfolios across growth, crisis and recovery markets. The identification of each stage of the market, as well as the optimisation of the portfolios was conducted in sample, and a separate portfolio optimisation was undertaken for each market stage and each country. The comparison of the two portfolio optimisation risk factors was inconclusive, with the variance-optimised portfolio outperforming the semivariance portfolio in some geographies and market stages in terms of both risk and return, but not in others.

In conclusion, despite contentions that portfolios constructed using downside semivariance tend to perform better than those constructed using variance (Markowitz & Cowles, 1959), empirical results are varied. Examination of the efficient frontiers of a semivariance-optimised portfolio and a variance-optimised portfolio constructed using the same assets show that a semivariance-optimised portfolio can result in lower downside risk than the commensurate variance-optimised portfolio for the same expected return (Boasson et al., 2011; Stanković et al., 2020), as shown in Figure 2.2.

Figure 2.2

Variance-optimised vs semideviation-optimised efficient frontiers

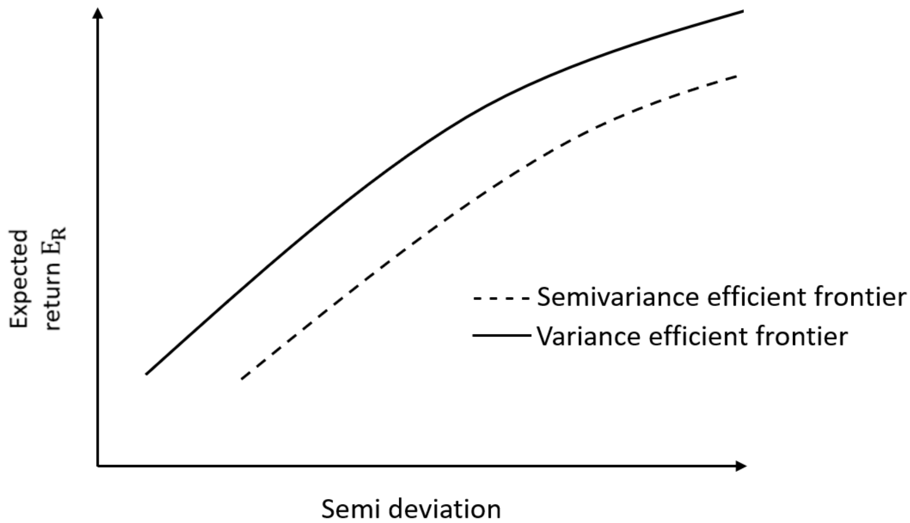

Note. Graphical depiction of the statement by Boasson et al that constructing an efficient frontier using downside semivariance will result in lower downside risk than a variance-optimised portfolio for the same level of returns.
Source: Boasson et al., 2011

However, some studies showed better aggregate returns for portfolios constructed using variance as a risk measure than those constructed with downside semivariance (Jacobsen, 2005; Stanković et al., 2020), while other studies showed that better performance between semivariance-optimised portfolios and variance-optimised portfolios depended upon the geography and state of the market being tested (Hunjra et al., 2020).

2.6.3 Mean absolute deviation and semi-absolute deviation

Mean absolute deviation (MAD) is the average of the absolute values of the deviations from the mean, while lower semi-absolute deviation (SAD) is to MAD as downside semivariance is to variance, including in the measure only the absolute deviations from the mean return which are lower than the mean (Stanković et al., 2020). SAD is a similar measure to LPM(1,*E(R)*), since both sum un-squared observations below the mean return, albeit the former sums the absolute value of the deviations, while the latter sums the signed deviations, all of which will be negative, given the conditionality.

Optimisation of portfolios using SAD addresses the criticisms of MPT relating to minimisation of upside risk and MPT's assumption of normally-distributed returns, while the latter criticism is also addressed by optimisation using MAD. Similarly to the comparison of variance and downside semivariance, MAD includes upside deviations from the mean in the risk measure, while SAD does not.

MAD-optimised portfolios are more consistent with maximisation of expected utility than variance-optimised portfolios (Chen et al., 2022). MAD is less sensitive to outliers than variance (Byrne & Lee, 2004), since the deviations from the mean are not being squared. It is also computationally easier to optimise a portfolio using MAD than variance, while producing similar (Byrne & Lee, 2004) or more efficient (Hunjra et al., 2020) portfolio weights and trade-offs of risk and return. While some studies have shown that MAD-optimised portfolios sometimes result in lower returns variance (Hunjra et al., 2020), others show a contrary result (Silva et al., 2017). The use of MAD in portfolio optimisation may, however, result in a loss of information available in the asset returns (Byrne & Lee, 2004).

A challenge associated with MAD-optimised models is that of estimation error; the out-of-sample distribution of returns often differs substantially from the distribution of in-sample returns. Chen et al. addressed this issue by applying distributionally-robust methodologies to optimising MAD portfolios. In their out-of-sample tests on weekly data, a classic MAD-optimised portfolio outperformed a variance-optimised portfolio by a small margin, at the cost of slightly higher risk, while their distributionally robust MAD-optimised portfolio significantly outperformed both as well as a uniformly-allocated benchmark portfolio. Testing on a daily data sample indicated poor performance of the MAD- and variance-optimised portfolios relative to the uniformly-allocated portfolio, due to conservative treatment of high volatility in the out-of-sample period by the risk-averse optimisation methods (Chen et al., 2022).

Silva et al. (2017) investigated the efficacy of, *inter alia*, MAD-optimised portfolios in highly volatile financial crisis periods, using data from the Brazilian Stock Exchange between 2004 and 2013. They found that MAD-optimised models outperformed variance-optimised models in four of the six test portfolios they created, but at the cost of higher risk (Silva et al., 2017).

Stanković et al. (2020) concluded that MAD- and SAD-optimised portfolios substantially outperformed a variance-optimised portfolio on the Belgrade Stock Exchange over a 250-day study, out of sample, with the variance-optimised portfolio achieving a return of 3.3%, as compared to 12.87% achieved by the MAD- and SAD-optimised portfolios. The MAD- and SAD-optimised portfolios resulted in much smaller numbers of shares being held (five) compared to the variance-optimised portfolio (27). The additional returns came at the cost of very slightly higher additional risk, but with the Sharpe ratio of the MAD- and SAD-optimised portfolios circa four times that of the variance-optimised portfolio due to the significantly higher returns.

In Hunjra et al.'s (2020) in-sample study on the Pakistan, Bombay and Dhaka stock exchanges for the period 2003 to 2015, MAD-optimised portfolios outperformed variance-optimised portfolios across all three exchanges in growth, crisis and recovery markets in all instances except one. The study did not show any clear outperformance between MAD-optimised portfolios and variance-optimised portfolios with respect to risk, using variance as a risk measure.

Jacobsen (2005) compared the performance of portfolios constructed of two assets, being the Standard & Poor's 500 and the Dow Jones Two-Year Corporate Bond Index, allocating between the two indices daily based on the past 60 days of trading history. A portfolio constructed using the first lower partial moment centred about a zero return (i.e LPM(1,0)) outperformed a variance-optimised portfolio for the sample period between 1996 and 2005 by a relatively narrow margin, while showing fractionally higher risk.

Kasenbacher et al. (2017) found that mean absolute deviation optimised portfolios outperformed variance optimised portfolios based on historical data from the S&P 500, with the mean absolute deviation optimised portfolio also displaying higher Sharpe ratios.

In summary, empirical studies comparing MAD-optimised portfolios to variance-optimised portfolios have typically reported higher total returns (Hunjra et al., 2020; Stanković et al., 2020; Chen et al., 2022). Similarly, empirical studies of SAD-optimised portfolios report higher total returns than variance-optimised portfolios (Stanković et al., 2020), and empirical studies comparing the performance of an LPM(1,0) report total return exceeding that of variance-optimised portfolios (Jacobsen, 2005), albeit noting that the study uses a deviation from zero, rather than the mean return.

2.6.4 Value at risk

The value at risk (VaR) is the n^{th}-percentile lowest return over a period, where n is a risk parameter (Lohre, 2010). Consequently, assuming past experience is a reasonable estimator for future returns, there will be an approximately 1-n% probability that the returns over each successive period will exceed the n%-VaR. VaR Is widely used by banks and financial regulators to calculate the appropriate amount of capital to set aside against operational losses (Lwin et al., 2017). Optimisation of portfolios using VaR addresses the criticisms of MPT relating to minimisation of upside risk and MPT's assumption of normally-distributed returns.

There are three main methods of calculating VaR, being the nonparametric historical simulation approach, the parametric approach which typically fits an assumed distribution to the returns and the Monte Carlo simulation approach. The historical simulation approach is that which is most commonly followed by major firms (Lwin et al., 2017).

Efficient frontiers calculated using variance provide a poor approximation for efficient frontiers calculated using VaR as the risk measure, with the consequence that investors wishing to minimise VaR should optimise over mean-VaR, rather than using other risk measures (Gaivoronski, 2005). However, the number of nonparametric VaR portfolio optimisation studies is small (Lwin et al., 2017). VaR has the advantage that it can be estimated empirically and does not rely on any assumption of normality of returns (Bodie et al., 2017).

VaR also has disadvantages as a risk measure. In contrast to portfolio variance, it violates the subadditivity requirement for a coherent risk measure: diversifying a portfolio does not necessarily reduce the VaR, and VaR does not provide any information about the extent of losses past the specified threshold. Portfolios optimised using VaR are also sensitive to the confidence level used (Boasson et al., 2011), to the extent that, for very wide confidence intervals, a minimum-VaR portfolio may not exist (Alexander, 2002). Another problem with using VaR to optimise portfolios is that it is not smooth (that is, the efficient frontier between expected returns and VaR is not smooth) and can result in local minima, making mean-VaR portfolio optimisation computationally intractable (Gaivoronski, 2005, Lwin et al., 2017).

The returns on Jacobsen's (2005) two-index portfolio were slightly higher when optimised using standard deviation than VaR for the sample period between 1996 and 2005, while the VaR-optimised portfolio showed slightly lower risk.

Andreu et al. (2009) also observed that VaR-optimised minimum-risk indices exhibited lower risk than equivalent low-risk indices in the Spanish, United States and Argentinian stock markets, with higher returns than the indices in the Spanish and Argentinian markets.

Stanković et al. (2020) concluded that a VaR-optimised portfolio underperformed a variance-optimised portfolio on the Belgrade Stock Exchange over a 250-day study, out of sample, with the variance-optimised portfolio achieving a return of 3.3%, as compared to 2.68% achieved by the VaR-optimised portfolio. The VaR-optimised portfolios resulted in slightly smaller numbers of shares being held (23) compared to the variance-optimised portfolio (27). The variance-optimised portfolio also had a lower risk than the VaR-optimised portfolio, with the Sharpe ratio of the VaR-optimised portfolios being circa 80% of that of the variance-optimised portfolio.

In conclusion, empirical studies comparing portfolios optimised using VaR to those optimised using variance have reported mixed results, although the difficulty in comparing studies because of differences in VaR parameters between different studies is acknowledged. Some studies have resulted in broadly similar aggregate returns (Jacobsen, 2005), while others resulted in poorer mean-VaR returns (Stanković et al., 2020).

2.6.5 Conditional value at risk

Where VaR is defined as the n^{th}-percentile lowest return over a period, Conditional value at risk (CVaR) is the expected loss given the return is below the n^{th}-percentile lowest return over a period (Lohre, 2010). As with VaR-optimised portfolios, optimisation of portfolios using CVaR addresses the criticisms of MPT relating to minimisation of upside risk and MPT's assumption of normally-distributed returns.

Unlike VaR, CVaR is a coherent risk measure, satisfying the subadditivity requirement which VaR does not (Lim et al., 2011). Optimising a portfolio with respect to CVaR is more effective than doing so with respect to VaR (Lohre, 2010), and usually lead to well-optimised portfolios in terms of VaR; that is, CVaR-optimised portfolios should have low VaR as well (Stanković et al., 2020, Di Clemente, 2002). This is because the CVaR will never exceed the VaR which suggests that portfolios with low CVaR must necessarily have low VaR as well (Rockafellar & Uryasev, 2000).

CVaR-optimised portfolios may be appropriate when the underlying assets are not symmetrically distributed, as in the case of put and call option portfolios (Dao, 2014) and in distressed markets (Nguyen et al., 2018).

As with VaR, optimising a portfolio with respect to CVaR has certain disadvantages and challenges. It similarly requires a considerable amount of data to calculate, particularly for wider confidence intervals, or alternatively an assumption about the distribution of returns (Boasson et al., 2011; Lohre, 2010). Again as with VaR, CVaR-optimised portfolios are sensitive to the confidence interval used: Nguyen et al. report poorer performance of CVaR-optimised portfolios where historical data is used as input to the optimisation problem and the confidence level is very high (Nguyen et al., 2018). Hunjra *et al.* (2020) reported dramatically different performance between CVaR-optimised portfolios using a 95% vs 99% confidence interval, with significantly different returns and no outright winner between the two across three geographic markets and three market states. More specifically, optimising using CVaR with a 99% confidence interval resulted in higher returns relative to optimisation with a 95% confidence interval for some geographic markets in crisis, but lower for others. Optimisation with a 95% confidence interval resulted in higher returns for markets in recovery or growth stages across all three geographic markets investigated. Use of a 99% confidence interval resulted in lower variance for markets in recovery stages, but higher variance for markets in crisis or growth stages, across all geographic markets.

Another possible disadvantage of CVaR-optimised portfolios is an increased trading frequency, resulting in higher fees and commensurately lower returns (Nguyen et al., 2018).

Hafsa (2015) undertook a study using 20 stocks from the French SBF250 market index between 2005 and 2009, comparing variance optimised portfolios to CVaR- optimised portfolios. The study noted that the return distributions of the 20 stocks examined were significantly non normal. The study concluded that in most months the monthly return of the CVaR-optimised portfolio exceeded that of the variance-optimised portfolio.

Nguyen et al. (2018) conducted a study on 23 different selection criteria of stocks from the NYSE, AMEX and NASDAQ based on different criteria. For each stock selection criteria, they compared performance of variance-optimised against CVaR-optimised portfolios, finding that 90% CVaR-optimised portfolios outperformed variance-optimised portfolios in approximately 93% of the selection criteria, ranked by Sharpe ratio. After adding the impact of trading fees, the percentage of CVaR-optimised portfolios which outperformed variance-optimised portfolios decreased to circa 78%.

In summary, empirical studies comparing portfolios optimised using CVaR to those optimised using variance also reported mixed results, with some CVaR-optimised portfolios outperforming variance-optimised portfolios (Lohre, 2010) and other studies showing CVaR-optimised portfolios outperforming variance-optimised portfolios in some instances, but not in others (Hunjra et al., 2020; Nguyen et al., 2018; Hafsa, 2015).

2.7 CONCLUSION

Modern Portfolio Theory allows the investor to identify an efficient frontier of portfolios, each of which, for each given expected return, has a minimum risk; or equivalently, for each given level of risk, has the greatest possible expected return (Markowitz, 1952). The inclusion of a risk-free asset into the hypothetical investment universe allows the investor to improve his risk or expected return by investing into a combination of the risk-free asset and the optimal portfolio, being the portfolio at the point at which the Capital Market Line touches the efficient frontier (Sharpe, 1964).

MPT is the subject of a number of criticisms. Amongst these, MPT assumes the probability distribution of returns is constant through time; it assumes risk and return persist into the future; it minimises upside variance; it assumes returns are normally distributed; it fails to capture systemic events; it ignores transaction costs; and it assumes investor behaviour is rational.

The criticisms relating to the assumption of normal returns and minimisation of upside variance have been addressed through variations of MPT's portfolio optimisation methodologies using alternative risk measures. These alternative risk measures have included, inter alia, downside semivariance, value at risk and Conditional value at risk, Mean-Absolute Deviation and Semi-absolute Deviation. The construction of portfolios using each risk measure has advantages and disadvantages relative to that carried out by Markowitz using variance as a risk measure.

Typically, portfolios are constructed using past performance in the form of expected returns and various risk measures with the assumption that these measures will persist into the future, thereby retaining the optimal ratio of risk and return desirable to the investor, evident in the portfolio at the date it is constructed. However, risk is known to be heteroskedastic. Consequently, the persistence and heteroskedasticity of risk must be balanced against each other to obtain a risk measure which achieves this balance.

In addition to balancing persistence and heteroskedasticity, the risk measures used to construct the optimal portfolio must also take into account the period of the available returns data and the frequency with which the portfolio weights are reset.

In general, using a shorter sample interval to construct a forecast of a risk metric will result in more robust forecasts; and resetting portfolio weights periodically leads to improved performance. Application thereof goes some way towards mitigating the criticisms that MPT assumes the probability distribution of returns is constant through time and that MPT assumes risk and return persist into the future.

CHAPTER 3 - RESEARCH DESIGN AND METHODS

3.1 INTRODUCTION

The preceding chapter discussed the available literature relating to MPT and described how it fits into broader Portfolio Theory. It went on to discuss the literature on the risk measures which will be used in the study.

This chapter discusses how the research will be carried out. In particular, it details how each risk measure will be calculated; how the EWMA of each risk measure will be calculated; how each risk measure will be used to construct an optimal portfolio and the concomitant investment strategy; what performance measures will be used to measure each investment strategy; what data will be used in the study; how it will be collected; and how and with which tools the data will be processed.

3.2 RESEARCH PARADIGM / PHILOSOPHY

The study is based on a positivist research paradigm in terms of which the objective efficacy of different risk measures in portfolio construction will be investigated (Abdul Rehman & Alharthi, 2016). This study will undertake quantitative research which is characterised by the view that a single reality can be measured using objective scientific principles by researchers who are separated from their research subjects (Soiferman, 2010). It will be based on publicly-available data, which will allow it to be replicated and verified.

3.3 DESCRIPTION OF INQUIRY STRATEGY AND BROAD RESEARCH DESIGN

Time series data is defined as "a collection of observations on a variable that is measured over time at regular intervals" (Montgomery et al., 2015). Accordingly, the study will make use of an empirical enquiry strategy based on numeric time series data, specifically financial time series data collated by a third party. The data will be analysed using a combination of statistical and numerical methods. The type of data and the combination of statistical and numerical methods proposed for the study are consistent with those used in similar studies (Boasson et al., 2011; Byrne & Lee, 2004; Gaivoronski, 2005; Hunjra et al., 2020; Lohre, 2010; Sortino, 1994; Stanković et al., 2020).

3.4 SAMPLING AND DATA COLLECTION

The study sample comprises daily index data for a 10-year period for the Standard and Poor's (S&P) Global 1200 sector indices, between 29 February 2012 and 25 March 2022. Accordingly, the data will be time series, quantitative data, with the indices calculated in US Dollars. The sample length covers multiple states of the market, including periods of both low- and high volatility, stable markets, bull- and bear markets. In addition, the sample period covers the CoVID19 crisis period. The indices aggregate stock returns over many geographical areas, thereby reducing the possibility of bias due to the peculiarities of a specific geographical market.

The sectors and the constituents thereof represented in the S&P Global 1200 sector indices are defined per the top (most aggregated) tier of MSCI's Global Industry Classification Standard (GICS) methodology. These constitute Energy, Materials, Industrials, Consumer Discretionary, Consumer Staples, Health Care, Financials, Information Technology, Communication Services, Utilities and Real Estate (Morgan Stanley Capital International, 2020), although the Real Estate index is not included in the primary sample due to lack of data availability. The S&P Global 1200 index will be used as a benchmark against which the constructed portfolios will be measured. The S&P Global 1200 covers circa 70% of world market capitalisation. The S&P Global 1200 indices data for the study was downloaded directly from its creator, S&P. The data for each index is provided in Excel files and has been stored in source form.

Since, unlike logarithmic returns, discrete simple portfolio returns are a linear function of the weights and simple returns of the constituent assets (Miskolczi, 2017), the study will use simple returns, calculated by dividing the index values in successive periods and subtracting one. This has the advantage that the partial derivative of the portfolio's simple return with respect to the weighting of each constituent asset is a function of that asset only, while this is not the case using logarithmic returns. Further, logarithmic returns, which hold the benefit of closer-to-normally-distributed returns, are not required because the study uses non-parametric methods to calculate sample risk measures.

Each index is calculated by S&P in real time, but rebalanced quarterly on a float-adjusted market capitalisation weighted basis (Standard & Poor's, 2021). While the index provides both price and total return tickers, the study will take its sample from the total return index data.

The ten indices each have investable ETFs based on them, meaning that any portfolio optimisation strategy developed may be implemented as an investment strategy. Note that it is not required that the ETFs display the same risk characteristics as their constituents. The portfolio optimisation takes place based on the risk and return of the ETFs (represented by the indices) themselves.

To corroborate the results of the primary analysis, the study will make use of a sample of the MSCI daily Industry Group- and Industry-level index data available from Refinitiv Eikon. Accordingly, the data will be time series data, with the indices calculated in US Dollars.

The MSCI daily industry-level data comprises indices for 69 industries grouped into 24 Industry Groups and the 10 sectors in the primary sample data. The GICS classifications broken down into Sectors, Industry Groups and Industries can be found in Appendix A.
The industry-level indices do not all have investable ETFs, and as such, their use is not suitable for building an investable portfolio.

Since the portfolio optimisation is carried out at the high-level sectors, and each sector index already represents a level of diversification across many companies' shares, the study uses the indices for the 24 Industry Groups and the 69 Industries developed by MSCI at the next two levels down in the classification system to (partially or fully) replicate and corroborate the results. This is intended to establish whether the results of the optimisation carried out are pervasive across less aggregated or diversified assets. The portfolios created by this analysis are not tradeable, since ETFs do not exist to cover all the indices for the Industry Groups and Industries. Where the Industry Groups or Industries indices have repeating values for in excess of fifty periods, these indices are not used in the analysis, on the basis that indices with incomplete data will not be used in the study. This excludes two of the 69 Industry Indices (Diversified Consumer Services and Thrifts and Mortgage Finance).

The increase in the number of input indices increases run times considerably. Consequently, while daily, weekly and monthly runs are conducted in the Sector indices, only weekly and monthly runs are conducted for the Industry Group and Industry input strategies.

Note further than a direct comparison of the returns of the strategies using the S&P indices as inputs and the MSCI indices as inputs is problematic, since the S&P indices are total return indices (that is, they include reinvested dividends) whereas the MSCI indices are price indices (that is, they track only share prices and do not take account of dividends). However, a comparison of each strategy to their appropriate benchmarks allows indirect comparison of their efficacy.

Descriptive statistics for the data are as follows:

Table 3.1

Descriptive statistics for the primary data sample

	Energy	*Financials*	*Consumer Staples*	*Consumer Discretionary*	*Health Care*	*Industrials*	*IT*	*Materials*	*Utilities*	*Communications*
Mean	0.0002	0.0004	0.0004	0.0005	0.0005	0.0004	0.0007	0.0003	0.0004	0.0003
Standard Error	0.0003	0.0002	0.0001	0.0002	0.0002	0.0002	0.0002	0.0002	0.0002	0.0002
Median	0.0002	0.0007	0.0005	0.0009	0.0007	0.0006	0.0011	0.0004	0.0008	0.0005
Standard Deviation	0.0151	0.0109	0.0072	0.0095	0.0086	0.0094	0.0115	0.0105	0.0091	0.0089
Sample Variance	0.0002	0.0001	0.0001	0.0001	0.0001	0.0001	0.0001	0.0001	0.0001	0.0001
Kurtosis	25.1026	18.8054	13.9923	12.6219	9.3263	17.1441	12.9207	10.5556	24.3023	10.8187
Skewness	- 0.7006	- 0.7877	- 0.6956	- 0.7324	- 0.3732	- 0.5546	- 0.4020	- 0.4474	- 0.4734	- 0.6229
Range	0.3597	0.2222	0.1420	0.1761	0.1428	0.2004	0.2211	0.2041	0.2068	0.1473
Minimum	- 0.1917	- 0.1096	- 0.0852	- 0.0931	- 0.0777	- 0.0993	- 0.1203	- 0.1015	- 0.1115	- 0.0863
Maximum	0.1680	0.1126	0.0567	0.0830	0.0652	0.1011	0.1008	0.1026	0.0953	0.0610
Sum	0.6315	1.0790	0.9442	1.2751	1.4224	1.0778	1.8888	0.8445	0.9925	0.8598
Count	2606	2606	2606	2606	2606	2606	2606	2606	2606	2606

Note. Descriptive statistics for the ten sector indices used in the study.

Source: Author’s calculations

The data covers ten indices, each of which provides a daily end-of-day index value (for business days). Table 3.1 shows descriptive statistics for the daily returns calculated from these index values. The mean return per workday ranges from 2.4 basis points for the Energy sector to 7.3 basis points for the IT sector. Consumer Staples shows the lowest standard deviation of daily returns at 0.7%, while Energy shows the highest standard deviation at 1.5%, also evident in Energy having the smallest and largest minima and maxima respectively. The distribution of returns for Energy and Utilities display the heaviest tails, with kurtosis of 25 and 24 respectively, while Healthcare and Materials show lighter tails with kurtosis values of nine and 10.5 respectively. Energy, Financials and Consumer Discretionary are the most skewed to negative returns, while Healthcare and IT are the least skewed, in line with these two sectors having the highest mean return.

All three of the Shapiro-Wilk test, the D’Agostino K^2 test and the Anderson-Darling test for normalcy reject the null hypothesis that the daily returns are distributed normally for each index. The absence of normally-distributed returns has little impact on the study, since all of the risk measures in the study, with the exception of the EWMA of value at risk and the EWMA of conditional value at risk, are calculated using non-parametric methods,

The performance of the sector indices over the study period is displayed in Figure 3.1. As is evident from Table 3.1, the IT sector index shows the greatest return over the study period, with Energy showing the lowest return.

Figure 3.1

S&P Global 1200 Sector Indices based from 1.0 at start of study

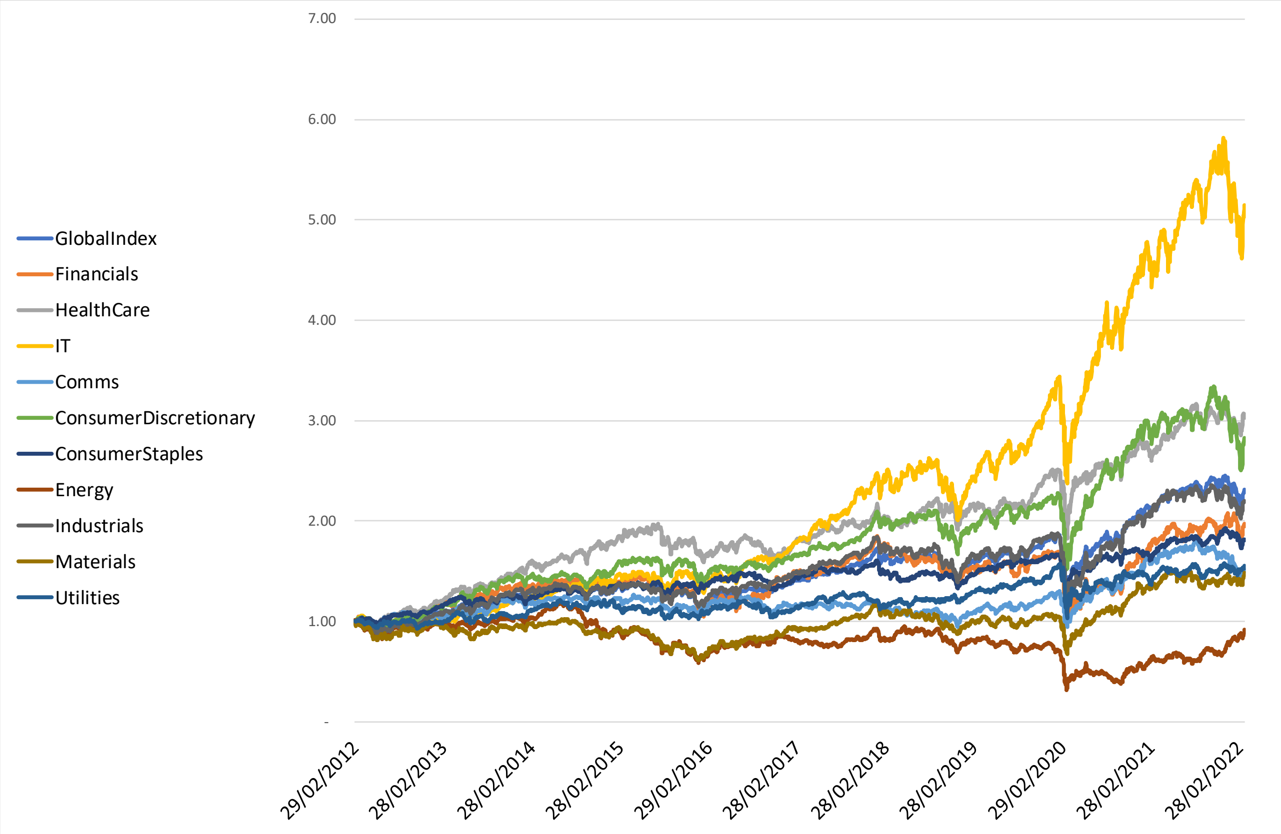

Note. Graph of the S&P Global 1200 Sector Indices performance over the study period. Source: S&P, 2022.

In summary, all sector-level results tabulated in Section 4 are based on the S&P Global 1200 sector index data, while all industry group- and industry-level results are based on the MSCI industry group and industry index data.

3.5 DATA ANALYSIS

This subsection describes how the index data was analysed. Existing research relating to the effect of sample period length and recalculation period is discussed in the context of the study's data. This is followed by a description of how data was partitioned to calculate successive portfolios in the study and what results were calculated. Thereafter, each risk measure used in the study is discussed, followed by discussion of existing research relating to forecasting risk and how this is applied in the study. Finally, the numerical methods algorithm used to calculate each successive optimal portfolio is described.

3.5.1 Sampling periods and rebalancing

Volatility forecasts are typically calculated using data sampled at the same interval as the forecast horizon. However, the use of low-frequency sample data can introduce stale information into the forecast. For example, monthly risk forecasts typically make use of 60 sample periods, such that five-year-old information will have an influence on the current forecast. While this may be mitigated by weighting newer observations more heavily than older observations (Gosier, 2005), empirical studies of the persistence of volatility imply that five-year-old data provides very little information about present volatility. Engle (2001) estimated the half-life of volatility on the Dow Jones Industrial Index, being the period over which volatility will, on average, move half-way back from its current level towards its long-term mean, to be 73 days (Engle & Patton, 2001), implying that a five-year-old observation will have only a 0.000003% impact on a volatility forecast.

Another problem with the use of long sample intervals is that information suggested by intra-period returns is lost. The use of shorter sample intervals makes risk measures more robust, even for forecasts with a horizon longer than the sample interval (Gosier, 2005).

The study differences the index data at different time steps to derive daily (differencing each row), weekly (differencing every fifth row) and monthly simple returns (differencing every twenty-second row), thereby enabling comparison of the performance between portfolios created using each timestep. To obtain the 22-day monthly period, the number of daily observations in the 10-year sample is divided by 120 months to obtain 21.84 days per month.

Related to the sample interval is the amount of time elapsed between successive recalculation of the optimal portfolio weights and resetting thereof. An investor can modify his portfolio utilising the same optimisation strategy, by either recalculating the optimal portfolio in successive time periods and modifying the weights of the assets (reweighting), or by rebalancing the weights, which will have changed as a result of the difference in returns between the assets since the last rebalancing, back to the original weights, referred to as rebalancing. A portfolio cannot be rebalanced continuously (Maeso & Martellini, 2020), and studies as to whether rebalancing results in an increased return, that is a rebalancing premium, are limited by, *inter alia*, transaction costs, the rebalancing period and initial weighting schemes (Maeso & Martellini, 2020). Nonetheless, Maeso and Martellini (2020) concluded that periodic rebalancing did indeed result in a positive rebalancing premium.

Recalculating the optimal portfolio weights periodically addresses the criticism of MPT that it assumes that the probability distributions of asset returns is static in time, since each successive optimisation allows the assumption of a new probability distribution.

The study splits the returns data into rolling windows of three months each. Each window overlaps the previous window in the first *n*-1 of its *n* daily periods, as shown in Figure 3.5.11.

Figure 3.2

Illustration of rolling windows used to calculate each optimal portfolio

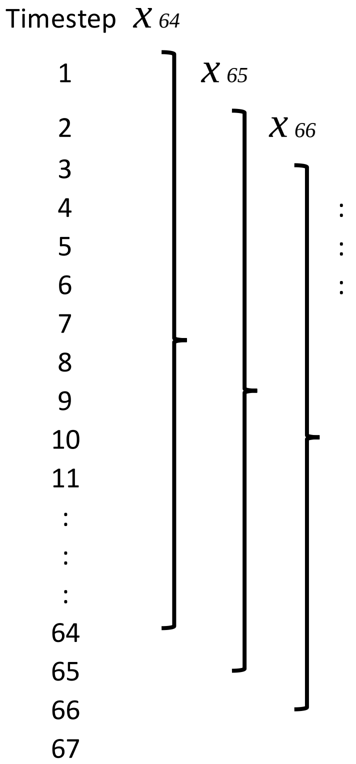

Note. Each portfolio is calculated using a 3-month (66 business days) rolling window of sector returns for each asset. For strategies recalculating portfolios daily, the rolling window is moved forward one day to calculate each subsequent portfolio.

Source: Author's calculations

Applying a set of asset weightings to each set of returns creates a notional portfolio which has a set of portfolio returns. The portfolio returns over the three-month period are in turn used to calculate the risk measures, expected returns, λ (the rate of decay for the EWMA), and the EWMAs of the risk measures. By varying the weights, the portfolio return, mean return and risk measure over the window will change, and may then be maximised or minimised by numerical methods.

Accordingly, expected [Daily/Weekly/Monthly] return for timestep $t + N$ =

$$\sum_{d=t+1}^{t+N} \frac{\sum_{i=1}^{X} x_i P_{i,d+n}}{\sum_{i=1}^{X} x_i P_{i,d}} \tag{1}$$

Where:

n = 1 for daily expected return; 5 for weekly expected return; 22 for monthly expected return
t = the starting time step
N = the number of days in the three-month estimation period (approximately 64)
X = the number of constituents assets
x_i = the portfolio weight for the i^{th} constituent asset
$P_{i,d}$ = the price of asset i on day d

And successive expected returns are calculated on the rolling windows as portrayed in Figure 3.1.

The asset weights which result in the optimal mean return and risk measure combination represent the asset allocation for the portfolio as at the last day in the rolling window.

The portfolio strategy being assessed is calculated by rolling the window forward a period at a time and repeating the optimisation at each timestep.

Calculation of the expected return, each risk measure and the EWMA of each risk measure is undertaken by a custom function which outputs an array consisting of the mean and the specified risk measure, having taken the rolling window and a set of asset weights as an input. The calculation of each risk measure can be corroborated by calculating it independently or manually in Excel based on the input data, and subsequent comparison to the function output. This allows the numerical methods-based optimisation of the expected return-risk measure ratio to call the function repeatedly, specifying the risk measure, the returns rolling window and the asset weights, thereby allowing it to be agnostic as to what risk measure it is using to optimise.

The numerical methods-based optimisation is also implemented as a custom function, which takes as an input the full returns data set, the desired timestep and the risk measure to optimise over. This allows optimisation across all timesteps and risk measures by looping through lists of each and calling the optimisation function with the relevant input parameters. The optimisation of the expected return-risk measure ratio is based on MPT, but is carried out with numerical methods, as this will allow the same procedure to be applied irrespective of the risk measure, and avoids, in most cases, the criticism of MPT that it assumes returns are normally distributed (Boasson et al., 2011; Stanković et al., 2020). This has the advantage of reducing complexity and ensure that the study is consistent in how it optimises across different risk measures, albeit noting that the function requires different hyperparameters for different risk measures.

The use of functions allows reuse of code at each stage in the analysis, thereby reducing the time and complexity of the required analysis.

Each portfolio's returns are graphed for visual inspection and have various measures calculated upon them. To assess the success of each portfolio strategy, the study calculates annualised returns, risk measures and Sharpe ratios using the returns of the portfolio and compare these measures across portfolios to establish to what extent it holds that a portfolio optimised using a particular risk measure results in low risk and/or high returns, relative to portfolios optimised using other risk measures. The same custom function used to calculate the risk measures for the portfolio construction stage is used to calculate the risk measures used in the portfolio returns analysis stage of the study. Sharpe ratios are calculated using the full study period.

A modular code design was chosen for the study in order to reduce development time and to reduce run time as many iterations of the portfolio optimisation process were carried out using different hyperparameters to assess the impact thereof. The modular code design allows this with minimal human intervention barring changes to input parameters, thereby saving time and allowing wider analysis.

3.5.2 Estimation of risk measures

Each risk measure is calculated using only information from the current and preceding time steps, so as to ensure that the study does not incorporate any forward-looking bias. That is, the portfolios calculated in the study are forecasts, calculated out of sample. For each risk measure, the study constructs portfolios using three recalculation intervals, being one day, one week and one month. Daily total return data is used to calculate the risk measures and returns for each, and each measure is scaled to accommodate the length of the interval, to the extent required. The sample period for each reweighting interval is set at three months and is consistent between different risk measures.

3.5.2.1 *Variance*

Where return data for the constituents of the portfolio are available, the portfolio variance can be calculated as follows (Bodie et al., 2017):

$$\sigma^2_p = \sum_{i=1}^{n} \sum_{j=1}^{n} w_i w_j Cov(r_i, r_j) \qquad (2)$$

Where:

- w_i and w_j are the portfolio weight of each asset
- $Cov_{i,j}$ is the covariance of the i^{th} and j^{th} assets.

Alternatively, portfolio variance may also be calculated directly from the returns of the portfolio:

$$\sigma^2_p = \frac{\sum_{t=1}^{n} (r_t - \bar{r})^2}{(n-1)} \qquad (3)$$

Where:

- r_t is the return on the portfolio at period t
- $\bar{r}$ is the average return on the portfolio from periods 1 to n.

The study uses the direct calculation.

3.5.2.2 *Downside semivariance*

Downside semivariance is the second LPM with parameters τ being the mean return over the sample period (Lohre, 2010):

$$\mathrm{LPM}_{\tau,k}(\mathrm{R}) = E\ (\tau - \mathrm{R})^{k}\ |\ \mathrm{R} < \tau).P(\mathrm{R} < \tau) \tag{4}$$

Where:

- k determines which lower partial moment is being calculated
- τ is the target return
- R is the return distribution of the portfolio

3.5.2.3 *Mean absolute deviation and semi-absolute deviation*

Mean absolute deviation is defined as follows (Byrne & Lee, 2004):

$$\mathrm{MAD} = \frac{1}{\mathrm{T}}\sum_{t=1}^{\mathrm{T}} |\mathrm{R_t} - E(\mathrm{R})| \tag{5}$$

Where:

- R is the return distribution of the portfolio

Semi-absolute deviation is similarly calculated but instead incorporates only those observations where $\mathrm{R_t} - \mathrm{E(R)} < 0$.

3.5.2.4 *Value at risk*

The value at risk (VaR) is the n^{th}-percentile lowest return over a period, where *n* is a risk parameter (Bodie et al., 2017). It is calculated by ranking the returns from lowest to highest in the sample and extracting the n^{th}-percentile return in the sample. The study uses parameters for *n* of the 5^{th} percentile and the 10^{th} percentile.

As noted by Lwin et al., there are three main methods of calculating VaR, being the nonparametric historical simulation approach, the parametric approach which typically fits an assumed distribution to the returns and the Monte Carlo simulation approach. The historical simulation approach is that which is most commonly followed by major firms (Lwin et al., 2017) and is the approach followed in this study for the calculation of raw VaR (and for all other raw risk measures used in the study).

However, since application of an EWMA to return observations *n* periods in the past will reduce them by a factor of $(1-\lambda)^n$, the EWMAs for VaR (and CVaR) use a parametric calculation approach with an assumption of Normal returns. The distribution is fitted using the EWMAs of the return and the variance of returns, and the corresponding VaR or CVaR calculated from the tail of the distribution.

3.5.2.5 *Conditional value at risk*

Conditional value at risk is defined as follows (Stanković et al., 2020):

$$\mathrm{CVaR}_{\alpha}(\mathrm{R}) = E\left(\mathrm{R}\middle|\mathrm{R} \leq \mathrm{VaR}_{\alpha}(\mathrm{R})\right) \tag{6}$$

Where:

- R is the portfolio return
- α is the confidence interval parameter, expressed as a percentage

CVaR is calculated by ranking the returns from lowest to highest in the sample and then averaging the returns below the α-percentile return in the sample. The study uses parameters for α corresponding to the 5^{th} percentile and the 10^{th} percentile.

3.5.3 Forecasting Risk Measures

The final criticism of MPT which is addressed in the study is that MPT assumes that risk and return persist into the future, in that it seeks to optimise the present risk and return of a portfolio in the hope that such optimisation will persist into the future. The study addresses this criticism by optimising forecast risk and forecast return. This section discusses the characteristics of risk measures which impact on how they may be forecast.

3.5.3.1 *Persistence and heteroskedasticity*

Volatility forecasts are crucial to portfolio construction (Gosier, 2005). Only when the measures used to determine the optimal portfolio are persistent (Maeso & Martellini, 2020) or predictable (Lohre, 2010) can investors benefit from optimising their portfolio with respect to those measures.

However, risk is heteroskedastic (Rachev et al., 2008), with the consequence that risk measures for the same portfolio will change over time. Time-varying volatility is commonly observed in financial returns (Fan & Lee, 2017), with volatility clustering into periods of high- and low volatility (Dachraoui, 2018).

The Autoregressive Conditional Heteroskedastic (ARCH) model, introduced by Engle in 1982, introduced a means for volatility to change over time (Engle, 1982).

Moreno and Olmeda discuss the use of time-varying models in MPT, in particular the ARCH model and the Generalised Autoregressive Conditional Heteroskedastic (GARCH) model, which allow the forecast of variance (or risk) to vary in time. They optimised the risk and return of portfolios comprising thirteen years of 35 MSCI country market index data, with their portfolio optimised using a GARCH variance forecasting model achieving marginally lower returns, but also marginally lower risk, than a standard homoscedastic variance-optimised portfolio. They concluded that there was little material difference between the use of a heteroskedastic vs a homoscedastic variance forecasting model for the purposes of optimising a portfolio using MPT (Moreno et al., 2005).

Nguyen et al. use an Autoregressive-GARCH (AR-GARCH) conditional location-scale model of monthly returns to calculate conditional means and covariance matrices as *inputs* into the calculation of CVaR-optimised portfolios. This overcomes the problem of stale data by updating variance and mean based upon the most recent observations. They compared performance of CVaR-optimised portfolios using AR-GARCH inputs to that of CVaR-optimised portfolios using raw inputs. The AR-GARCH input portfolios perform better than the raw input portfolios in approximately 80% of the selection criteria, ranked by Sharpe ratio. However, after adding the impact of trading fees, the percentage of AR-GARCH portfolios which outperformed raw input portfolios decreased to circa 46% (Nguyen et al., 2018).

3.5.3.2 *Exponentially-weighted moving averages of risk measures*

An exponentially-weighted moving average assigns exponentially smaller weights to observations further back in time by employing a decay factor, lambda (λ) (Danielsson, 2011).

However, some risk measures may lose efficacy through smoothing. For example, the Basel Accords do not allow the use of EWMA in calculating VaR (Danielsson, 2011) since application of an EWMA to return observations *n* periods in the past will reduce them by a factor of $(1-\lambda)^n$. Despite this, EWMAs are commonly applied to variance, which is in turn used as an input into a parametric VaR estimate.

A potential disadvantage of using EWMA as a risk forecast is that λ is constant and identical for all assets, while this assumption is not borne out in practice. Although other GARCH models are able to relax these assumptions, the difference between the EWMA forecasts and other GARCH forecasts can in practice be very small (Danielsson, 2011). Suganuma (2000) tested the performance of a variety of simple moving average models, EWMA models with different assumptions for λ, and a GARCH model against a benchmark EWMA volatility prediction model, and found that none of the models investigated consistently outperformed the benchmark.

The “true” volatility of the underlying data series is required to calculate λ. This may be calculated for a specific sample period as the sum of the squared returns. Lambda may then be calculated by minimising a function of the difference between the true volatility in the series and the EWMA of the series. This minimisation can occur either in sample or out of sample and may use a variety of different functions of this difference to estimate λ, including Root Mean Square Error (RMSE), Mean Absolute Error (MAE), heteroskedasticity-adjusted RMSE (HRMSE) and the heteroskedasticity-adjusted MAE (HMAE). Bollen (2015) estimated λ over rolling out-of-sample 36-month windows on the S&P 500 index between 1957 and 2013, and his results lend credibility to the contention that the use of a constant and identical assumption for λ may be inappropriate.

3.5.4 Calculating the EWMA

One of the criticisms of MPT is that it assumes that probability distributions are constant through time. The use of an EWMA, and in particular, an EWMA with time-dependent rate of decay optimised using a rolling window of data preceding the date of portfolio optimisation, seeks to address this criticism.

The EWMA forecast for variance at time *t* may be calculated recursively as follows (Danielsson, 2011):

$$\sigma_t^2 = (1-\lambda)r_{i,t-1}^2 + \lambda\sigma_{t-1}^2 \tag{7}$$

Where:

- $r_{i,t}$ is the return for portfolio *i* in the period *t*
- λ is the decay factor

The study assumes that the daily mean return is zero for the purposes of calculating the EWMA of each risk measure, since incorporating a non-zero mean into each deviation will result in different weightings of the same observations. This is because the first non-zero-mean deviation is a factor of λ, the second, a factor of $\lambda(1-\lambda)$, the third of $\lambda(1-\lambda)^2$ and so on. Consequently, if an unweighted mean is used to calculate each deviation, the first deviation will incorporate a term equal to $\frac{1}{n}\sum_{t=1}^{n} r_{i,t}$. It follows that every observation of the return incorporated into the mean term will then be weighted by a factor of λ, rather than by the exponentially-reducing factors expected as a result of applying an EWMA. Conversely, the $r_i^2,_t$ terms in the EWMA will be weighted by the expected exponentially-reducing factors. As a result, using a non-zero mean in an EWMA of deviations from the mean will result in the same observations being weighted by different factors (that is, different multiples of λ). This discrepancy holds for each subsequent iteration of the EWMA.

Further, if an EWMA-weighted mean is used to calculate each deviation, the mean term incorporated into the deviation will already be a function of multiples of λ. When the deviation is incorporated into the EWMA by weighting it by multiples of λ, the mean term portion of the deviation will then be weighted by λ again, resulting in the double-weighting of some observations (the observations used to calculate the mean will be weighted first in calculating the mean and again when calculating the risk measure) and single-weighting of others (the observations of the returns which are differenced from the means will be weighted once).

A zero mean is a frequent assumption for the calculation of risk measures in daily time-series data (Alexander, 2008).

Consequently, the exponentially-weighted moving average of each risk measure *r* at time step *t*, $\bar{r}_t$, is calculated recursively as follows:

$$\bar{r}_t = \lambda x_t + (1 - \lambda)\bar{x}_{t-1} \quad (8)$$

Where:

- $\bar{r}_t$ is the exponentially-weighted moving average of the risk measure at time *t*
- x_t is the value/sample observation of the risk measure at time *t*. Following from the assumption of zero mean return, and where the observed return is denoted as *r*, for:
    - Variance, standard deviation, downside semivariance and downside semi standard deviation, $x_t = r^2$
    - MAD and SAD, $x_t = |r|$

  - VaR and CVaR, $x_t = r^2$, where the EWMA will forecast the variance for a normal distribution of returns, which will in turn be used to determine parametric values for VaR and CVaR rather than empirical estimates (Bollen, 2015)
- λ is a smoothing factor with values between 0 and 1

The study determines appropriate rolling values for λ at each calculation period by applying the methodology put forward by Bollen (2014) over the sample period, using numerical methods to minimise the heteroskedasticity-adjusted mean absolute error (HMAE). Each observation (that is, each term) in the HMAE requires two sample periods to calculate. The first period is used to calculate an estimate for the risk measure, *r*, which is proposed to be one month, or 22 periods. The second period is used to calculate the forecast for the risk measure (that is, the EWMA, $\bar{r}_t$), and is proposed to be a minimum of two months, or 44 periods, but may be longer, going back to the beginning of the full set of sample data. For the calculation of λ at time *t*, the observation of the HMAE at time *t* is therefore:

$$\mathrm{HMAE\_Obs_t} = |1 - \frac{\mathrm{r_{t-22:t-1}}}{\mathrm{\bar{r}_{1:t-23}}}| \tag{9}$$

Where:

- $r_{\mathrm{t}-22:t-1}$is the value of the risk measure calculated on the sample returns between periods *t*-22 and -1 (inclusive); and
- $\bar{r}_{1:t-23}$ is the value of the EWMA of the risk measure calculated on sample returns between the first period of the full sample and period -23 (inclusive)

The full HMAE is calculated by successively moving backwards one period in time (from time *t* to *t*-1) and repeating, to calculate $|1 - \frac{\mathrm{r_{t-23:t-2}}}{\mathrm{\bar{r}_{1:t-24}}}|$, then $|1 - \frac{\mathrm{r_{t-24:t-3}}}{\mathrm{\bar{r}_{1:t-25}}}|$ and so on, and averaging the iterations. Using subscript *p* to denote the number of times this is repeated, and the subscript *j* to denote the iteration over the periods between *t*-*p* and *t*:

$$\mathrm{HMAE}_t = \frac{1}{\mathrm{p}} \sum_{\mathrm{j= t-p}}^{\mathrm{t}} \mathrm{HMAE\_Obs_t} \tag{10}$$

The $HMAE_t$ is then calculated for possible values of λ between 0.05 and 0.995 with a 0.005 increment. The final estimate for λ is that value of λ between 0 and 1 (exclusive) which results in the lowest $HMAE_t$.

Evaluating a new value for λ in each reweighting interval addresses the criticism of EWMA that the λ parameter is constant despite changes in the underlying risk, and adjusts for possible differences in optimal λ between different risk measures.

The periods used (that is, hyperparameters) are:

- *p* = 3 months
- The period over which the estimate of the risk measure is calculated is 1 month (22 periods)
- The period over which the EWMA is calculated is a minimum of 2 months (44 periods), up to the full length of the sample data

A longer period for calculating the risk measure estimate results in a relatively staler estimate for the true value of the sample observation at time t, but will smooth progression thereof, and consequently also that of λ. It seems feasible that the former may perform better in stable market periods, while the latter may perform better in more volatile markets.

Similarly, a longer total estimation period, *p*, results in a smoother progression of λ over time, while a shorter period for p results in a faster and larger deviation in λ. Again, a longer total estimation period for calculating λ results in a relatively staler estimate for λ at time *t*, but will smooth progression thereof.

Since the weekly and monthly portfolios use the daily data to calculate risk measures, it is not necessary to calculate different values for λ for each interval.

3.5.5 Determining the optimal portfolio at a point in time

Having calculated the expected return and risk measures at each time step, Markowitz's (1952) illustrative portfolio may now be constructed as follows:

- Let random variable R_i be the return on the i^{th} security and random variable R the portfolio return
- Let X_i be the percentage of the investor's portfolio invested into the i^{th} security, with $\Sigma X_i = 1$

Then:

Return $R = \Sigma R_i X_i$ (11)

Expected Return = $\Sigma_{i=1}^{N} X_i \mu_i$ (12)

Variance = $\Sigma_{i=1}^{N} \Sigma_{j=1}^{N} \sigma_{ij} X_i X_j$ (Markowitz, 1952) (13)

Per Markowitz's (1952) formulation, any portfolio for which, given the return of the portfolio, there is no other portfolio with a smaller variance, will lie on the efficient frontier. For the purposes of calculating each risk-measure-optimised portfolio, variance is replaced with each risk measure in turn.

Per Sharpe (1964), the portfolio on the efficient frontier which touches a tangential line connecting a risk-free asset and efficient frontier (the CML) is the optimal portfolio in terms of allocation of risk and return. Finding the portfolio which maximises the gradient of the CML finds the optimal portfolio.

The study uses numerical methods to find the portfolio which maximises the gradient of the CML. Using numerical methods ensures that the same method can be used to find the optimal portfolio for any risk measure without introducing unintentional bias. However, the slope of the CML as a function of the weights of the portfolio constituents is discontinuous for some risk measures, which results in the application of numerical methods being intractable. Conversely, the use of numerical methods also allows calculation of an optimal portfolio where no tractable closed form solution exists and allows the imposition of constraints on the portfolio's stock weightings.

The use of non-parametric numerical methods to calculate each risk measure investigated, with the exception of EWMA-weighted value at risk and EWMA-weighted Conditional value at risk, also addresses the criticism of MPT that it assumes returns are normally distributed.

This methodology is applied for each risk measure, each EWMA forecast and each reweighting interval (that is, rolling from period to period) to find the relevant optimal portfolio as follows:

1. Select a portfolio which contains all assets in the investment universe with weightings X_i, but subject to the constraint that $\Sigma X_i = 1$. In the first reweighting interval, the weights X_i may be arbitrary or equal. However, at each subsequent reweighting interval the weights will rather be the optimal portfolio found at the previous iteration (that is, the previous period in which the optimal portfolio was calculated), thereby simulating the performance of the continuously-recalculated portfolio. This will have the advantage of reducing the computing requirement, as well as reducing the possibility that the portfolio will oscillate between differently-weighted portfolios with similar mean-risk measure characteristics, thereby avoiding unnecessary trading fees.
2. Calculate the expected return and the risk measure (or EWMA forecast) as detailed in Sections 0 and 3.5.4 using the preceding three months of returns as the inputs.
3. Optimise the portfolio using numerical methods by:
    a. Estimating a partial derivative, $\frac{dS}{dX_i}$, of the slope, S = expected return / risk measure, with respect to each X_i by incrementing each X_i in turn by a small amount and then recalculating S.
    b. Once every partial derivative has been calculated, incrementing or decrementing X_i by a small amount in the direction which increases S.
    c. Repeating with successively smaller increments and decrements until the change in S falls below an acceptable threshold.
4. Resize each X_i proportionately to ensure that $\Sigma X_i = 1$

The study imposes constraints on each X_i such that no short positions are allowed.

3.5.6 Recalculating successive optimal portfolios

Having identified the optimal portfolio at a point in time, the return of the portfolio between the given timestep *t* and the next timestep *t*+1 may be calculated using the weights at timestep *t* and the asset prices or total returns at timestep *t* and *t*+1.

The process will be repeated at timestep *t*+1 and each subsequent timestep.

The study incorporates the effect of fees on the aggregate return as follows:

- Total expense ratio (TER) for each ETF of 0.1% per annum, as per the actual TERs charged by the ETFs based on the indices used;
- Trading fees of USD 0.003 per share traded, based on the trading fees of the NASDAQ and NYSE. For the purposes of the study, it will be assumed that the starting value of each share will be the starting index value multiplied by 1.00 US Dollar; and
- Zero commission.[1]

Given the low fee levels there is little difference in trading fees between the different strategies, with fees across all strategies resulting in a circa 10bp reduction in annual return, except for the daily EWMA-optimised strategy which was closer to 12bp. The trading fees are typically low, given their very small amount per share traded, with the result that the total fees are typically close to the TER.

While the applicable fees are subtracted from each portfolio at each time step in which the portfolio is recalculated, the impact of fees shown in the results tables in Section 4.3 (Strategy Performance), is expressed as an annualised reduction in yield. Consequently, since the reduction in yield is a compounded figure, but the fees added to the portfolios are not compounded, the figures in the tables can be less than the 10bp.

[1] Based on the fees schedules of Merrill Edge, Fidelity Investments, TD Ameritrade and Interactive Brokers.

The incorporation of fees into the calculation of total returns addresses the criticism of MPT that it does not include fees, although fees are not included as a constraint into the optimisation problem itself.

As indicated, the above procedure has been repeated for each risk measure, and for the EWMA of each risk measure.

The output of this process is a set of returns for each of the portfolios or strategies. The process is summarised in Figure 3.5.12.

Figure 3.3

Optimisation process flow diagram

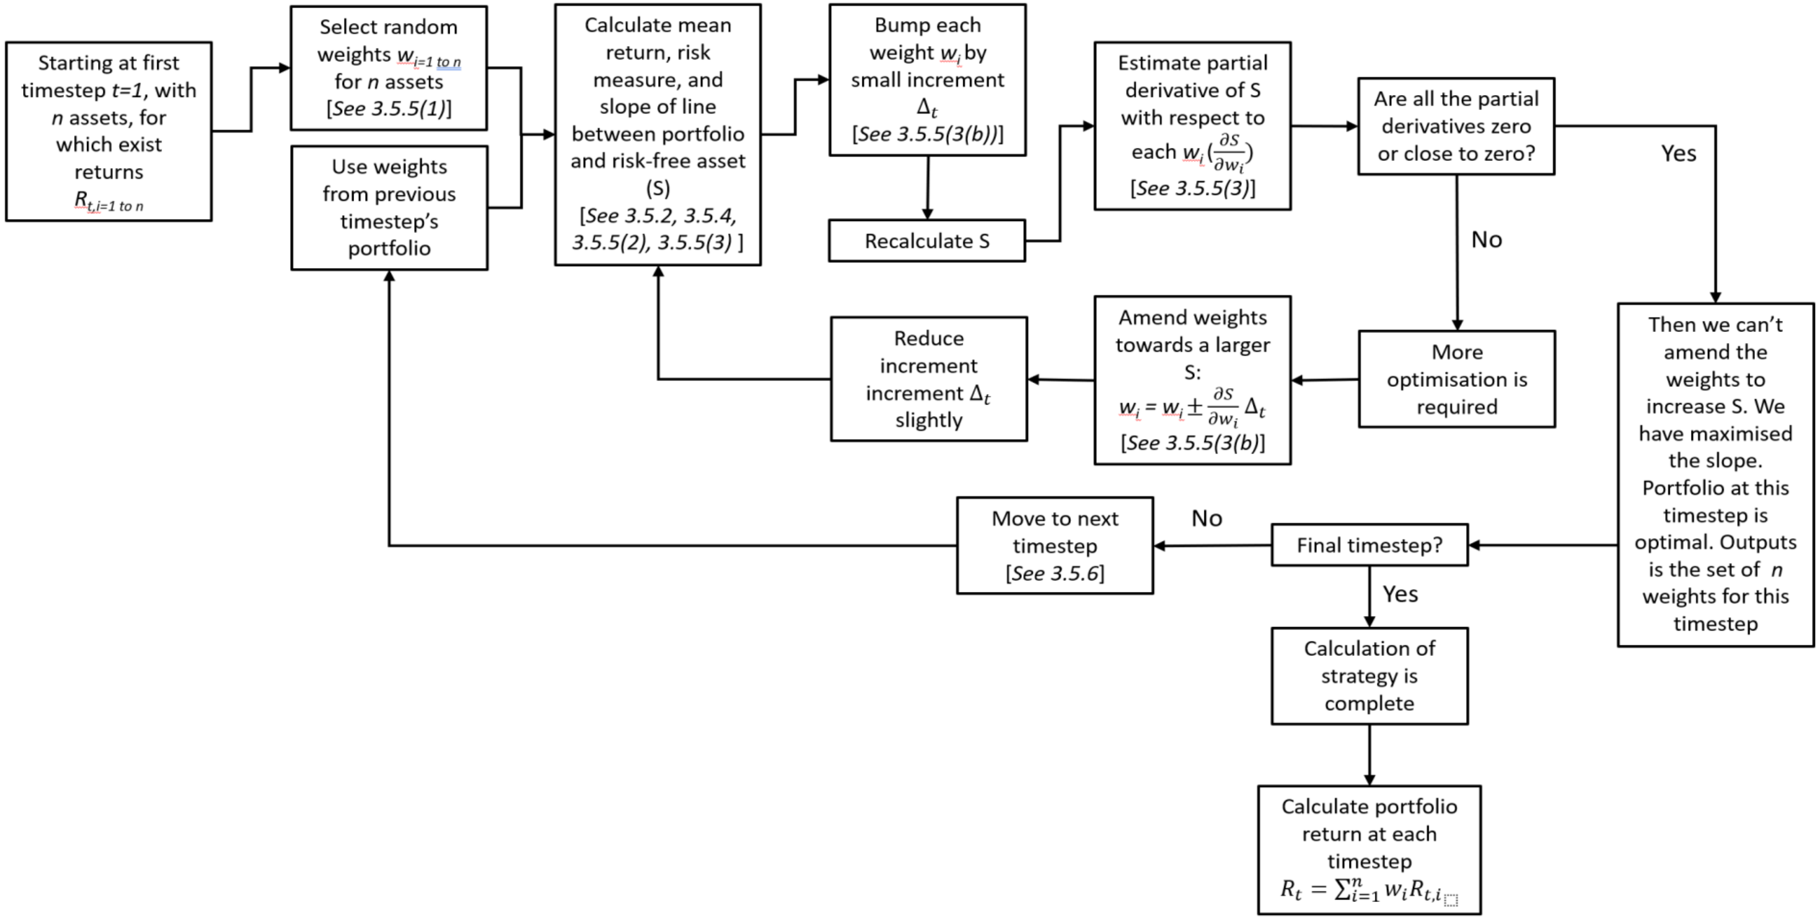

Note. Process flow followed to create each investment strategy.

3.6 ROBUSTNESS CHECKS

Possible sources of bias or error include the following:

1) Asset sample size: The availability of tradeable ETFs based on the indices used in the study reduces the study sample size to only ten assets plus the global index. While other studies have used an asset universe as small as only two assets (Jacobsen 2005), the results will be corroborated or refuted by conducting a similar optimisation exercise on a subset of the portfolios investigated using the next-level-down 24 MSCI Industry Group indices and the 69 MSCI Industry indices.
2) In-sample testing: A portfolio construction technique which uses future information to construct a portfolio is likely to perform extremely well, but is also impossible to implement in reality. The study will ensure that only information prior to each timestep will be used to construct the asset weights to be applied at that timestep. Ensuring this is facilitated by checking that the returns inputs into the functions described in section 3.6 (Data analysis) do not include any information from after the timestep being optimised for. This also impacts on the choice of training period: Although there are training periods which result in better strategy outcomes than the three months selected, the training period was not altered to accommodate this.
3) Overfitting of hyperparameters: If enough hyperparameters are assessed, it is possible that one or more will result in a particularly optimal output portfolio by chance, the results of which would not be replicable with different assets or a different period of sample data. This risk may be minimised by reporting on the process of identifying the study hyperparameters, and by using the portfolios optimised using the MSCI Industry Group and Industry indices as an out-of-sample test case.

3.7 RESEARCH ETHICS

The data used in the study is collected and analysed as described, without manipulation or falsification of data or results. The findings are reported in a transparent manner and avoid the portrayal of any misleading outcomes or misrepresentations.

All information from external sources is appropriately paraphrased and cited.

The necessary approvals have been obtained from the Research Ethics Committee of the Faculty of Economic and Management Sciences.

3.8 SUMMARY

The study uses a positivist research paradigm based on an empirical enquiry strategy. It uses quantitative and statistical methods applied to time series numeric secondary data sourced from S&P corroborated with data from MSCI. These methods facilitate the construction of several investment strategies based on portfolios optimised using different risk measures; as well as the comparison of the risk-adjusted performance of the investment strategies.

CHAPTER 4 – RESULTS AND DISCUSSION

4.1 INTRODUCTION

Identifying an optimal investment strategy is built up from constituent steps, each of which will be discussed before the resulting strategies are compared and possible optimal strategies are examined in more detail. The steps followed in the study start with establishing that the optimisation algorithm employed is fit for purpose – that is, that the numerical methods optimisation does indeed result in an efficient portfolio. This is followed by the process of establishing appropriate hyperparameters, in line with one of the stated purposes of the study. Comparisons of the results of strategies implemented with different hyperparameters must, however, be weighed against the risk of introducing forward-looking bias. Hyperparameters examined in the study include the training period, the recalculation interval, the size of the asset universe, the choice of whether to optimise over raw risk measures or forecast risk measures, and the choice of lambda as an input parameter to the EWMA.

The study calculated the returns, each risk measure used in the study, drawdown percentage and risk-adjusted returns, for the 140 different strategies investigated. While returns and strategy risk measures for all 140 strategies are included in Appendix B, only the summarised results and the best-performing strategy are discussed in detail in this chapter. In particular, the returns and Sharpe ratio of all strategies are compared to identify the best and worst performing strategies.

4.2 POINT-IN-TIME PORTFOLIO CONSTRUCTION

If the numerical methods algorithm used to calculate each successive optimal portfolio is ineffective in doing so, the results obtained from each strategy will be misleading. To guard against such a possibility, the algorithm was tested for point-in-time portfolios optimised using each risk measure and the EWMA of each risk measure[2]. The conclusion of the

[2] Since the results of the validation are sufficient to confirm the validity of the results coming out of the optimisation methods, it is not necessary to describe the full validation in the main body of the study. Readers interested in the process of validation can find the process and analysis described in Appendix B.

validation testing is that, firstly, certain risk measures are easier to optimise over than others, in particular, symmetrical risk measures (variance, standard deviation, mean absolute deviation) are typically more tractable for optimisation than single-sided risk measures (semi variance, semideviation, semi-absolute deviation, value at risk). The reasons for this are explained in Appendix B. Irrespective of this, the optimised portfolios are sufficiently well optimised that the resulting strategies are suitable for drawing inference about portfolios optimised using these risk measures. Secondly, it is clear that application of an EWMA to single-sided risk measures renders them tractable to optimisation using numerical methods.

4.3 STRATEGY PERFORMANCE

Having established that the optimisation algorithm does indeed result in efficient portfolios, the portfolios must be calculated in each time step to create an investment strategy. However, prior to doing so, an intermediate step is required, being to optimise each of the parameters which are required as inputs into the investment strategy; that is, to establish optimal hyperparameters. The remainder of this section touches briefly on the process of identifying optimal hyperparameters, before examining a possible solution to the criticism of MPT that past performance does not persist into the future. Thereafter, the performance of the different investment strategies examined in the study is compared taking into account the chosen hyperparameters, and finally more comprehensive results for the best-performing strategies are examined in more detail.

4.3.1 Hyperparameters

The study required the optimisation of different hyperparameters before using them in the final investment strategies, being recalculation period, length of training period and GICS tier.

In order to avoid forward-looking bias, the hyperparameters were optimised using past data, or alternatively, where this was not feasible, a parameter was chosen and the entire sample was used to corroborate, but not fit, the choice of parameter.

It is preferable to select hyperparameters for which adjacent parameter values have a positive but similar effect on the outcome of a strategy in preference to those which result in better strategy performance, but poor results for adjacent parameter values. Doing so means that a small shift in the investment environment with respect to the effect of the parameter is less likely to result in significant changes to the strategy performance.

The choice of hyperparameters has a profound effect on the portfolios fitted from the data. Poorly-optimised or overfitted hyperparameters are likely to result in models with poor predictive capabilities, invalidating the results of the study. Conversely, hyperparameters with forward-looking bias might provide spuriously positive results which could never be obtained in reality - which has no access to future data - thereby equally invalidating the results of the study. Consequently, the process of calculating and choosing the hyperparameters is a key input into the results of the study. The outcome of the hyperparameterisation process is that the results of the study are calculated as follows: weekly and monthly portfolio recalculation periods provide the best returns; a training period of three months works well for the study; and using sector and industry group indices as inputs results in portfolios with better performance than using less-aggregated industry indices.

The process of optimising hyperparameters is documented in Appendix C, while the financial results of portfolio optimisation using each different risk measure are documented in the remainder of this chapter.

4.3.2 Comparison of all strategies

Once the hyperparameters were established, the 140 risk-return optimised strategies were run incorporating each risk measure, GICS tier, length of recalculation period, and EWMA/raw risk measure combination. The annualised return for each strategy is shown in Table 4.1, while the Sharpe ratio is shown in Table 4.2 The gradation of performance by colour (green is better, red is worse) is displayed horizontally between risk measures in both tables. All results are out-of-sample.

Table 4.1

Annualised returns for all investment strategies

			Risk measure optimised against									
Tier	Calculation Period	Raw/EWMA	V	SD	DSV	DSSD	MAD	SAD	VaR0.1	VaR0.05	CVaR0.1	CVaR0.05
Sector	Daily	Raw	9.61%	10.98%	10.88%	10.91%	10.94%	12.60%	9.31%	8.84%	9.92%	10.52%
		EWMA	11.26%	8.71%	7.10%	7.72%	8.95%	8.72%	8.82%	9.01%	9.12%	9.64%
	Weekly	Raw	10.66%	11.72%	9.73%	13.94%	12.20%	13.50%	9.57%	8.06%	11.61%	11.87%
		EWMA	10.71%	12.34%	12.70%	11.45%	16.18%	15.74%	12.63%	12.60%	5.11%	7.49%
	Monthly	Raw	11.34%	10.99%	10.50%	13.07%	12.61%	15.18%	9.27%	11.14%	10.90%	11.81%
		EWMA	7.61%	11.43%	7.70%	10.64%	11.50%	11.62%	10.90%	11.26%	8.89%	11.94%
Industry Group	Weekly	Raw	9.23%	14.48%	9.98%	13.71%	12.64%	11.60%	10.95%	14.57%	15.42%	16.27%
		EWMA	10.53%	18.88%	14.93%	19.18%	20.01%	19.48%	18.69%	18.71%	14.90%	15.76%
	Monthly	Raw	11.78%	14.51%	12.41%	13.13%	13.75%	14.17%	12.63%	12.16%	15.51%	15.31%
		EWMA	6.56%	10.53%	11.10%	13.03%	11.04%	12.29%	10.16%	10.17%	11.94%	10.29%
Industry	Weekly	Raw	10.66%	11.72%	9.73%	13.94%	12.20%	13.50%	9.57%	8.06%	11.61%	11.87%
		EWMA	10.71%	12.34%	12.70%	11.45%	16.18%	15.74%	12.63%	12.60%	5.11%	7.49%
	Monthly	Raw	11.34%	10.99%	10.50%	13.07%	12.61%	15.18%	9.27%	11.14%	10.90%	11.81%
		EWMA	7.61%	11.43%	7.70%	10.64%	11.50%	11.62%	10.90%	11.26%	8.89%	11.94%
Raw			10.66%	12.20%	10.53%	13.11%	12.42%	13.67%	10.08%	10.57%	12.27%	12.78%
EWMA			9.29%	12.24%	10.56%	12.02%	13.62%	13.60%	12.10%	12.23%	9.14%	10.65%
Sector			10.20%	11.03%	9.77%	11.29%	12.06%	12.89%	10.08%	10.15%	9.26%	10.54%
Industry Group			9.52%	14.60%	12.11%	14.76%	14.36%	14.38%	13.11%	13.90%	14.44%	14.41%
Industry			10.08%	11.62%	10.16%	12.28%	13.12%	14.01%	10.59%	10.76%	9.13%	10.78%
Overall			9.97%	12.22%	10.55%	12.56%	13.02%	13.64%	11.09%	11.40%	10.70%	11.71%

Note. Annualised returns for each investment strategy investigated. V = variance, SD = standard deviation, DSV = downside semivariance, DSSD = downside semideviation, MAD = mean absolute deviation, SAD = semi-absolute deviation, VaR0.1 = 10% value at risk, VaR0.05 = 5% value at risk, CVaR0.1 = 10% conditional value at risk, CVaR0.05 = 5% conditional value at risk.

Source: Author's calculations

Table 4.2

Sharpe ratios for all investment strategies

			Risk measure optimised against									
Tier	Calculation Period	Raw/EWMA	V	SD	DSV	DSSD	MAD	SAD	VaR0.1	VaR0.05	CVaR0.1	CVaR0.05
Sector	Daily	Raw	0.65	0.70	0.68	0.69	0.69	0.82	0.58	0.54	0.62	0.65
		EWMA	0.77	0.55	0.45	0.48	0.57	0.56	0.55	0.57	0.58	0.62
	Weekly	Raw	0.71	0.73	0.62	0.95	0.76	0.85	0.60	0.49	0.73	0.73
		EWMA	0.71	0.77	0.81	0.71	1.03	1.00	0.79	0.79	0.30	0.44
	Monthly	Raw	0.91	0.79	0.82	1.07	0.92	1.11	0.70	0.82	0.78	0.85
		EWMA	0.59	0.86	0.56	0.76	0.87	0.89	0.82	0.85	0.71	0.84
Industry Group	Weekly	Raw	0.53	0.81	0.57	0.75	0.74	0.70	0.62	0.83	0.87	0.91
		EWMA	0.66	1.08	0.89	1.16	1.15	1.12	1.07	1.07	0.78	0.86
	Monthly	Raw	0.79	0.88	0.87	0.82	0.84	0.88	0.77	0.76	0.96	0.96
		EWMA	0.43	0.64	0.68	0.76	0.70	0.76	0.62	0.62	0.59	0.52
Industry	Weekly	Raw	0.71	0.73	0.62	0.95	0.76	0.85	0.60	0.49	0.73	0.73
		EWMA	0.71	0.77	0.81	0.71	1.03	1.00	0.79	0.79	0.30	0.44
	Monthly	Raw	0.91	0.79	0.82	1.07	0.92	1.11	0.70	0.82	0.78	0.85
		EWMA	0.59	0.86	0.56	0.76	0.87	0.89	0.82	0.85	0.71	0.84
Raw			0.75	0.78	0.71	0.90	0.80	0.90	0.65	0.68	0.78	0.81
EWMA			0.64	0.79	0.68	0.76	0.89	0.89	0.78	0.79	0.57	0.65
Sector			0.72	0.73	0.66	0.78	0.81	0.87	0.67	0.68	0.62	0.69
Industry Group			0.60	0.85	0.75	0.87	0.86	0.87	0.77	0.82	0.80	0.81
Industry			0.73	0.79	0.70	0.87	0.90	0.96	0.73	0.74	0.63	0.72
Overall			0.69	0.78	0.70	0.83	0.85	0.90	0.72	0.74	0.67	0.73

Note. Sharpe ratios for each investment strategy investigated. V = variance, SD = standard deviation, DSV = downside semivariance, DSSD = downside semideviation, MAD = mean absolute deviation, SAD = semi-absolute deviation, VaR0.1 = 10% value at risk, VaR0.05 = 5% value at risk, CVaR0.1 = 10% conditional value at risk, CVaR0.05 = 5% conditional value at risk.

Source: Author’s calculations

Across all GICS tiers, all recalculation periods and both raw risk measure- and EWMA-optimised strategies, semi-absolute deviation-optimised strategies achieved the highest overall annualised returns (Table 4.1) and Sharpe ratios (Table 4.2), followed by mean-absolute deviation-optimised strategies and downside semideviation-optimised strategies. The performance of semi-absolute deviation-optimised strategies was also *consistently* strong, with, barring two outliers for the weekly raw risk measure-optimised strategy and the daily EWMA-optimised strategy, a relatively small difference between the best and worst performance. In combination with the principle relating to hyperparameters of the benefit of adjacent parameter values having a positive but similar effect on the outcome of a strategy, this shows semi-absolute deviation to be the best performing risk measure in the study.

Significantly, given variance is the risk measure originally put forward in MPT, variance-optimised strategies performed most poorly in terms of absolute returns and second most poorly in terms of Sharpe ratio.

The performance of semi-absolute deviation-optimised strategies compared to variance-optimised strategies is consistent with other empirical studies in the literature. Stanković et al. (2020) concluded that semi-absolute deviation-optimised portfolios substantially outperformed variance-optimised portfolios on the Belgrade Stock Exchange, albeit at the cost of higher risk. Similarly, Jacobsen (2005) showed that the performance of a semi-absolute deviation-optimised portfolio outperformed a variance-optimised portfolio consisting of the Standard & Poor's 500 and the Dow Jones Two-Year Corporate Bond Indices for the sample period between 1996 and 2005.

One of the factors impacting the performance of semi-absolute deviation-optimised strategies is asset allocation concentration. Stanković et al. (2020) observed that semi-absolute deviation-optimised portfolios resulted in much smaller numbers of shares being held compared to variance-optimised portfolios. The same is true in this study, with average asset allocation weights in semi-absolute deviation-optimised strategies across all periods being more concentrated than other strategies. Using sector-based raw monthly strategies as an example, semi-absolute deviation-optimised strategies had the highest and second most concentrated asset allocations of any strategy, allocating 37.5% of assets to IT and 26.4% to Health Care on average, compared to the average most concentrated allocation across all the other strategies of 17.7%. This contributed to the success of the semi-absolute deviation-optimised strategies, since the average monthly return of the IT index was the highest of all the sector indices, at 1.7% per period, and the Health Care index achieved the third-highest monthly return at 1.1% per period.

The improvement in annualised returns over the variance-optimised strategies evident in strategies optimised using downside semivariance, downside semideviation and semi-absolute deviation, as evident in Table 4.1, is at least partially attributable to the use of these measures addressing the criticism made my Markowitz (1952) that MPT minimises upside risk as well as downside risk. Any constituent asset which exhibited large upward shifts in return relative to average returns will have been penalised in variance-optimised portfolios, but not in strategies optimised using downside semideviation or semi-absolute deviation.

Comparing individual variance-optimised strategies to downside semivariance-optimised strategies, and analogously, standard deviation-optimised strategies to downside semideviation-optimised strategies, the annual return results in Table 4.1 are mixed as to whether one materially outperforms the other or not, just as experienced by Jacobsen (2005), Stanković et al., (2020) and Hunjra et al. (2020) in each of their studies comparing variance-optimised strategies to downside semivariance-optimised strategies. On average, however, downside semivariance-optimised strategies outperformed variance-optimised strategies slightly, downside semideviation-optimised strategies outperformed standard deviation-optimised strategies, and semi-absolute deviation strategies each outperformed mean absolute deviation strategies by circa 1% p.a.

Consequently, it appears that this particular criticism of MPT may be addressed through the use of single-sided risk measures in preference to symmetrical risk measures.

The high annualised returns in Table 4.1 of mean absolute deviation- and semi-absolute deviation-optimised strategies relative to other strategies, may also be attributable to these risk measures conferring a penalty to inclusion of an asset into the portfolio proportional to the difference between each period's return and the mean return, rather than the square thereof. As contended by Byrne & Lee (2004), this makes these risk measures less sensitive to outliers than variance, and similarly less sensitive to outliers than standard deviation, downside semivariance and downside semideviation. A higher sensitivity to outliers would penalise the inclusion into the optimised portfolio of assets with bigger deviations in return away from the mean return. This would be particularly problematic for symmetrical risk measures such as variance, reducing the weighting in assets with large upside return deviations. However, assuming at least some symmetry of return distributions, penalising the portfolio weightings of assets with large downside return deviations would also impact negatively on the portfolio return, since the large upside return deviations associated with those same assets would also be excluded. This effect would be exacerbated by an environment in which the return on assets was generally positive, as is the case in the study period, since, on average, excluding large deviations in return would be likely to exclude more large upside deviations than large downside deviations.

Comparing variance-optimised strategies to standard deviation-optimised strategies, it is clear from Table 4.1 that the latter strategy results in higher annualised returns. This is a result of variance being the square of standard deviation. Where variance is the denominator of the function being optimised, being return/risk, the optimisation algorithm attributes more importance to risk than where standard deviation is the denominator, such that assets with lower risk and lower returns will be more heavily weighted in a variance-optimised strategy, while assets with proportionately higher risk and higher returns will be more heavily weighted in a standard deviation-optimised strategy. The same relationship is evident between downside semivariance-optimised strategies and downside semideviation-optimised strategies in Table 4.1, and for the same reason.

The relatively low annualised returns in Table 4.1 of value at risk- and conditional value at risk-optimised strategies is a result of two factors. Firstly, while no similar discussion was evident in journal articles on this topic, these two risk measures penalise the weighting of assets with downside outliers, since to minimise value at risk or conditional value at risk, asset combinations resulting in downside returns lower than the n^{th} percentile must be excluded from the portfolio. Again, to the extent that the assets in the portfolio display at least some symmetry of return distributions, penalising the portfolio weightings of assets with large downside return deviations would also impact negatively on the portfolio return, since the large upside return deviations associated with those same assets would also be excluded. Secondly, the returns of the indices generally display a reversal after achieving the minimum return amongst the assets. Using the raw monthly returns as an example, the impact of excluding the asset with the minimum return from the next period’s portfolio is a reduction in yield for eight of the ten assets, with an average reduction in yield of 1% per annum. Put more broadly, the value at risk and conditional value at risk-optimised strategies penalise ‘bounce back’ or mean reversion, which the indices’ returns show sufficiently to engender poor returns relative to other strategies.

Boasson et al. (2011) noted that portfolios optimised using VaR were sensitive to the confidence level used. This was also reported by Hunjra *et al.* (2020), who noted that portfolios optimised using a higher confidence interval outperformed those optimised using a lower confidence interval in some instances, and *visa versa* in other instances. This is apparent in the study results, with the strategies differing only in confidence level yielding very different results. Averaging both annualised returns and Sharpe Ratio across all strategies (see Tables 4.1 and 4.2), the strategies with a lower confidence level (5%) yielded better performance on average than those with a higher confidence level (10%). This is attributable to the strategies with a higher confidence level capturing a greater proportion of downside periods followed by a subsequent upwards correction in return for which the strategy has then already exited its position in said asset, as discussed in more detail in the previous paragraph.

While semi-absolute deviation is optimal in terms of annualised returns and Sharpe ratio, it does not minimise the risk of the strategy relative to portfolios optimised using other risk measures. Again, this is consistent with the empirical studies of Stanković et al (2020) and Jacobsen (2005), which concluded that semi-absolute deviation-optimised portfolios result in higher risk than variance-optimised portfolios. Table 4.3 shows the average risk in terms of each risk measure used in the study for each risk measure-optimised portfolio, averaging across weekly and monthly recalculated EWMA- and raw risk measure-optimised strategies.

Table 4.3

Annualised risk measures for each investment strategy

Risk measure used for optimisation	Strategy risk: V	SD	DSV	DSSD	MAD	SAD	VaR0.1	VaR0.05	CVaR0.1	CVaR0.05	Average	Max Drawdown
V	0.08%	2.77%	0.11%	3.28%	1.82%	2.09%	2.38%	3.49%	4.44%	6.08%	2.65%	23.5%
SD	0.09%	3.01%	0.12%	3.47%	1.99%	2.22%	2.47%	3.66%	4.73%	6.51%	2.83%	23.1%
DSV	0.08%	2.90%	0.11%	3.37%	1.90%	2.10%	2.56%	3.69%	4.75%	6.56%	2.80%	23.8%
DSSD	0.08%	2.91%	0.11%	3.35%	1.92%	2.08%	2.20%	3.40%	4.52%	6.33%	2.69%	22.1%
MAD	0.09%	3.00%	0.12%	3.46%	1.98%	2.20%	2.46%	3.57%	4.62%	6.36%	2.79%	22.9%
SAD	0.09%	2.99%	0.12%	3.48%	1.98%	2.21%	2.42%	3.66%	4.56%	6.29%	2.78%	24.6%
VaR0.1	0.09%	2.93%	0.11%	3.38%	1.96%	2.15%	2.5925%	3.67%	4.67%	6.37%	2.79%	23.4%
VaR0.05	0.09%	2.98%	0.12%	3.42%	1.99%	2.18%	2.46%	3.5743%	4.64%	6.44%	2.79%	24.0%
CVaR0.1	0.09%	2.93%	0.11%	3.38%	1.94%	2.16%	2.3997%	3.81%	4.76%	6.55%	2.81%	22.1%
CVaR0.05	0.10%	3.11%	0.12%	3.49%	2.05%	2.21%	2.61%	3.7446%	4.93%	6.83%	2.92%	23.3%

Note. Annualised risk measures for the returns generated by each investment strategy investigated. V = variance, SD = standard deviation, DSV = downside semivariance, DSSD = downside semideviation, MAD = mean absolute deviation, SAD = semi-absolute deviation, VaR0.1 = 10% value at risk, VaR0.05 = 5% value at risk, CVaR0.1 = 10% conditional value at risk, CVaR0.05 = 5% conditional value at risk. Each risk measure is a sample statistic calculated based on the full strategy returns over the 10-year study period. Descriptions of the calculations for each are described in Section 3.5.2.

Source: Author's calculations

The results of optimising over the different risk measures displayed in Table 4.3 are not counterintuitive when considering the penalties imposed on the optimisation by minimising each different risk measure. The underlying reasons for the relative riskiness of portfolios optimised using different risk measures is not discussed in the literature. However, consideration of the mathematical formulation of each risk measure shows which assets result in an increase in the final value of the risk measure and which do not, with particular reference to whether minimising a risk measure results in stricter or more lenient exclusion of assets in the portfolio. If the exclusion of assets from the portfolio as a result of minimising risk is viewed as a penalty, then the most extreme penalty applied is likely to be that of variance, which penalises all large movements away from the mean return, with larger movements away from the mean penalised more than smaller movements due to variance being the sum of the square of the differences. Downside variance is similar, but does not penalise the upside deviations which variance penalises. Conversely, value at risk and conditional value at risk have penalties which apply only to the tail, resulting in the fewest penalties being applied during the optimisation. Somewhere in between the two, mean absolute deviation applies penalties to both upside and downside deviations from the mean return, but has no increased weighting applied to returns far away from the mean, and semi-absolute deviation is similar, but only penalises downside deviations from the mean. The application of these heavier or lighter penalties is evident in Table 4**Error! Reference source not found.**.3, where the risk of portfolios optimised using each risk measure is generally inversely proportional to the extremity of each risk measure's penalty.

Stanković et al. (2020) suggested that optimising a portfolio with respect to conditional value at risk usually leads to well-optimised portfolios in terms of value at risk. While this seems to be true when comparing the value at risk of strategies optimised using conditional value at risk to that of strategies optimised using value at risk at a 10% confidence level (these are equal to each other; see Table 4.3), this breaks down at a 5% confidence level, where the value at risk-optimised strategy shows lower value at risk than the conditional value at risk-optimised strategy. One possibility is that this is attributable to methodology differences between the calculation of the two risk measures. In the study, raw value at risk and raw conditional value at risk are calculated directly from historical data, whereas for their EWMAs, a distribution is fitted to losses and the risk measure estimated from the parameters thereof. However, this does not seem to be the case, non-parametric conditional value at risk-optimised strategies do not *all* have lower value at risk than non-parametric value at risk-optimised strategies, and similarly for the parametric cases. Despite this, the average value at risk for conditional value at risk-optimised strategies and value at risk-optimised strategies is only very slightly different.

A number of conclusions can be drawn from the relative risk of the different strategies. Firstly, while the classical MPT variance-optimised strategy did not perform well in terms of returns, it performs extremely well in terms of minimising risk across all risk measures used to measure the risk of the strategy. The opposite is true of semi-absolute deviation, which has intermediate levels of risk as indicated by the risk measures, and the highest maximum draw down of any strategy. This implies that its high Sharpe ratio is more a function of high returns than of low risk. A second conclusion evident from Table 4.3 is that a strategy's impact on risk is generally consistent across all risk measures, not just the risk measure being optimised over.

The tight range of maximum drawdown between different strategies is also noteworthy. Maximum drawdown was between 21.1% and 23.6% for all strategies and occurred during the CoVID19 market crash. This corroborates the contentions of Lukomnik and Hawley (2021) when they argued that modern portfolio theory fails to capture ambiguous events, black swan events and systemic risk, citing COVID-19 as an example, which was not captured by preceding returns data and therefore was not an investment input into MPT models. Intuitively, the choice of risk measure, all of which are backwards-looking, should not provide any significant mitigation against such events. This is also corroborated by the study. While capturing black swan events may not be possible, allowance for systemic events may be possible through the introduction of short selling and market-neutral investment strategies. However, since short selling is outside of the scope of this study, so is systemic risk.

While the merits of replacing variance with an alternative risk measure to optimise a portfolio seem evident from the results in Table 4.1 and Table 4.2, Table 4.4 displays the percentage of instances in which the risk-adjusted returns of strategies optimised using alternative risk measures outperformed the risk-adjusted returns of strategies optimised using variance.

Table 4.4

Percentage of alternative risk measure-optimised strategies which outperformed variance-optimised strategies (by Sharpe ratio)

Data	Calculation Period	Raw/EWMA	Outperformed Variance
Sector	Daily	Raw	60%
		EWMA	0%
	Weekly	Raw	60%
		EWMA	70%
	Monthly	Raw	37%
		EWMA	91%
Industry Group	Weekly	Raw	100%
		EWMA	100%
	Monthly	Raw	80%
		EWMA	97%
Industry	Weekly	Raw	60%
		EWMA	70%
	Monthly	Raw	37%
		EWMA	91%
Average			68%

Note. Outperformed Variance shows the percentage of the nine strategies in each category optimised using a non-variance risk measure which outperformed the variance-optimised strategy in that category.

Source: Author’s calculations

Table 4.4 shows that 68% of strategies optimised using alternative risk measures outperformed variance-optimised strategies, where better performance is categorised as having a higher Sharpe Ratio within each grouping of GICS tier, recalculation period and EWMA/Raw risk measure. Consequently, using an alternative risk measure as compared to variance to optimise a portfolio has significant merit. Speculatively, because of very short return cycles posited by Dow Jones theory (Schannep, 2012), strategies recalculated daily may suffer from the introduction of an EWMA, since the smoothing introduced by the EWMA would allow the introduction of assets with high short-term volatility which would otherwise be penalised by an unsmoothed risk measure. Further, since MPT is momentum-based, since weighting of investment into an asset is proportional to its previous-period return, using an EWMA to smooth returns may simply result in too-slow reaction to a short-duration return cycle.

4.3.3 The benefits of using an EWMA

The study attempted to address the criticism of MPT that risk and return do not persist into the future through estimation of the next period's risk and return using an EWMA. These estimates were then used in lieu of the mean return and previous period's risk measure as inputs into the calculation of the efficient frontier.

Table 4.5 shows the difference in return between EWMA-optimised strategies and raw risk measure optimised strategies across the 70 strategies of GICS tier, calculation period and the risk measure. Categories in which EWMA-optimised strategies outperformed raw risk measure-optimised strategies are shaded.

Table 4.5

Difference in annualised return between EWMA-optimised and equivalent raw risk measure-optimised strategies

		Risk measure optimised against										Average
Tier	Calculation Period	V	SD	DSV	DSSD	MAD	SAD	VaR0.1	VaR0.05	CVaR0.1	CVaR0.05	
Sector	Daily	1.66%	-2.27%	-3.78%	-3.19%	-2.00%	-3.88%	-0.49%	0.17%	-0.80%	-0.89%	-1.55%
	Weekly	0.05%	0.62%	2.97%	-2.49%	3.98%	2.24%	3.06%	4.54%	-6.50%	-4.37%	0.41%
	Monthly	-3.72%	0.44%	-2.80%	-2.43%	-1.11%	-3.56%	1.63%	0.12%	-2.01%	0.14%	-1.33%
Industry Group	Weekly	1.30%	4.39%	4.95%	5.47%	7.37%	7.88%	7.75%	4.15%	-0.53%	-0.50%	4.22%
	Monthly	-5.22%	-3.98%	-1.31%	-0.11%	-2.70%	-1.88%	-2.47%	-1.99%	-3.57%	-5.03%	-2.83%
Industry	Weekly	0.05%	0.62%	2.97%	-2.49%	3.98%	2.24%	3.06%	4.54%	-6.50%	-4.37%	0.41%
	Monthly	-3.72%	0.44%	-2.80%	-2.43%	-1.11%	-3.56%	1.63%	0.12%	-2.01%	0.14%	-1.33%
Average		-1.37%	0.04%	0.03%	-1.10%	1.20%	-0.08%	2.02%	1.66%	-3.13%	-2.13%	-0.01%

Key:
EWMA return higher than Raw return
Raw return higher than EWMA return

Note. Figures shown are, for each investment strategy, the EWMA strategy annualised return less the raw strategy annualised return. Strategies for which EWMA strategy return exceeded the raw strategy return, and by implication for which application of an EWMA was beneficial, are shaded in grey.
Source: Author's calculations

Overall, there is very little difference between portfolios optimised using an EWMA and those optimised using raw risk measures, with raw risk measure-optimised portfolios outperforming EWMA-optimised portfolios by an average of 0.01%. However, it is evident that certain risk measures benefit from the use of an EWMA, while others do not. Value at risk-optimised strategies appear to benefit greatly, while those optimised using conditional value at risk or downside semi deviation are prejudiced by the use of an EWMA. Strategies optimised using the remainder of the risk measures show mixed results, with some GI CS tiers and calculation periods performing better then their raw risk measure equivalent, and others worse.

The lack of improvement in the EWMA forecast-based strategies' returns or risk-adjusted returns for the majority of risk measures is likely due to the forecasting methodology (EWMA) providing little improvement in forecasts as compared to the previous period's risk or return (the raw risk measure and mean return).

Despite the lack of improvement in the forecasts themselves, portfolios optimised using an EWMA of value at risk and conditional value at risk show significantly different returns from those optimised using their raw equivalents. Lohre (2010) stated that optimising a portfolio with respect to conditional value at risk is more effective than doing so with respect to value at risk, usually leading to well optimised portfolios in terms of value at risk. Part of the reason for this is that conditional value at risk, unlike value at risk, is a coherent risk measure, satisfying the subadditivity requirement (Lim et al., 2011). Conditional value at risk therefore presents as a smoothed value at risk and is more tractable and effective for optimising portfolios. Corroborating this, the study found that, on average, raw conditional value at risk-optimised portfolios outperformed raw value at risk-optimised portfolios by an average of 1%. However, application of an EWMA also smooths the risk measure, albeit in a different way. Again, applying such a smoothing mechanism to value at risk when optimising portfolios results in better returns, with EWMA value at risk-optimised portfolios outperforming raw value at risk optimised portfolios by an average of 1.8%. Applying a smoothing mechanism to value at risk also ameliorates the impact its lack of smoothness has on optimisation (Gaivoronski, 2005, Lwin et al., 2017)[3]. Despite the benefits of smoothing on value at risk, oversmoothing results in loss of information and eventually poorer optimisation[4]. Such oversmoothing may be the reason that EWMA conditional value at risk optimised portfolios perform poorly, since they represent a smoothing (through the application of the EWMA) of an already-smoothed value at risk (that is, the conditional value at risk), and therefore a double loss of information.

Also evident from Table 4.5, is that strategies recalculated weekly benefit from the application of an EWMA, while those calculated daily or monthly do not. This might be attributable to two reasons. Firstly, they may simply be a sweet spot for smoothing around a recalculation period of one week. Short duration return cycles, as proposed in Dow Jones theory (Schannep, 2012), may result in poor returns when trades are based on a smoothed, slightly longer-duration signal provided by the EWMA. Similarly, the inverse may be true in the case of monthly portfolio recalculation, if trades are based on a shorter-duration EWMA than the rebalancing period. Secondly, the calculation of squared returns required for the calculation of most risk measures' EWMAs is based on the difference between the return and zero, rather than between the return and the mean return[5]. Particularly in the case of long recalculation periods, this might result in loss of information when the mean return over the period has in fact been significantly different to zero.

4.3.4 Semi-absolute deviation-optimised strategy performance

Having established that the best-performing risk measure in the study for optimising an investment strategy is semi-absolute deviation, this section examines the performance thereof in more detail. The strategy is based on sector-level data, and compared to the S&P Total Return Global Index as a benchmark. The Global Index is the market capitalisation-weighted aggregation of the companies which comprise the sector indices and therefore represents a reasonable alternative to optimisation-based diversification across the

---

[3] This impact is evident in Figure 7.1 in Appendix B, which shows the relative tractability of Value at Risk and its EWMA in optimising a single portfolio.
[4] An example of this can be seen in Table 8.3 in Appendix C, wherein longer training periods result in poorer-performing portfolios.
[5] See 3.5.4 (Calculating the EWMA) for the reason this is the case.

constituent sector-based indices. The Global Index achieved an annualised return of 12.1% and a Sharpe ratio of 0.74 over the study period. The strategy is also compared to strategies optimised using the traditional risk measure used in MPT, variance, to examine in more detail the impact of using a different risk measure.

Figure 4.1

Semi-absolute deviation-optimised strategy performance

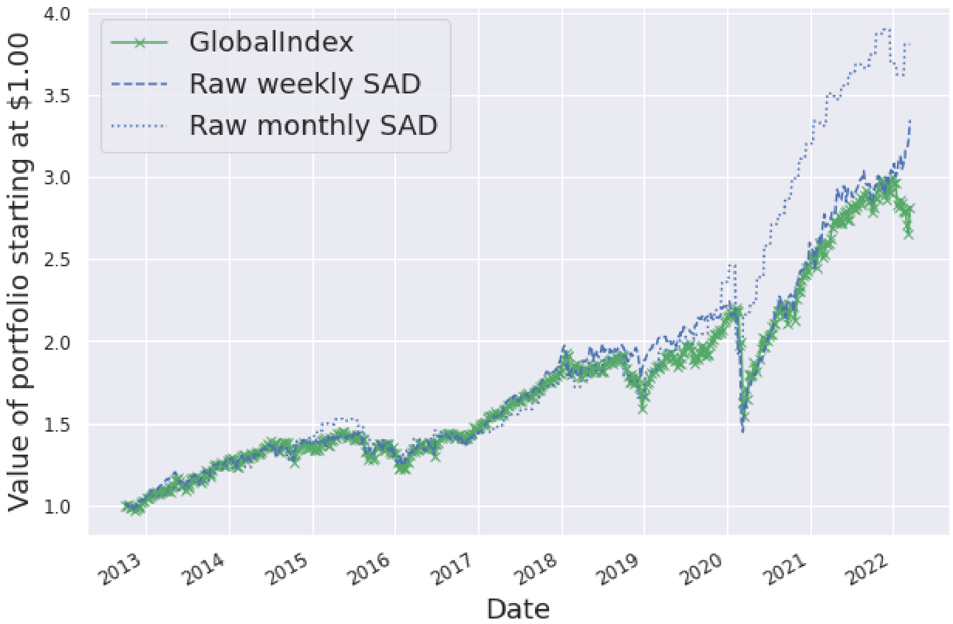

Panel A: Performance of raw risk measure semi-absolute deviation-optimised strategies

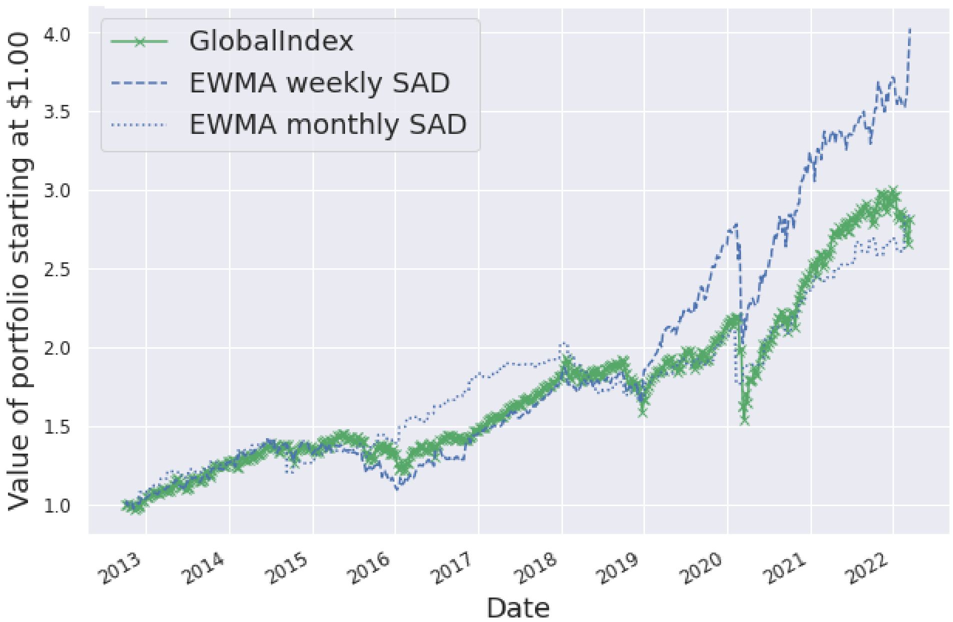

Panel B: Performance of EWMA semi-absolute deviation-optimised strategies

Note. Graphs of strategy returns compared to the benchmark S&P 1200 Global Index. The top panel shows returns for strategies calculated weekly and monthly using raw risk returns, while the bottom panel shows those for strategies calculated using EWMAs.

Source: Author's calculations

Comparing first to the benchmark, Figure 4.1 shows the return on a $1 portfolio on the y-axis, while the horizontal axis shows the date. The performance of raw semi-absolute deviation and EWMA-optimised strategies is shown for both a weekly and monthly recalculation period. As shown in Table 4.6, the annualised return of both weekly and monthly recalculated raw risk measure-optimised strategies outperforms the benchmark, as does the weekly EWMA-optimised strategy. However, the benchmark outperforms the monthly EWMA-optimised strategy by approximately 17bp (annualised). The Sharpe ratio of all four semi-absolute deviation-optimised strategies exceeds that of the benchmark.

As can be seen from the period between early 2020 and the beginning of 2022 in Figure 4.1, the semi-absolute deviation-optimised strategies generally performed particularly well during the recovery period of the market following the CoVID-19 pandemic. This conforms with the expectation that portfolios optimised using semi-absolute deviation will not penalise the weighting of assets with high upside deviations in return relative to the mean return, which is clearly advantageous in a bullish market.

Table 4.6

Annualised returns, fees, Sharpe ratios and risk measures for variations of SAD-optimised strategies

						Annualised risk measures										
		Mean	Annualised Return	Fees	Sharpe Ratio	Max Drawdown	V	SD	DSV	DSSD	MAD	SAD	VaR0.1	VaR0.05	CVaR0.1	CVaR0.05
	GlobalIndex	12.08%	11.56%		0.74	29.87%	2.17%	17.12%	14.44%	19.22%	25.34%	33.65%	14.74%	2.93%	9.72%	11.15%
Raw	Weekly	13.88%	13.50%	0.100%	0.85	35.46%	2.27%	17.58%	13.37%	19.59%	25.18%	34.46%	15.06%	3.09%	9.55%	10.56%
	Monthly	15.21%	15.18%	0.099%	1.11	18.00%	1.71%	15.13%	10.46%	18.33%	20.24%	27.64%	13.07%	2.29%	9.46%	10.59%
EWMA	Weekly	15.83%	15.74%	0.104%	1.00	27.29%	2.25%	16.58%	13.84%	19.68%	23.95%	32.15%	15.01%	2.75%	9.89%	10.25%
	Monthly	11.91%	11.62%	0.100%	0.89	17.83%	1.52%	14.61%	9.97%	13.40%	19.18%	27.35%	12.31%	2.14%	8.52%	9.94%

Note. Returns, fees, Sharpe ratio and risk measures for the four semi-absolute deviation-optimised strategies. V = variance, SD = standard deviation, DSV = downside semivariance, DSSD = downside semideviation, MAD = mean absolute deviation, SAD = semi-absolute

deviation, VaR0.1 = 10% value at risk, VaR0.05 = 5% value at risk, CVaR0.1 = 10% conditional value at risk, CVaR0.05 = 5% conditional value at risk.
Source: Author’s calculations

In addition to annualised return and Sharpe ratio, Table 4.6 shows the annualised returns, fees and annualised risk measures for the four semi-absolute deviation-optimised strategies. Annualised risk measures were obtained by multiplying the risk measure by the number of periods per year (variance and downside semivariance) or the square root thereof (all other risk measures), while fees are expressed as an annualised reduction in return.

Observing the small difference in the fees between weekly and monthly strategies, it seems that the increase in cost of more frequent trading inherent in a weekly, as compared to a monthly, recalculation of the strategy, is of relatively little impact, and the majority of the fee is comprised of the ETF’s total expense ratio, rather than trading fees.

Comparing the risk of the different strategies, the observation made in Section 4.3.2 (Comparison of all strategies), that a strategy’s impact on risk is generally consistent across all risk measures, not just the risk measure being optimised over, is corroborated. Accordingly, the monthly EWMA-optimised strategy displays the lowest risk across all risk measures, albeit at the cost of much lower returns than the other strategies. This is followed by the monthly raw risk measure-optimised strategy and then the weekly EWMA-optimised strategy. The riskiest strategies are the weekly raw risk measure-optimised strategy and the benchmark, displaying similar levels of risk; however, of the two, the weekly raw risk-measure optimised strategy has a much higher return.

Comparing the results of semi-absolute deviation-optimised strategies against those optimised using variance, the traditional risk measure used in modern portfolio theory, the former tends to outperform the latter. Figure 4.2 shows the performance of weekly semi-absolute deviation-optimised strategies compared to variance-optimised strategies. While between 2013 and 2017 the variance-optimised strategies matched or outperformed the semi-absolute deviation-optimised strategies, the latter delivered exceptional returns thereafter, particularly during the recovery period following the CoVID-19 pandemic, doubling the portfolio in two years. As stated above, this conforms with the expectation that portfolios optimised using semi-absolute deviation will not penalise the weighting of assets with high upside deviations in return relative to the mean return. Conversely, optimising using variance will confer a penalty on these assets because it includes upside deviations relative to the mean return into the risk measure. Further, those assets with the largest upside deviations will be most penalised due to variance squaring the deviations. The impact of this doubly-disadvantageous effect during a bull market is clearly evident in the difference in portfolio growth during the post-CoVID recovery period.

Figure 4.2

Weekly semi-absolute deviation- vs variance-optimised strategy performance

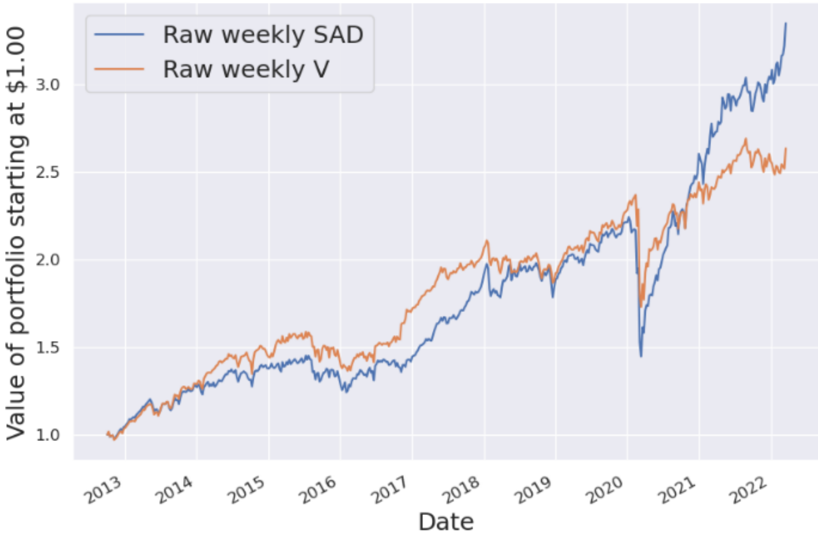

Panel A: Performance of weekly raw risk measure semi-absolute deviation-optimised strategies

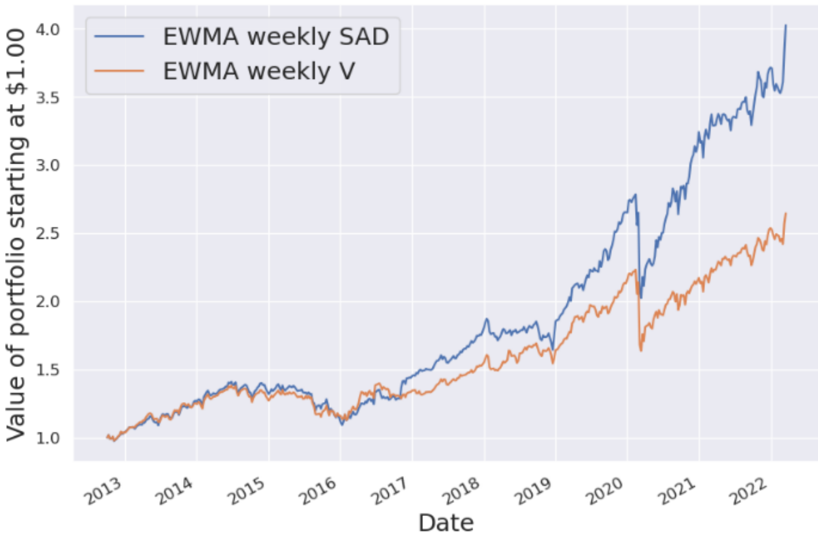

Panel B: Performance of weekly EWMA semi-absolute deviation-optimised strategies

Note. Graphs comparing the returns of weekly semi-absolute deviation-optimised strategies to those of variance-optimised strategies. The top panel shows returns for strategies calculated weekly and monthly using raw risk returns, while the bottom panel shows those for strategies calculated using EWMAs.

Source: Author's calculations

The same is true of monthly strategies. Figure 4.3 shows the performance of the same strategies recalculated on a monthly instead of weekly basis. Again, the semi-absolute deviation strategies outperform the variance-optimised strategies. Unlike the weekly strategies, the variance-optimised strategy never materially exceeds the performance of the semi-absolute deviation-optimised strategy.

Figure 4.3

Monthly semi-absolute deviation- vs variance-optimised strategy performance

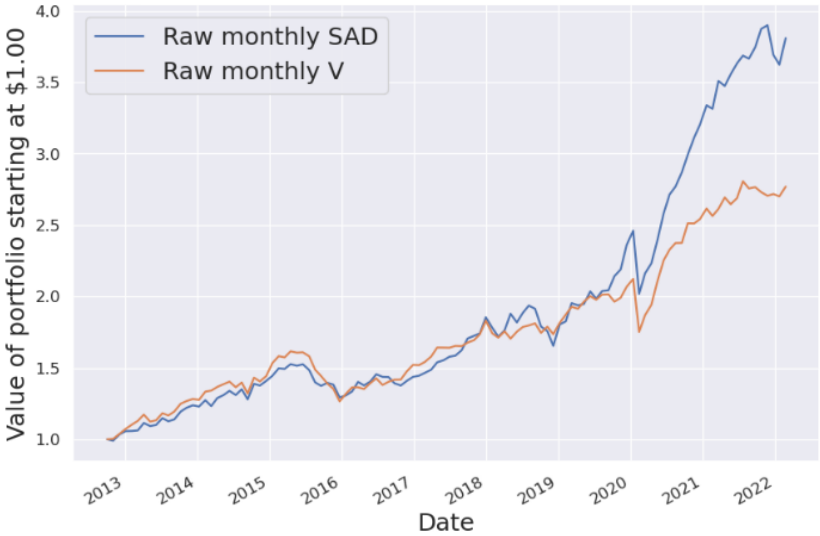

Panel A: Performance of monthly raw risk measure semi-absolute deviation-optimised strategies

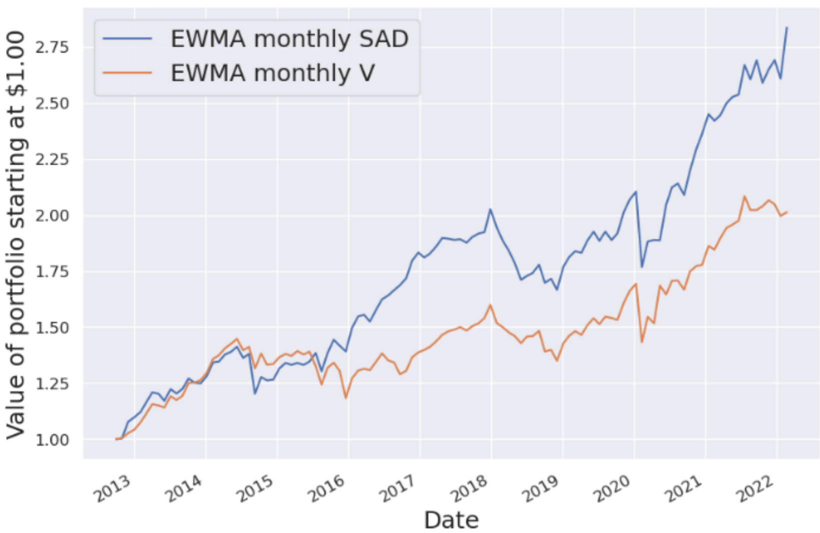

Panel B: Performance of monthly EWMA semi-absolute deviation-optimised strategies

Note. Graphs comparing the returns of monthly semi-absolute deviation-optimised strategies to those of variance-optimised strategies. The top panel shows returns for strategies calculated weekly and monthly using raw risk returns, while the bottom panel shows those for strategies calculated using EWMAs.

Source: Author's calculations

While variance-optimised strategies are disadvantaged during a bull market, the opposite is certainly not true: the performance of semi-absolute deviation-strategies during normal market (that is, non-bull market) conditions does not appear to be materially worse than that of variance-optimised strategies, and better in some cases. For example, while weekly raw variance-optimised strategies outperform weekly raw semi-absolute deviation-optimised strategies in the normal market period leading up to 2020, the opposite is true for monthly EWMA-optimised strategies, and weekly EWMA- and monthly raw-optimised strategies show very similar performance during this period.

The conjunction of better bull-market performance and similar or slightly better normal-market performance results in semi-absolute deviation-optimised strategies providing better overall performance than variance-optimised strategies.

4.4 SUMMARY OF FINDINGS

One hundred and forty investment strategies were run across the combinations of ten different risk measures, three GICS tiers, three lengths of recalculation period and two applications of EWMA vs raw risk measure. Of all the risk measures, strategies optimised using semi-absolute deviation performed best in terms of absolute returns and Sharpe ratios, followed by mean absolute deviation-optimised strategies and semi-absolute deviation-optimised strategies.

Strategies optimised using the traditional measure for MPT, variance, tended to perform poorly in terms of both absolute returns and Sharpe ratio.

The characteristics of each risk measure influence the penalty applied to the choice of including indices into the portfolio, which would otherwise be based solely on return. In particular, portfolios optimised using asymmetrical risk measures outperformed portfolios optimised using symmetrical risk measures on average, lending credibility to the criticism of MPT that it minimises upside risk alongside minimising downside risk. Further, optimising portfolios using risk measures which incorporate a square of deviations penalised large movements away from the mean return, as compared to risk measures which incorporated unsquared deviations. In a market environment with more positive large swings than negative large swings, or where large positive market corrections followed large negative market movements, portfolios optimised using unsquared risk measures benefited. A similar effect was evident in portfolios optimised using value at risk and conditional value at risk, which also penalised the inclusion of indices with larger volatility and therefore more outliers.

Application of an EWMA to forecast returns and risk measures provided inconclusive results, with the average annualised performance and Sharpe ratios across all categories of raw risk measure-optimised strategies being only slightly higher than the average annualised return across all categories of EWMA-optimised strategies.

CHAPTER 5 – CONCLUSION

5.1 INTRODUCTION

MPT suffers a number of criticisms, including that risk and return do not persist into the future, that returns are not statically normally distributed and that the risk measure used to optimise portfolios minimises upside risk as well as downside risk. This study examines the use of alternative risk measures to optimise portfolios as well as the use of different risk forecasting techniques to those used by MPT for forecasting those risk measures.

The literature contains studies of alternative risk measures applied to MPT, examples of which include downside semivariance and downside semi deviation, mean absoute deviation and semi-absolute deviation, value at risk and conditional value at risk. Although the results of these studies are varied, they show that the use of alternative risk measures to optimise portfolios often results in better returns than the use of variance.

This study corroborates these findings: variance is not an ideal risk measure to optimise against. In over two thirds of the 140 strategies examined in the study, using an alternative risk measure to variance to optimise portfolios resulted in better annualised returns, with the single instance in which variance-optimised strategies performed best being daily EWMA-optimised strategies on sector-level data.

Portfolios optimised using alternative risk measures outperformed portfolios optimised using variance for a variety of reasons. Firstly, the criticism of MPT that using variance to optimise portfolios minimises upside risk as well as downside risk (Boasson et al., 2011; Byrne & Lee, 2004; Cardoso et al., 2019; Hunjra et al., 2020) appears to have merit: On average, downside semivariance-optimised strategies in the study outperformed variance-optimised strategies, downside semideviation-optimised strategies outperformed standard deviation-optimised strategies, and semi-absolute deviation strategies outperformed mean absolute deviation strategies by circa 1% p.a.

Secondly, the use of variance to optimise portfolios penalises return profiles which contain more outliers relative to the use of risk measures which are a function of unsquared differences between returns and average return. This leads to higher average returns for standard deviation-optimised strategies than variance-optimised strategies, and similarly for semivariance-optimised strategies and downside semideviation-optimised strategies. The impact is even more noticeable when comparing the performance of mean absolute deviation optimised strategies to variance optimised strategies, and that of semi-absolute deviation optimised strategies to downside semivariance optimised strategies.

Having established that the use of alternative risk measures to optimise portfolios is indeed worthwhile, the question then becomes one of identifying which risk measures perform best. The study shows that of all the risk measure-optimised strategies, semi-absolute deviation-optimised strategies had the best risk-adjusted returns most frequently. In particular, absolute deviation-optimised strategies performed well when coupled with an EWMA.

Addressing the criticism of MPT that risk and return do not persist into the future (Rigamonti and Lučivjanská, 2022; Iyiola et al., 2012), the study investigated the use of an EWMA to better forecast the risk measures optimised against. The results were inconclusive, with the average annualised performance across all categories of raw risk measure-optimised strategies being only 0.07% higher than the average annualised return across all categories of EWMA-optimised strategies, and the difference between their Sharpe ratios being only 0.017. At a more detailed level, weekly EWMA-optimised strategies outperformed weekly raw risk-optimised strategies, while daily and monthly raw risk-optimised strategies outperformed their respective EWMA-optimised strategies.

The study attempted to avoid the criticism of MPT that it assumes that the probability distributions of asset returns are static in time by refitting the EWMA’s decay factor based upon a rolling 3-month window. However, it was unable to conclude that using a decay factor (lambda) which was constantly refitted to the past three month’s risk resulted in better performance of EWMA-optimised strategies than EWMA-optimised strategies which used a constant decay factor, since the risk-adjusted returns of weekly strategies using a variance decay factor outperformed those using a constant decay factor, while the opposite was true of monthly strategies.

The study showed that a daily recalculation period did not help absolute- or risk-adjusted returns. This was only partially because of higher trading frequency and fees. Strategies with weekly and monthly recalculation periods performed better on both an absolute- and risk-adjusted return basis. Linking to recalculation periods, training periods were typically optimal in the three- to six-month range. Longer and shorter training periods typically resulted in poorer annualised returns and Sharpe ratios for both raw-risk and EWMA-optimised strategies. Another factor investigated was the number of assets in the investment universe. Although it was not possible to directly compare strategies with Sector-level inputs to those with Industry Group- or Industry-level inputs because of differences in index construction, direct comparison of strategies with Industry Group-level and Industry-level inputs was possible. In most instances, Industry Group-level strategies achieved better performance than Industry-level strategies. Consequently, increasing the number of investible assets does not necessarily improve performance.

5.2 IMPLICATIONS

The study addresses some of the criticisms of MPT. In particular, it addresses the criticism that MPT uses a symmetrical risk measure (variance) thereby penalising upside risk, by optimising against different single-sided risk measures. In each instance, the annualised returns over the study period of strategies optimised using single-sided risk measures are higher than those of the strategies optimised against their analogous symmetrical risk measures, with downside semivariance-optimised strategies outperforming variance-optimised strategies slightly, downside semideviation-optimised strategies outperforming standard deviation-optimised strategies, and semi-absolute deviation strategies outperforming mean absolute deviation strategies by circa 1% p.a. each. This is a meaningful outcome in that it provides several robust and suitable alternatives to variance which can be used in MPT which each address this criticism. An implication of this outcome is that investors optimising risk and return using symmetrical risk measures, whether they be the traditional variance or another risk measure, may well be obtaining suboptimal results and should instead use the asymmetrical equivalent risk measure in their optimisation in order to improve their returns.

The study also addresses the criticism of MPT that it assumes normally-distributed returns, by, for the most part, demonstrating that a fully non-parametric approach to optimising risk and return based on historical data achieves reasonable results. This confirms that investors using MPT or MPT-like optimisation have the option of using numerical methods algorithms to optimise their portfolios instead of closed-form mathematical solutions. Doing so will engender two advantages upon investors: firstly, it is possible to apply the same methods to different risk measures, allowing investors a wider, more flexible and quicker-to-develop toolkit; and secondly, it is possible to include market-related complexity into the optimisation which is not possible with closed-form solutions, such as the impact of fees. By implication, given the availability of commercial optimisation software, investors using MPT-like optimisation should use numerical-methods-based optimisation and not closed-form solutions.

The study attempted to address the criticism of MPT that it assumes risk and return persist into the future by comparing the traditional approach of optimising risk and return in the most recent historical period to an approach in which the optimisation occurs on the next period's risk and return forecast. This was not altogether successful, with no consistent improvement in the forecast-based strategies' returns or risk-adjusted returns. This was likely due to the forecasting methodology providing little improvement in forecasts of risk and return themselves, as compared to the previous period's risk or return. Given these results, better forecasting techniques than an EWMA for either or both of risk and return may be required in order to improve the efficacy of MPT-like methods and in so doing, address the criticism of MPT that it incorrectly assumes risk and return persist into the future.

The study set out to investigate strategies optimised using several common risk measures for a common period and set of constituent assets, given that most similar studies investigate only a few risk measures. This will afford future researchers a direct point of comparison between different strategies with minimal additional influences impacting upon such comparisons. Further, the study applied MPT-like optimisation to a sub asset class which has been the subject of relatively few similar exercises, providing future researchers with wider data coverage. Presently studies on alternative risk measures applied to MPT cover stocks, bonds and options, but very few instances of indices (as proxies for ETFs). Demonstrating the efficacy of MPT-like optimisation using alternative risk measures for a hither-to under-researched asset class improves the case for investors using such optimisation, as it demonstrates that it is more generally applicable than was indicated prior to the research having been undertaken.

Lastly, the study detailed the process by which it arrived at the strategy hyperparameters, which will assist future researchers undertaking similar studies.

5.3 FUTURE AVENUES OF RESEARCH

While the study achieved most of its purposes, it failed to show a definitive improvement in returns by optimising strategies based on forecast risk and returns, likely due to the forecasting methods used providing poor forecasts of risk and return themselves. This provides an opportunity for future research, as to whether better forecasts may be incorporated into the optimisation to improve the strategy results. An example of this may be to use GARCH forecasting models, or to build machine learning models to forecast risk and return as a function of qualitative inputs and use the outputs thereof as inputs into MPT-like optimisation and diversification. If intra-day data is available (perhaps for different asset classes), shortening the prediction horizon might yield better predictive accuracy of risk and return, leading to suitable trading strategies, if not investment strategies.

A mix of value-based investing and quantitative investment techniques could be investigated. In particular, purely quantitative methods are subject to the criticism that risk and return do not persist into the future, which the study attempted to address through the application of an EWMA. By using value-based investment techniques which take the characteristics of the underlying asset itself into account, rather than only the historical movements of its share price, a distribution for next-period returns might be constructed which could, in turn, be an input into MPT-like quantitative methods. Similarly, other proxies or drivers for future returns, such as arbitrage pricing could be incorporated into the model.

As discussed in the study, the application of MPT like methods to as-yet uninvestigated asset classes improves understanding of how and whether these methods are effective. Further corroboration of the relative efficacy of each risk measure in building investment strategies could be sought by applying the methods used in the study to different asset classes. A possible example could be to apply it to several different random subsets of the S&P 500, in order to avoid the impact of having chosen a specific set of assets which happens to result in a specific but spurious outcome.

The study used long-only and fully-invested constraints when constructing each portfolio. These constraints could be relaxed to investigate whether this improves strategy performance. Different rule-based mechanisms for balancing long-and short positions could be investigated. Similarly, the impact of making capital invested in the next period a function of expected risk/ return could be investigated.

Another avenue for future research relates to further optimisation of hyper parameters and analysis of how and why each hyperparameter impacts on the success of each strategy. In particular, rebalancing period and training period could be investigated in significantly more detail. If intra-day data is available (perhaps for different asset classes), moving the strategy from an investment strategy to a trading strategy can be investigated.

REFERENCES

Abdul Rehman, A., & Alharthi, K. (2016). An introduction to research paradigms. *International Journal of Educational Investigations, 3,* 51-59.

Ahmadi-Javid, A., & Fallah-Tafti, M. (2019). Portfolio optimization with entropic value-at-risk. *European Journal of Operational Research, 279*(1), 225-241.

Alexander, C. (2008). *Market risk analysis*. John Wiley & Sons.

Alexander, G. J., & Baptista, A. M. (2002). Economic implications of using a mean-VaR model for portfolio selection: A comparison with mean-variance analysis. *Journal of Economic Dynamics and Control*, *26*(7-8), 1159-1193.

Andreu, J., & Torra, S. (2009). Optimal market indices using value-at-risk: a first empirical approach for three stock markets. *Applied financial economics, 19*(14), 1163-1170.

Boasson, V., Boasson, E., & Zhou, Z. (2011). Portfolio optimisation in a mean-semivariance framework. *Investment Management and Financial Innovations*, *8*, 58-68.

Bodie, Z., Kane, A., & Marcus, A., J. (2017). *Investments: tenth edition*. McGraw-Hill.

Bollen, B. (2015). What should the value of λ be in the exponentially weighted moving average volatility model? *Applied Economics*, *47*(8), 853-860.

Byrne, P., & Lee, S. (2004). Different risk measures: different portfolio compositions? *Journal of Property Investment & Finance*, *22*(6), 501-511.

Cardoso, R., Barroso, B., Oliveira, M., & Paiva, F. (2019). Analysis of risk measures in multiobjective optimisation portfolios with cardinality constraint. *Brazilian Review of Finance*, *17*(3), 26-46.

Chang, T.-J., Yang, S.-C., & Chang, K.-J. (2009). Portfolio optimization problems in different risk measures using genetic algorithm. *Expert Systems with Applications, 36*(7), 10529-10537.

Chen, D., Wu, Y., Li, J., Ding, X., & Chen, C. (2022). Distributionally robust mean-absolute deviation portfolio optimisation using wasserstein metric. *Journal of Global Optimisation*, 1-23.

Dachraoui, K. (2018). On the Optimality of Target Volatility Strategies. *The Journal of Portfolio Management*, *44*, 58-67.

Danielsson, J. (2011). *Financial risk forecasting: the theory and practice of forecasting market risk with implementation in R and Matlab*. John Wiley & Sons.

Dao, B. (2014). Optimal Portfolio with Options in a Framework of Mean-CVaR. *SSRN Electronic Journal*. Retrieved from https://dx.doi.org/10.2139/ssrn.2543722

Di Clemente, A. (2002). The empirical value at risk-expected return frontier: A useful tool of market risk managing. *Centro Interdipartimentale sul Diritto e l'Economia dei Mercati* (11).

Engle, R. (2001). GARCH 101: The Use of ARCH/GARCH Models in Applied Econometrics. *Journal of Economic Perspectives*, *15*(4), 157-168.

Engle, R., & Patton, A. (2001). What good is a volatility model? *Quantitative Finance*, *1*, 237-245.

Engle, R. F. (1982). Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation. *Econometrica*, *50*(4), 987-1007.

Estrada, J. (2006). Downside Risk in Practice. *Journal of Applied Corporate Finance, 18*(1), 117-125.

Fan, R., & Lee, J. H. (2017). Predictive Quantile Regressions Under Persistence and Conditional Heteroskedasticity. *SSRN Electronic Journal*. Retrieved from https://dx.doi.org/10.2139/ssrn.3016449

French, K. R., Schwert, G. W., & Stambaugh, R. F. (1987). Expected stock returns and volatility. *Journal of Financial Economics*, *19*(1), 3-29.

Gaivoronski, A. A. P., Georg. (2005). Value-at-Risk in Portfolio Optimisation Properties and Computational Approach. *Journal of Risk*, *7*, 1-31.

Garcia, R. C., González, V., Contreras, J., & Custodio, J. E. (2017). Applying modern portfolio theory for a dynamic energy portfolio allocation in electricity markets. *Electric Power Systems Research, 150*, 11-23.

Gosier, K. M., Ananth; Serbin, Vitaly; Yang, Jian. (2005). Towards Better Risk Forecasts. *Journal of Portfolio Management*, *31*(3), 82-91.

Guo, X., Chan, R., Wong, W.-K., & Zhu, L. (2019). Mean–variance, mean–VaR, and mean–CVaR models for portfolio selection with background risk. *Risk Management, 21*. Retrieved from https://doi.org/10.1057/s41283-018-0043-2

Hafsa, H. (2015). CVaR in Portfolio Optimisation: An Essay on the French Market. *International Journal of Financial Research*, *6*, 101-111.

Hawley, J., & Lukomnik, J. (2018). The third, systems stage of corporate governance: Why institutional investors need to move beyond modern portfolio theory. *SSRN Electronic Journal*. Retrieved from https://dx.doi.org/10.2139/ssrn.3127767

Hox, J., & Boeije, H. (2005). Data collection, primary versus secondary. *Encyclopedia of Social Measurement*, *1*, 593-599.

Hunjra, A. I., Alawi, S. M., Colombage, S., Sahito, U., & Hanif, M. (2020). Portfolio Construction by Using Different Risk Models: A Comparison among Diverse Economic Scenarios. *Risks*, *8*(4), 1-23.

Iyiola, O., Munirat, Y., & Nwufo, C. (2012). The modern portfolio theory as an investment decision tool. *Journal of Accounting and Taxation*, *4*, 19-28.

Jacobsen, B. (2005). The Use of Downside Risk Measures in Portfolio Construction and Evaluation. *Computing in Economics and Finance*, *5*. Retrieved from http://repec.org/sce2005/up.5184.1102515146.pdf

Kasenbacher, G., Lee, J., & Euchukanonchai, K. (2017). Mean-variance vs. mean-absolute deviation: A performance comparison of portfolio optimization models. [Thesis, University of British Columbia]. Retrieved from https://www.researchgate.net/profile/Geoffrey-Kasenbacher/publication/330358884_Mean-Variance_vs_Mean-Absolute_Deviation_A_Performance_Comparison_of_Portfolio_Optimization_Models/links/5c3c4239a6fdccd6b5ab3e4e/Mean-Variance-vs-Mean-Absolute-Deviation-A-Performance-Comparison-of-Portfolio-Optimization-Models.pdf

Kolm, P. N., Tütüncü, R., & Fabozzi, F. J. (2014). 60 Years of portfolio optimisation: Practical challenges and current trends. *European Journal of Operational Research*, *234*(2), 356-371.

Konno, H., & Yamazaki, H. (1991). Mean-absolute deviation portfolio optimization model and its applications to Tokyo stock market. *Management science, 37*(5), 519-531.

Lam, W. H., Jaaman, S. H., & Isa, Z. (2010). An empirical comparison of different risk measures in portfolio optimization. *Business and Economic Horizons, 1*. Retrieved from https://doi.org/10.15208/beh.2010.06

Li, P., Han, Y., & Xia, Y. (2016). Portfolio optimization using asymmetry robust mean absolute deviation model. *Finance Research Letters, 18*, 353-362.

Lim, A. E. B., Shanthikumar, J. G., & Vahn, G.-Y. (2011). Conditional value-at-risk in portfolio optimisation: Coherent but fragile. *Operations Research Letters*, *39*(3), 163-171.

Liu, F., Tang, X., & Zhou, G. (2019). Volatility-Managed Portfolio: Does It Really Work? *The Journal of Portfolio Management*, *46*, 38-50.

Lohre, H. N., Thorsten, & N., Winterfeldt, T. (2010). *Portfolio Construction with Downside Risk.* SSRN Electronic Journal, Retrieved from http://dx.doi.org/10.2139/ssrn.1112982

Lukomnik, J., & Hawley, J. (2021). *Moving Beyond Modern Portfolio Theory: Investing That Matters*. Routledge.

Lwin, K., Qu, R., & MacCarthy, B. (2017). Mean-VaR Portfolio Optimisation: A Nonparametric Approach. *European Journal of Operational Research*, *260*, 1-29.

Maeso, J.-M., & Martellini, L. (2020). Measuring Portfolio Rebalancing Benefits in Equity Markets. *The Journal of Portfolio Management*, *46*, 94-109.

Malkiel, B. G. (1973). *A random walk down Wall Street: the time-tested strategy for successful investing*. W.W. Norton.

Markowitz, H. (1952). Portfolio Selection. *The Journal of Finance*, *7*, 77-91.

Markowitz, H., & Cowles, F. (1959). *Portfolio selection: efficient diversification of investments* (1st , 2nd printing. ed.). Yale University Press.

Miskolczi, P. (2017). Note on simple and logarithmic return. *Applied Studies In Agribusiness And Commerce*, *11*, 127-136.

Montgomery, D., Jennings, C., & Kulahci, M. (2015). *Introduction to Time Series Analysis and Forecasting*. Wiley.

Moreno, D., Marco, P., & Olmeda, I. (2005). Risk forecasting models and optimal portfolio selection. *Applied Economics*, *37*, 1267-1281.

Morgan Stanley Capital International. (2020). *Global Industry Classification Standard (GICS) Methodology*. Retrieved from https://www.msci.com/documents/1296102/11185224/GICS+Methodology+2023.pdf

Nguyen, L., Nguyen, L., & Adegbite, E. (2018). Does Mean-CVaR Outperform Mean-Variance? A Practical Perspective. *SSRN Electronic Journal*. Retrieved from https://dx.doi.org/10.2139/ssrn.3143827

Omisore, I., Yusuf, M. O., & Christopher, N. (2012). The modern portfolio theory as an investment decision tool. *Journal of Accounting and Taxation*, *4*(2), 19-28.

Ortobelli, S., Rachev, S. T., Stoyanov, S., Fabozzi, F. J., & Biglova, A. (2005). The proper use of risk measures in portfolio theory. *International Journal of Theoretical and Applied Finance, 8*(08), 1107-1133.

Ployhart, R., & Vandenberg, R. (2010). Longitudinal Research: The Theory, Design, and Analysis of Change. *Journal of Management*, *36*, 94-120.

Rachev, S., Ortobelli, S., Stoyanov, S., Fabozzi, F., & Biglova, A. (2008). Desirable properties of an ideal risk measure in portfolio theory. *International Journal of Theoretical and Applied Finance (IJTAF)*, *11*, 19-54.

Ramos, H. P., Righi, M. B., Guedes, P. C., & Müller, F. M. (2023). A comparison of risk measures for portfolio optimization with cardinality constraints. *Expert Systems with Applications, 228*. Retrieved from https://dx.doi.org/10.2139/ssrn.4141301

Richardson, M., & Smith, T. (1993). A Test for Multivariate Normality in Stock Returns. *The Journal of Business*, *66*, 295-321.

Rigamonti, A., & Lučivjanská, K. (2022). Mean-semivariance portfolio optimisation using minimum average partial. *Annals of Operations Research*. Retrieved from https://doi.org/10.1007/s10479-022-04736-x

Righi, M., & Borenstein, D. (2018). A simulation comparison of risk measures for portfolio optimization. *Finance Research Letters, 24*, 105-112.

Rocha, E. (2016). Security selection in post-modern portfolio theory: an application to the European stock market. [Masters Thesis, Instituto Superior de Economia e Gestão]. Retrieved from https://www.repository.utl.pt/handle/10400.5/13094

Rockafellar, R., & Uryasev, S. (2000). Optimisation of Conditional Value-At-Risk. *Journal of Risk*, *2*, 21-42.

Roman, D., & Mitra, G. (2009). Portfolio selection models: a review and new directions. *Wilmott journal: the international journal of innovative quantitative finance research, 1*(2), 69-85.

Sharpe, W. (1964). Capital Asset Prices: A Theory of Market Equilibrium Under Conditions of Risk. *The Journal of Finance*, *19*, 425-442.

Schannep, J. (2012). *Dow Theory for the 21st Century*. Wiley.

Silva, L., Alem, D., & Carvalho, F. (2017). Portfolio optimisation using Mean Absolute Deviation (MAD) and Conditional Value-at-Risk (CVaR). *Production*, *27,* 1-14.

Soiferman, L. K. (2010). Compare and Contrast Inductive and Deductive Research Approaches. Retrieved from https://files.eric.ed.gov/fulltext/ED542066.pdf

Sortino, F. A. (1994). Performance Measurement in a Downside Risk Framework. *The Journal of Investing*, *3*, 59-64.

Standard & Poor's. (2022). *S&P Global 1200*. Retrieved from https://www.spglobal.com/spdji/en/index-family/equity/global-equity/sp-global-1200/#indices.

Stanković, J. Z., Petrović, E., & Denčić-Mihajlov, K. (2020). Effects of Applying Different Risk Measures on the Optimal Portfolio Selection: The Case of the Belgrade Stock Exchange. *Facta Universitatis, Series: Economics and Organization*. *17*(1)*,* 17-26.

Suganuma, R. (2000). Reality Check for Volatility Models. Retrieved from https://www.econ.puc-rio.br/uploads/adm/trabalhos/files/suganuma.pdf

Tavakoli Baghdadabad, M. R. (2015). An empirical analysis of funds' alternative measures in the mean absolute deviation (MAD) framework. *International Journal of Emerging Markets, 10*(4), 726-746.

Vercher, E., & Bermúdez, J. D. (2015). Portfolio optimization using a credibility mean-absolute semi-deviation model. *Expert Systems with Applications, 42*(20), 7121-7131.

Wilson, L. (2014). Modern Portfolio Theory with VaR Objective Functions. *SSRN Electronic Journal*. Retrieved from https://dx.doi.org/10.2139/ssrn.2516022

APPENDIX A: GICS CLASSICFICATIONS

Table 6.1

GICS Sector, Industry Group and Industry classifications

Sector	Industry Group	Industry
Energy	Energy	Energy Equipment & Services
		Oil, Gas & Consumable Fuels
Materials	Materials	Chemicals
		Construction Materials
		Containers & Packaging
		Metals & Mining
		Paper & Forest Products
Industrials	Capital Goods	Aerospace & Defense
		Building Products
		Construction & Engineering
		Electrical Equipment
		Industrial Conglomerates
		Machinery
		Trading Companies & Distributors
	Commercial & Professional Services	Commercial Services & Supplies
		Professional Services
	Transportation	Air Freight & Logistics
		Airlines
		Marine
		Road & Rail
		Transportation Infrastructure
Consumer Discretionary (Consumer Cyclical)	Automobiles & Components	Auto Components
		Automobiles
	Consumer Durables & Apparel	Household Durables
		Leisure Products
		Textiles, Apparel & Luxury Goods
	Consumer Services	Hotels, Restaurants & Leisure
		Diversified Consumer Services
	Retailing	Distributors
		Internet & Direct Marketing Retail
		Multiline Retail
		Specialty Retail
Consumer Staples (Consumer Defensive)	Food & Staples Retailing	Food & Staples Retailing
	Food, Beverage & Tobacco	Beverages
		Food Products
		Tobacco
	Household & Personal Products	Household Products
		Personal Products
Health Care	Health Care Equipment & Services	Health Care Equipment & Supplies
		Health Care Providers & Services
		Health Care Technology
	Pharmaceuticals, Biotechnology & Life Sciences	Biotechnology
		Pharmaceuticals
		Life Sciences Tools & Services
Financials	Banks	Banks
		Thrifts & Mortgage Finance
	Diversified Financials	Diversified Financial Services
		Consumer Finance
		Capital Markets
		Mortgage Real Estate Investment Trusts (REITs)
	Insurance	Insurance
Information Technology	Software & Services	IT Services
		Software
	Technology Hardware & Equipment	Communications Equipment
		Technology Hardware, Storage & Peripherals
		Electronic Equipment, Instruments & Components
	Semiconductors & Semiconductor Equipment	Semiconductors & Semiconductor Equipment
Communication Services	Telecommunication Services	Diversified Telecommunication Services
		Wireless Telecommunication Services
	Media & Entertainment	Media
		Entertainment
		Interactive Media & Services
Utilities	Utilities	Electric Utilities
		Gas Utilities
		Multi-Utilities
		Water Utilities
		Independent Power and Renewable Electricity Producers
Real Estate	Real Estate	Equity Real Estate Investment Trusts (REITs)
		Real Estate Management & Development

APPENDIX B: POINT-IN-TIME PORTFOLIO CONSTRUCTION

The purpose of this appendix is to demonstrate the efficacy of the numerical methods algorithm used to find the efficient portfolio when optimising risk and return for each risk measure. If the portfolios returned are not optimal, or close to optimal, then the conclusions drawn for the resulting strategies will not be correct.

Figure 7.1

Examples of point-in-time optimised vs random portfolio

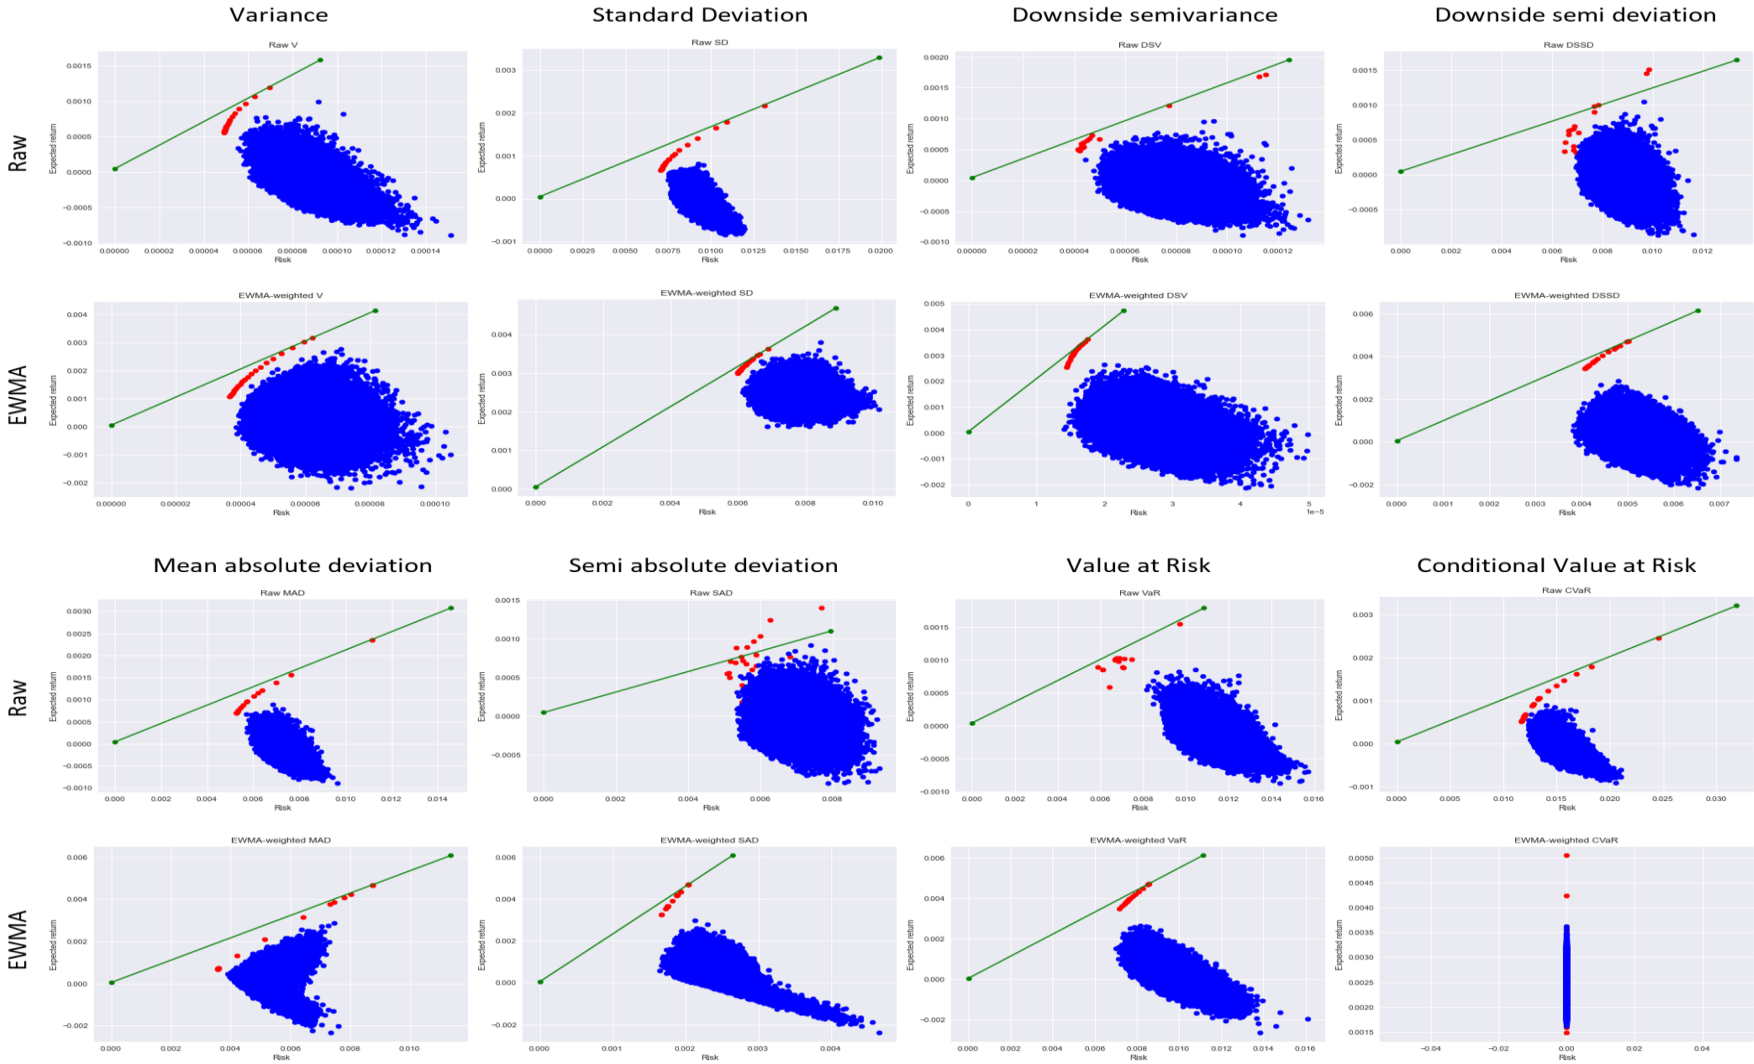

Note. Risk vs return plots for random (blue), and optimised (red) portfolios optimised using each risk measure investigated in the study for a single period. The Capital Market Line (green) used to optimise portfolios is plotted between the risk and return of the risk-free asset and an optimised portfolio given the return of the particular risk-free asset.

Source: Author’s calculations

Figure 7.1 plots the risk (x-axis) and mean return (y-axis) for point-in-time portfolios, with each graph showing portfolios for each different raw and EWMA-weighted risk measure. Each dot in the point-in-time portfolio figures show the mean return and variance for a portfolio estimated over 100 days of data. Using a Monte Carlo simulation approach, the blue dots represent 50,000 portfolios with randomly assigned constituent weights, subject to the constraint that the weights sum to one and no short selling (random portfolios), while the red dots represent portfolios on the efficient frontier (optimised portfolios). The optimised portfolios are calculated by choosing different risk-free rates (that is, points on the y-axis) and maximising the slope of the CML for each using the numerical methods procedure described above. The green line represents the CML, joining the actual risk-free rate at the end of the 100-day period and the efficient frontier.

For portfolios optimised using the majority of the raw risk measures and all of the EWMA-weighted risk measures, the gap between the 50,000 random portfolios and the efficient frontier indicates the efficacy of the optimisation procedure: the random portfolios are not as efficient as the optimised portfolios, even across 50,000 attempts.

In contrast, portfolios optimised using raw downside semivariance, downside semi deviation, semi-absolute deviation and value at risk are not perfectly tractable risk measures for portfolio optimisation using the numerical methods employed in the study. This issue is noted by Gaivoronski (2005) and Lwin et al. (2017) when discussing the value at risk-optimised portfolios. They note that the efficient frontier between expected returns and VaR is not smooth and can result in local minima, making mean-VaR portfolio optimisation computationally intractable.

Semi-absolute deviation provides a further example, since the risk measure only includes an observation into its sum when the observation is negative. When optimising the weights of a portfolio, there are certain weights which, when nudged upwards or downwards, result in an observation moving from being included into the risk measure to being excluded, or *visa versa*. However, the numerical methods algorithm relies on the principle that successive small changes in each portfolio weight by the same amount should result in a similar change to the risk measure.

Figure 7.2

Illustration of numerical methods algorithm premise

	Asset A	Asset B	Risk Measure
Weighting 1	50%	50%	100
Weighting 2	49%	51%	99
Change	-1%	1%	-1

implies that

	Asset A	Asset B	Risk Measure
Weighting 2	49%	51%	99
Weighting 3	48%	52%	~98
Change	-1%	1%	~-1

Note. Example used to demonstrate how the numerical methods algorithm used to optimise each portfolio uses gradient descent to optimise.

By way of example, Figure 7.2 illustrates the premise used in the numerical methods algorithm that in a two asset universe, if a move in the portfolio weights from 50% / 50% to 49% / 51% results in a change in the risk measure from 100 to 99, then a move from 49% / 51% to 48% / 52% should result in a change in the risk measure from 99 to approximately 98. However, when the weights being optimised are near these thresholds of including or excluding return observations from the risk measure calculation, the move from inclusion to exclusion or *visa versa* invalidates this premise, thereby causing the numerical methods algorithm to over- or underestimate the impact of successive weighting changes.

Because downside semivariance, downside semi deviation, semi-absolute deviation all exhibit the property of including or excluding return observations into the calculation of the risk measure dependent on the return of the portfolio in the period, and therefore upon the weights of the portfolio's constituent assets over the period, none of these risk measures are perfectly tractable to optimise over.

To address this, for these risk measures, the algorithm selects the portfolio weights resulting in the best risk/return ratio from a thousand randomised portfolio weightings as a starting point for the gradient descent, and repeats the full optimisation more than once, from which it chooses the best results.

The graph for the portfolios optimised using a risk measure most affected by this characteristic, semi-absolute deviation, yields the following observations: firstly, because the algorithm selects the best of many random starting points, the final level of optimisation may differ, with the result that one optimised portfolio may be more optimal than another. This results in the CML in the semi-absolute deviation graph clearly not being tangential, since the portfolio found by optimisation to find the CML is less optimal than others found starting from other random starting points. Secondly, while the resulting efficient frontier for these risk measures are not as smooth as that of, for example, variance-optimised portfolios, the optimised portfolios are nonetheless better optimised than the 50,000 random portfolios generated in terms of the historical ratio of risk and return. Consequently, despite the portfolios optimised using these risk measures not necessarily being the *most* efficient portfolios possible, they are nonetheless considerably more efficient than even the best of the random portfolios, and the resulting strategies should be suitable for drawing inference about portfolios optimised using these risk measures. Thirdly, it is clear that application of an EWMA to these risk measures renders them tractable to optimisation using numerical methods.

APPENDIX C: HYPERPARAMETER OPTIMISATION

The study required the optimisation of different hyperparameters before using them in the final investment strategies. In order to avoid forward-looking bias, the hyperparameters may only be optimised using past data, or alternatively, where this is not feasible, a parameter was chosen and the entire sample was used to corroborate the choice of parameter.

It is preferable to select hyperparameters for which adjacent parameter values have a positive but similar effect on the outcome of a strategy in preference to those which result in better strategy performance, but poor results for adjacent parameter values. Doing so means that a small shift in the investment environment with respect to the effect of the parameter is less likely to result in significant changes to the strategy performance.

The hyperparameters requiring calculation are recalculation period, length of training period and GICS tier. The first of these addresses is recalculation period.

8.1 RECALCULATION PERIOD

Comparing the returns and Sharpe ratios of daily, weekly and monthly strategies optimised using each raw risk measure, it is clear that a daily recalculation period is not helpful in optimising either annualised return or risk-adjusted returns. It is not entirely clear why this should be the case. Part of the reduced performance is attributable to higher trading costs due to more frequent trading, but this increase in costs was not sufficient to explain the full difference. As noted in Section 3.5.1 (Sample periods and rebalancing), the risk measures for weekly and monthly periods was calculated using daily data and scaled to fit their respective periods. It may be that such scaling understates the risk measures, resulting in a greater emphasis upon returns in the optimisation of return / risk for longer recalculation periods and concomitant higher annualised returns, albeit at the cost of higher risk. Annualised returns across weekly and monthly recalculation periods performed similarly, with five of the ten risk measure-optimised strategies recalculated weekly having the highest annualised return and four of the ten risk measure optimised-strategies recalculated monthly having the highest annualised return. On average, the annualised return for monthly strategies was slightly higher than that of weekly strategies. Moving to Sharpe ratios, the monthly recalculation period achieved the highest Sharpe ratio across all ten risk measure-optimised strategies. These results are summarised in table 8.1.

Table 8.1

Performance of raw-risk optimised strategies across different recalculation periods

	Annualised Return			Sharpe Ratio		
	Daily	Weekly	Monthly	Daily	Weekly	Monthly
V	10.1%	11.2%	11.6%	65%	71%	91%
SD	11.5%	12.3%	11.4%	70%	73%	79%
DSV	11.5%	10.4%	10.9%	68%	62%	82%
DSSD	11.5%	14.1%	13.2%	69%	95%	107%
MAD	11.5%	12.7%	12.9%	69%	76%	92%
SAD	12.9%	13.9%	15.2%	82%	85%	111%
VaR0.1	10.0%	10.3%	9.7%	58%	60%	70%
Var0.05	9.6%	8.9%	11.5%	54%	49%	82%
CVaR0.1	10.6%	12.2%	11.4%	62%	73%	78%
CVar0.05	11.2%	12.5%	12.2%	65%	73%	85%
Average	**11.0%**	**11.8%**	**12.0%**	**66%**	**72%**	**88%**

Note. Annualised returns and Sharpe ratios for strategies calculated using each raw risk measure compared across daily, weekly and monthly recalculation periods.
Source: Author's calculations

Undertaking a similar comparison for daily, weekly and monthly strategies optimised using the EWMA of each risk measure, daily recalculation periods are again excluded due to poor performance. Weekly recalculation periods overwhelmingly perform best in terms of annualised returns, with seven of the ten risk measure-optimised strategies recalculated weekly having the highest annualised return, as well as weekly strategies having the highest average return by a significant margin. However, six of the ten EWMA risk measure optimised-strategies recalculated monthly have the highest Sharpe ratio, and monthly strategies also have the highest average Sharpe ratio. These results are summarised in table 8.2.

Table 8.2

Performance of EWMA-risk optimised strategies across different recalculation periods

	Annualised Return			Sharpe Ratio		
	Daily	Weekly	Monthly	Daily	Weekly	Monthly
V	11.6%	11.2%	8.1%	77%	71%	59%
SD	9.4%	12.9%	11.8%	55%	77%	86%
DSV	7.9%	13.1%	8.3%	45%	81%	56%
DSSD	8.5%	12.1%	11.1%	48%	71%	76%
MAD	9.6%	16.2%	11.8%	57%	103%	87%
SAD	9.4%	15.8%	11.9%	56%	100%	89%
VaR0.1	9.6%	13.1%	11.3%	55%	79%	82%
Var0.05	9.7%	13.1%	11.6%	57%	79%	85%
CVaR0.1	9.8%	6.2%	9.3%	58%	30%	71%
CVar0.05	10.3%	8.5%	12.4%	62%	44%	84%
Average	**9.6%**	**12.2%**	**10.8%**	**57%**	**74%**	**77%**

Note. Annualised returns and Sharpe ratios for strategies calculated using each EWMA risk measure compared across daily, weekly and monthly recalculation periods.
Source: Author's calculations

In conclusion, a monthly rebalancing period performed best at optimising returns and risk-adjusted returns across the most risk-measure optimised strategies, even though weekly EWMA-optimised strategies performed best in terms of annualised returns. It is clear that a daily balancing period is generally suboptimal in maximising returns or risk-adjusted returns. Consequently, daily strategies were excluded from the analysis.

8.2 TRAINING PERIOD

The next hyperparameter examined is the length of the training period.

Table 8.3 provides the annualised returns and Sharpe ratios respectively, averaged across the ten risk measure-optimised strategies, for raw weekly-optimised strategies, raw monthly-optimised strategies, EWMA weekly-optimised strategies and EWMA monthly-optimised strategies. Each row represents a different period of training data, starting at one month and ending at eighteen months. The 'Average of averages' column is an average of the annualised returns and Sharpe ratios in the applicable row.

The training periods between three and nine months yield the best annualised returns and Sharpe ratios, with the exception of a seven-month training period.

Table 8.3

Average annualised returns and Sharpe ratios split by training period

	Average Annualised Return				Average Sharpe Ratio				Average of averages
	Raw		EWMA		Raw		EWMA		
Months training data	Weekly	Monthly	Weekly	Monthly	Weekly	Monthly	Weekly	Monthly	
1	11.2%	9.0%	11.5%	10.8%	0.71	0.68	0.67	0.72	0.40
2	9.3%	10.6%	6.7%	10.3%	0.58	0.80	0.41	0.67	0.35
3	11.0%	11.2%	11.7%	10.3%	0.73	0.90	0.78	0.81	0.46
4	14.2%	12.2%	10.3%	6.7%	1.04	0.98	0.75	0.46	0.46
5	11.3%	12.5%	13.8%	10.6%	0.88	1.05	0.99	0.71	0.51
6	12.2%	11.4%	9.5%	10.3%	0.82	0.97	0.57	0.76	0.44
7	10.8%	11.6%	9.2%	7.8%	0.68	0.93	0.63	0.54	0.40
8	12.3%	12.3%	11.8%	8.7%	0.81	1.04	0.75	0.63	0.46
9	12.8%	13.1%	10.8%	7.8%	0.93	1.11	0.81	0.59	0.49
10	12.5%	12.2%	8.4%	7.2%	0.96	1.01	0.57	0.52	0.43
11	11.1%	11.5%	7.1%	6.5%	0.69	0.95	0.43	0.47	0.36
12	10.4%	10.0%	12.8%	7.9%	0.63	0.82	0.85	0.61	0.42
13	10.6%	10.2%	6.9%	5.7%	0.66	0.82	0.43	0.40	0.33
14	11.4%	10.5%	8.0%	4.9%	0.81	0.84	0.56	0.33	0.36
15	11.2%	10.8%	8.5%	5.8%	0.82	0.89	0.58	0.42	0.38
16	10.8%	11.0%	11.3%	5.1%	0.66	0.89	0.68	0.34	0.37
17	11.3%	11.1%	7.8%	7.1%	0.66	0.91	0.46	0.51	0.36
18	11.4%	11.0%	9.9%	6.9%	0.72	0.88	0.64	0.52	0.39

Note. Average annualised returns and Sharpe ratios for strategies calculated using each risk measure, EWMA vs raw risk measure and weekly and monthly recalculation periods using different lengths of training data, ranging from one month to eighteen months. Higher returns and higher Sharpe ratios indicate more effective strategies, coloured in green.
Source: Author's calculations

While the study made use of 3 months of training data for all strategies presented, this period was not selected based on the efficacy of the training period, since doing so would have introduced in-sample data into the decision-making process. However, the risk-adjusted returns for strategies optimised using a 3-month training period are reasonable for monthly and weekly strategies, as compared to both longer and shorter training periods.

8.3 INDUSTRY TIER

The third hyperparameter examined in the study is the Industry tier. As discussed in Section 3.4 (Sampling and data collection), the GICS has three tiers of categorisation, from the most aggregated Sector tier to the least aggregated Industry tier. While comparison between the strategies using the Industry Group and Industry indices as inputs is applicable, comparison to strategies with Sector inputs is not directly analogous, since the S&P Sector indices are total return indices, while the MSIC Industry Group and Industry indices are price indices. However, since only the Sector-level indices have investible ETFs available, the purpose of analysing the impact of the data tier on strategy performance is to provide reassurance that there is no material loss of performance associated with using the already-highly-aggregated indices. Examining the results laid out in table 8.5, this reassurance is indeed provided.

Table 8.4 shows the average annualised return and average Sharpe ratio across all ten risk measure-optimised strategies, for each GICS tier. The results are split by recalculation period and Raw/EWMA categorisation, and indicate strongest performance from the mid-tier Industry Group indices. This is closely followed by the performance of Sector-level indices, with strategies based on Industry-level indices performing poorly in comparison to those based on the two more aggregated sets of indices.

Table 8.4

Risk-adjusted returns for each GICS tier's monthly raw risk-optimised strategies

		Average Annualised Return				Average Sharpe Ratio			
		Weekly		Monthly		Weekly		Monthly	
		Raw	EWMA	Raw	EWMA	Raw	EWMA	Raw	EWMA
Tier	Sector	11.3%	11.7%	11.7%	10.3%	0.72	0.74	0.88	0.77
	Industry Group	12.9%	17.1%	13.5%	10.7%	0.73	0.98	0.85	0.63
	Industry	11.4%	7.7%	10.3%	6.4%	0.59	0.42	0.64	0.39

Note. Comparison of industry tier, split into sector, industry group and industry, showing the relative average performance of strategies grouped by recalculation period and raw vs EWMA risk measure. Higher returns and higher Sharpe ratios indicate more effective strategies, coloured in green.

Source: Author's calculations

The results lend themselves to the following conclusions: Firstly, diversifying across a less aggregated set of indices does not ensure that risk-return-optimised strategies perform better. Secondly, the use of Sector indices as inputs does not result in materially poorer absolute or risk-adjusted performance than Industry Group-level indices, particularly for a monthly recalculation period, although there is a material outlier in the form of very strong performance of weekly recalculated EWMA-optimised strategies.

Accordingly, the use of Sector-level indices as input data into the study was corroborated as valid, although, were investible ETFs available for slightly less aggregated, Industry Group-level indices available, they would have provided a better-performing investment universe.

8.4 EWMA PARAMETERISATION

The fourth hyperparameter investigated relates to the choice of the EWMA's rate of decay. This is required before attempting to address the criticism of MPT that risk and return do not persist into the future through estimation of the next period's risk and return using an EWMA.

Two different approaches to determining lambda were compared. The first approach, described in Section 3.5.4 (Calculating the EWMA), finds the value of lambda which minimises the absolute value of the difference between an estimate for the true risk measure and the EWMA, while for the second approach, lambda is a constant value of 0.97. The results of this comparison are displayed in table 8.5, which, for each approach, shows the average annualised return and average Sharpe ratio across the ten risk measure-optimised strategies for the weekly and monthly recalculation periods. The results are not conclusive, with the constant-lambda approach resulting in better performance for strategies recalculated monthly by a small margin, and the variable-lambda approach performing better for weekly strategies by a larger margin.

Table 8.5

Comparison of variable λ vs constant λ for EWMA-optimised strategies

	Average Annualised Return			Average Sharpe Ratio		
	Weekly	Monthly	Average	Weekly	Monthly	Average
Constant lambda	13.1%	11.5%	12.3%	0.78	0.68	0.34
Variable lambda	17.1%	10.7%	13.9%	0.98	0.63	0.41

Note. Comparison between EWMA-optimised strategies where the EWMA decay factor is a constant vs where it is re-optimised each period. Higher returns and higher Sharpe ratios indicate more effective strategies, coloured in green.

Source: Author's calculations

Accordingly, the variable-lambda approach was used when calculating the EWMAs used in the study.

APPENDIX D: ADDITIONAL STRATEGY RESULTS

The results of strategies not discussed in detail are included in this appendix for reference purposes.

Table 9.1

Sector-based strategy annualised returns and annualised risk

Strategy							Strategy risk										
				Mean	Annualised Return	Fees	Max Drawdown	V	SD	DSV	DSSD	MAD	SAD	VaR0.1	VaR0.05	CVaR0.1	CVaR0.05
GlobalIndex				12.08%	11.56%		29.9%	0.0217	0.1712	0.1444	0.1922	0.2534	0.3365	0.1474	0.0293	0.0972	11.1%
Sector	Daily	Raw	V	10.13%	9.61%	0.103%	30.5%	0.0190	0.1464	0.1224	0.1895	0.2394	0.3283	0.1377	0.0214	0.0856	8.8%
			SD	11.51%	10.98%	0.103%	30.7%	0.0217	0.1568	0.1336	0.2141	0.2610	0.3546	0.1474	0.0246	0.0939	9.7%
			DSV	11.46%	10.88%	0.104%	33.2%	0.0225	0.1552	0.1284	0.2008	0.2541	0.3486	0.1501	0.0241	0.0914	9.3%
			DSSD	11.46%	10.91%	0.101%	33.0%	0.0218	0.1568	0.1360	0.2118	0.2610	0.3518	0.1475	0.0246	0.0942	9.8%
			MAD	11.49%	10.94%	0.103%	30.9%	0.0219	0.1573	0.1338	0.2161	0.2628	0.3566	0.1479	0.0247	0.0942	9.8%
			SAD	12.93%	12.60%	0.102%	33.6%	0.0211	0.1529	0.1340	0.2053	0.2545	0.3417	0.1452	0.0234	0.0927	9.5%
			VaR0.1	10.01%	9.31%	0.104%	35.9%	0.0220	0.1612	0.1328	0.2176	0.2647	0.3645	0.1484	0.0260	0.0927	9.8%
			VaR0.05	9.62%	8.84%	0.103%	35.2%	0.0227	0.1601	0.1346	0.2147	0.2665	0.3638	0.1507	0.0256	0.0943	9.8%
			CVaR0.1	10.59%	9.92%	0.103%	37.6%	0.0225	0.1610	0.1347	0.2169	0.2665	0.3651	0.1500	0.0259	0.0948	9.8%
			CVaR0.05	11.15%	10.52%	0.103%	34.9%	0.0228	0.1606	0.1366	0.2138	0.2660	0.3633	0.1509	0.0258	0.0957	9.9%
		EWMA	V	11.63%	11.26%	0.116%	30.6%	0.0190	0.1476	0.1253	0.1903	0.2444	0.3325	0.1378	0.0218	0.0881	9.2%
			SD	9.45%	8.71%	0.120%	36.2%	0.0216	0.1633	0.1395	0.2210	0.2665	0.3584	0.1471	0.0267	0.0964	10.2%
			DSV	7.90%	7.10%	0.120%	31.9%	0.0207	0.1531	0.1377	0.2053	0.2582	0.3473	0.1440	0.0234	0.0934	9.7%
			DSSD	8.54%	7.72%	0.120%	34.3%	0.0219	0.1621	0.1422	0.2220	0.2680	0.3606	0.1479	0.0263	0.0975	10.3%
			MAD	9.63%	8.95%	0.120%	35.7%	0.0210	0.1590	0.1394	0.2125	0.2621	0.3520	0.1448	0.0253	0.0956	10.1%
			SAD	9.42%	8.72%	0.121%	36.1%	0.0209	0.1600	0.1394	0.2151	0.2630	0.3529	0.1447	0.0256	0.0954	10.1%
			VaR0.1	9.56%	8.82%	0.120%	36.3%	0.0218	0.1638	0.1394	0.2184	0.2664	0.3591	0.1476	0.0268	0.0966	10.2%
			VaR0.05	9.73%	9.01%	0.120%	35.8%	0.0217	0.1621	0.1394	0.2196	0.2656	0.3570	0.1472	0.0263	0.0965	10.1%
			CVaR0.1	9.79%	9.12%	0.117%	29.3%	0.0211	0.1576	0.1277	0.2144	0.2663	0.3689	0.1451	0.0248	0.0924	9.6%
			CVaR0.05	10.27%	9.64%	0.117%	30.7%	0.0212	0.1546	0.1307	0.2136	0.2652	0.3633	0.1455	0.0239	0.0933	9.7%
	Weekly	Raw	V	11.19%	10.66%	0.101%	27.0%	0.0198	0.1641	0.1243	0.1913	0.2381	0.3235	0.1406	0.0269	0.0887	9.9%
			SD	12.30%	11.72%	0.101%	27.4%	0.0229	0.1762	0.1404	0.2164	0.2613	0.3580	0.1512	0.0310	0.0971	10.8%
			DSV	10.43%	9.73%	0.101%	26.9%	0.0215	0.1716	0.1425	0.2151	0.2570	0.3394	0.1466	0.0294	0.0934	10.5%
			DSSD	14.08%	13.94%	0.100%	23.5%	0.0194	0.1586	0.1422	0.1984	0.2418	0.3159	0.1391	0.0252	0.0947	10.5%
			MAD	12.73%	12.20%	0.101%	27.4%	0.0229	0.1734	0.1410	0.2119	0.2603	0.3571	0.1514	0.0301	0.0971	10.5%
			SAD	13.88%	13.50%	0.100%	35.5%	0.0227	0.1758	0.1337	0.1959	0.2518	0.3446	0.1506	0.0309	0.0955	10.6%
			VaR0.1	10.30%	9.57%	0.101%	27.4%	0.0219	0.1696	0.1454	0.2173	0.2563	0.3420	0.1481	0.0288	0.0948	10.2%
			VaR0.05	8.93%	8.06%	0.101%	27.7%	0.0224	0.1737	0.1461	0.2124	0.2621	0.3522	0.1495	0.0302	0.0975	10.8%
			CVaR0.1	12.18%	11.61%	0.101%	26.2%	0.0226	0.1722	0.1354	0.2180	0.2577	0.3523	0.1502	0.0297	0.0971	10.5%
			CVaR0.05	12.46%	11.87%	0.101%	26.5%	0.0234	0.1769	0.1349	0.2238	0.2652	0.3584	0.1530	0.0313	0.0993	11.0%
		EWMA	V	11.24%	10.71%	0.104%	26.7%	0.0198	0.1614	0.1367	0.1771	0.2288	0.3067	0.1407	0.0260	0.0912	9.8%
			SD	12.85%	12.34%	0.104%	26.8%	0.0229	0.1674	0.1443	0.2015	0.2475	0.3284	0.1513	0.0280	0.1000	10.4%
			DSV	13.12%	12.70%	0.104%	26.5%	0.0219	0.1643	0.1439	0.1861	0.2401	0.3197	0.1478	0.0270	0.0966	10.1%
			DSSD	12.06%	11.45%	0.104%	25.8%	0.0229	0.1708	0.1383	0.2107	0.2573	0.3428	0.1513	0.0292	0.1009	10.7%
			MAD	16.22%	16.18%	0.104%	27.3%	0.0227	0.1655	0.1369	0.1969	0.2389	0.3204	0.1507	0.0274	0.1004	10.3%
			SAD	15.83%	15.74%	0.104%	27.3%	0.0225	0.1658	0.1384	0.1968	0.2395	0.3215	0.1501	0.0275	0.0989	10.2%
			VaR0.1	13.10%	12.63%	0.104%	26.8%	0.0228	0.1670	0.1442	0.2015	0.2473	0.3278	0.1510	0.0279	0.0997	10.4%
			VaR0.05	13.08%	12.60%	0.104%	26.8%	0.0228	0.1667	0.1422	0.2015	0.2474	0.3286	0.1510	0.0278	0.0998	10.3%
			CVaR0.1	6.17%	5.11%	0.104%	27.5%	0.0226	0.1799	0.1435	0.1952	0.2665	0.3670	0.1503	0.0323	0.0979	11.1%
			CVaR0.05	8.47%	7.49%	0.104%	26.7%	0.0237	0.1808	0.1539	0.1953	0.2658	0.3573	0.1541	0.0327	0.1033	11.7%
	Monthly	Raw	V	11.58%	11.34%	0.100%	21.7%	0.0136	0.1386	0.1023	0.1370	0.1847	0.2567	0.1166	0.0192	0.0831	9.5%
			SD	11.44%	10.99%	0.100%	18.7%	0.0172	0.1556	0.1023	0.1684	0.2111	0.2918	0.1311	0.0242	0.0929	11.1%
			DSV	10.85%	10.50%	0.100%	19.1%	0.0145	0.1431	0.1036	0.1722	0.1998	0.2784	0.1205	0.0205	0.0851	9.7%
			DSSD	13.15%	13.07%	0.099%	18.7%	0.0134	0.1409	0.0678	0.1413	0.1830	0.2655	0.1157	0.0199	0.0794	9.1%
			MAD	12.90%	12.61%	0.100%	18.0%	0.0170	0.1513	0.1106	0.1538	0.2046	0.2792	0.1303	0.0229	0.0922	10.5%
			SAD	15.21%	15.18%	0.099%	18.0%	0.0171	0.1513	0.1046	0.1833	0.2024	0.2764	0.1307	0.0229	0.0946	10.6%
			VaR0.1	9.74%	9.27%	0.100%	19.8%	0.0150	0.1446	0.1089	0.1649	0.2025	0.2780	0.1226	0.0209	0.0879	10.0%
			VaR0.05	11.54%	11.14%	0.100%	22.3%	0.0163	0.1501	0.0997	0.1565	0.1980	0.2827	0.1278	0.0225	0.0926	10.6%
			CVaR0.1	11.36%	10.90%	0.100%	19.3%	0.0172	0.1469	0.0997	0.1711	0.2069	0.2924	0.1312	0.0216	0.0939	10.1%
			CVaR0.05	12.18%	11.81%	0.100%	18.9%	0.0171	0.1451	0.0997	0.1424	0.2001	0.2848	0.1307	0.0210	0.0920	9.5%
		EWMA	V	8.14%	7.61%	0.101%	18.4%	0.0140	0.1415	0.1022	0.1686	0.2048	0.2818	0.1185	0.0200	0.0815	9.9%
			SD	11.76%	11.43%	0.100%	19.4%	0.0156	0.1406	0.1032	0.1378	0.1980	0.2791	0.1248	0.0198	0.0881	9.5%
			DSV	8.30%	7.70%	0.101%	22.5%	0.0157	0.1443	0.1133	0.1453	0.2177	0.3124	0.1254	0.0208	0.0864	9.5%
			DSSD	11.13%	10.64%	0.100%	20.3%	0.0173	0.1460	0.1023	0.1329	0.2019	0.2938	0.1316	0.0213	0.0924	9.6%
			MAD	11.82%	11.50%	0.100%	18.9%	0.0153	0.1447	0.0956	0.1442	0.1940	0.2757	0.1238	0.0209	0.0867	9.9%
			SAD	11.91%	11.62%	0.100%	17.8%	0.0152	0.1461	0.0997	0.1340	0.1918	0.2735	0.1231	0.0214	0.0852	9.9%
			VaR0.1	11.28%	10.90%	0.100%	19.4%	0.0157	0.1434	0.1106	0.1416	0.2009	0.2813	0.1255	0.0206	0.0891	9.8%
			VaR0.05	11.60%	11.26%	0.100%	19.2%	0.0156	0.1411	0.1023	0.1391	0.1984	0.2813	0.1251	0.0199	0.0881	9.5%
			CVaR0.1	9.30%	8.89%	0.100%	15.4%	0.0134	0.1371	0.0982	0.1574	0.1993	0.2681	0.1159	0.0188	0.0806	9.4%
			CVaR0.05	12.35%	11.94%	0.100%	21.3%	0.0182	0.1498	0.1228	0.1743	0.2269	0.3165	0.1349	0.0224	0.0938	10.1%

Note. Annualised returns for each strategy based on sector-level indices. V = variance, SD = standard deviation, DSV = downside semivariance, DSSD = downside semideviation, MAD = mean absolute deviation, SAD = semi-absolute deviation, VaR0.1 = 10% value at risk, VaR0.05 = 5% value at risk, CVaR0.1 = 10% conditional value at risk, CVaR0.05 = 5% conditional value at risk.

Source: Author’s calculations

Table 9.2

Industry Group-based strategy annualised returns and annualised risk

			Strategy	Mean	Annualised Return	Fees	Strategy risk										
							Max Drawdown	V	SD	DSV	DSSD	MAD	SAD	VaR0.1	VaR0.05	CVaR0.1	CVaR0.05
	GlobalIndex			9.98%	9.10%		34.1%	0.0248	0.1773	0.1422	0.2370	0.2892	0.3966	0.1574	0.0314	0.1015	11.6%
Industry Group	Weekly	Raw	V	10.14%	9.23%	0.123%	34.9%	0.0262	0.1769	0.1475	0.2078	0.2617	0.3550	0.1618	0.0313	0.0988	11.4%
			SD	14.98%	14.48%	0.123%	35.1%	0.0288	0.1808	0.1514	0.2275	0.2758	0.3676	0.1697	0.0327	0.1058	11.6%
			DSV	10.86%	9.98%	0.125%	34.6%	0.0267	0.1757	0.1400	0.1986	0.2661	0.3673	0.1634	0.0309	0.0975	10.8%
			DSSD	14.39%	13.71%	0.122%	34.7%	0.0304	0.1867	0.1626	0.2269	0.2925	0.3873	0.1743	0.0348	0.1091	11.9%
			MAD	13.25%	12.64%	0.124%	35.4%	0.0263	0.1774	0.1484	0.2180	0.2759	0.3680	0.1622	0.0315	0.1048	11.3%
			SAD	12.20%	11.60%	0.121%	32.5%	0.0242	0.1812	0.1756	0.2307	0.2867	0.3724	0.1557	0.0328	0.1078	12.5%
			VaR0.1	11.81%	10.95%	0.129%	36.3%	0.0279	0.1826	0.1481	0.2286	0.2884	0.3919	0.1669	0.0333	0.1083	11.6%
			VaR0.05	15.02%	14.57%	0.125%	34.8%	0.0280	0.1814	0.1507	0.2328	0.2797	0.3669	0.1674	0.0329	0.1058	11.6%
			CVaR0.1	15.81%	15.42%	0.123%	34.6%	0.0290	0.1797	0.1467	0.2252	0.2777	0.3661	0.1703	0.0323	0.1062	11.5%
			CVaR0.05	16.56%	16.27%	0.127%	34.7%	0.0295	0.1840	0.1579	0.2185	0.2783	0.3684	0.1718	0.0339	0.1068	12.0%
		EWMA	V	11.14%	10.53%	0.159%	25.6%	0.0224	0.1606	0.1493	0.2248	0.2659	0.3523	0.1496	0.0258	0.0977	10.5%
			SD	18.74%	18.88%	0.150%	26.4%	0.0283	0.1760	0.1436	0.2303	0.2756	0.3694	0.1683	0.0310	0.1092	11.4%
			DSV	15.21%	14.93%	0.160%	25.7%	0.0256	0.1747	0.1534	0.2203	0.2763	0.3745	0.1599	0.0305	0.1066	11.7%
			DSSD	18.87%	19.18%	0.160%	25.5%	0.0255	0.1670	0.1369	0.2009	0.2562	0.3464	0.1597	0.0279	0.1045	10.7%
			MAD	19.69%	20.01%	0.156%	26.1%	0.0283	0.1712	0.1436	0.2145	0.2624	0.3534	0.1682	0.0293	0.1077	11.0%
			SAD	19.25%	19.48%	0.160%	26.1%	0.0283	0.1754	0.1424	0.2140	0.2641	0.3552	0.1682	0.0308	0.1078	11.4%
			VaR0.1	18.60%	18.69%	0.152%	26.6%	0.0286	0.1775	0.1458	0.2377	0.2764	0.3702	0.1691	0.0315	0.1099	11.7%
			VaR0.05	18.61%	18.71%	0.151%	26.4%	0.0284	0.1769	0.1456	0.2303	0.2760	0.3699	0.1684	0.0313	0.1095	11.6%
			CVaR0.1	15.58%	14.90%	0.182%	21.8%	0.0336	0.1789	0.1416	0.2301	0.2962	0.4142	0.1834	0.0320	0.1166	11.4%
			CVaR0.05	16.20%	15.76%	0.171%	20.8%	0.0310	0.1716	0.1392	0.2081	0.2760	0.3904	0.1761	0.0294	0.1092	10.8%
	Monthly	Raw	V	12.23%	11.78%	0.111%	26.2%	0.0196	0.1599	0.0854	0.1178	0.1989	0.2990	0.1402	0.0256	0.0894	9.3%
			SD	14.92%	14.51%	0.111%	26.9%	0.0246	0.1703	0.0908	0.1305	0.2210	0.3332	0.1570	0.0290	0.1014	10.0%
			DSV	12.71%	12.41%	0.111%	23.2%	0.0183	0.1471	0.0823	0.1346	0.1976	0.2936	0.1354	0.0216	0.0898	8.8%
			DSSD	13.61%	13.13%	0.112%	26.2%	0.0230	0.1647	0.0838	0.1341	0.2130	0.3341	0.1518	0.0271	0.0957	9.4%
			MAD	14.23%	13.75%	0.111%	26.9%	0.0242	0.1755	0.0879	0.1290	0.2202	0.3338	0.1555	0.0308	0.0983	10.4%
			SAD	14.54%	14.17%	0.108%	25.3%	0.0233	0.1680	0.0952	0.1229	0.2182	0.3328	0.1526	0.0282	0.1014	10.3%
			VaR0.1	13.29%	12.63%	0.111%	30.6%	0.0243	0.1763	0.0692	0.1421	0.2187	0.3455	0.1558	0.0311	0.0966	9.5%
			VaR0.05	12.72%	12.16%	0.112%	26.6%	0.0226	0.1669	0.0715	0.1556	0.2292	0.3501	0.1502	0.0279	0.0940	9.2%
			CVaR0.1	15.79%	15.51%	0.111%	27.8%	0.0241	0.1754	0.0838	0.1416	0.2216	0.3470	0.1554	0.0308	0.0983	10.0%
			CVaR0.05	15.57%	15.31%	0.111%	27.4%	0.0234	0.1656	0.0751	0.1412	0.2058	0.3160	0.1528	0.0274	0.1014	9.8%
		EWMA	V	7.38%	6.56%	0.117%	27.1%	0.0187	0.1650	0.0822	0.1375	0.2220	0.3375	0.1369	0.0272	0.0823	8.9%
			SD	11.33%	10.53%	0.117%	28.2%	0.0239	0.1705	0.1218	0.1484	0.2338	0.3335	0.1545	0.0291	0.0995	9.8%
			DSV	11.84%	11.10%	0.113%	27.7%	0.0236	0.1762	0.0784	0.1440	0.2401	0.3772	0.1537	0.0311	0.0957	9.7%
			DSSD	13.71%	13.03%	0.115%	28.9%	0.0263	0.1827	0.0811	0.1558	0.2342	0.3654	0.1622	0.0334	0.1035	10.5%
			MAD	11.68%	11.04%	0.113%	26.7%	0.0218	0.1665	0.0824	0.1279	0.2170	0.3417	0.1476	0.0277	0.0929	9.3%
			SAD	12.91%	12.29%	0.113%	27.9%	0.0236	0.1717	0.0744	0.1329	0.2194	0.3479	0.1537	0.0295	0.0963	9.6%
			VaR0.1	10.94%	10.16%	0.117%	27.7%	0.0231	0.1657	0.1175	0.1449	0.2306	0.3308	0.1520	0.0275	0.0973	9.4%
			VaR0.05	10.96%	10.17%	0.117%	27.4%	0.0233	0.1658	0.1236	0.1498	0.2324	0.3311	0.1525	0.0275	0.0980	9.5%
			CVaR0.1	13.21%	11.94%	0.120%	27.2%	0.0364	0.2136	0.1747	0.2384	0.3086	0.4188	0.1908	0.0456	0.1395	15.0%
			CVaR0.05	11.60%	10.29%	0.123%	27.2%	0.0343	0.2007	0.1747	0.2354	0.2945	0.3904	0.1852	0.0403	0.1339	13.9%

Note. Annualised returns for each strategy based on industry group-level indices. V = variance, SD = standard deviation, DSV = downside semivariance, DSSD = downside semideviation, MAD = mean absolute deviation, SAD = semi-absolute deviation, VaR0.1 =

10% value at risk, VaR0.05 = 5% value at risk, CVaR0.1 = 10% conditional value at risk, CVaR0.05 = 5% conditional value at risk.

Source: Author's calculations

Table 9.3

Industry-based strategy annualised returns and annualised risk

	Strategy			Mean	Annualised Return	Fees	Strategy risk: Max Drawdown	V	SD	DSV	DSSD	MAD	SAD	VaR0.1	VaR0.05	CVaR0.1	CVaR0.05
	GlobalIndex			9.98%	9.10%		34.1%	0.0248	0.1773	0.1422	0.2370	0.2892	0.3966	0.1574	0.0314	0.1015	11.6%
Industry	Weekly	Raw	V	11.19%	10.66%	0.101%	27.0%	0.0198	0.1641	0.1243	0.1913	0.2381	0.3235	0.1406	0.0269	0.0887	9.9%
			SD	12.30%	11.72%	0.101%	27.4%	0.0229	0.1762	0.1404	0.2164	0.2613	0.3580	0.1512	0.0310	0.0971	10.8%
			DSV	10.43%	9.73%	0.101%	26.9%	0.0215	0.1716	0.1425	0.2151	0.2570	0.3394	0.1466	0.0294	0.0934	10.5%
			DSSD	14.08%	13.94%	0.100%	23.5%	0.0194	0.1586	0.1422	0.1984	0.2418	0.3159	0.1391	0.0252	0.0947	10.5%
			MAD	12.73%	12.20%	0.101%	27.4%	0.0229	0.1734	0.1410	0.2119	0.2603	0.3571	0.1514	0.0301	0.0971	10.5%
			SAD	13.88%	13.50%	0.100%	35.5%	0.0227	0.1758	0.1337	0.1959	0.2518	0.3446	0.1506	0.0309	0.0955	10.6%
			VaR0.1	10.30%	9.57%	0.101%	27.4%	0.0219	0.1696	0.1454	0.2173	0.2563	0.3420	0.1481	0.0288	0.0948	10.2%
			VaR0.05	8.93%	8.06%	0.101%	27.7%	0.0224	0.1737	0.1461	0.2124	0.2621	0.3522	0.1495	0.0302	0.0975	10.8%
			CVaR0.1	12.18%	11.61%	0.101%	26.2%	0.0226	0.1722	0.1354	0.2180	0.2577	0.3523	0.1502	0.0297	0.0971	10.5%
			CVaR0.05	12.46%	11.87%	0.101%	26.5%	0.0234	0.1769	0.1349	0.2238	0.2652	0.3584	0.1530	0.0313	0.0993	11.0%
		EWMA	V	11.24%	10.71%	0.104%	26.7%	0.0198	0.1614	0.1367	0.1771	0.2288	0.3067	0.1407	0.0260	0.0912	9.8%
			SD	12.85%	12.34%	0.104%	26.8%	0.0229	0.1674	0.1443	0.2015	0.2475	0.3284	0.1513	0.0280	0.1000	10.4%
			DSV	13.12%	12.70%	0.104%	26.5%	0.0219	0.1643	0.1439	0.1861	0.2401	0.3197	0.1478	0.0270	0.0966	10.1%
			DSSD	12.06%	11.45%	0.104%	25.8%	0.0229	0.1708	0.1383	0.2107	0.2573	0.3428	0.1513	0.0292	0.1009	10.7%
			MAD	16.22%	16.18%	0.104%	27.3%	0.0227	0.1655	0.1369	0.1969	0.2389	0.3204	0.1507	0.0274	0.1004	10.3%
			SAD	15.83%	15.74%	0.104%	27.3%	0.0225	0.1658	0.1384	0.1968	0.2395	0.3215	0.1501	0.0275	0.0989	10.2%
			VaR0.1	13.10%	12.63%	0.104%	26.8%	0.0228	0.1670	0.1442	0.2015	0.2473	0.3278	0.1510	0.0279	0.0997	10.4%
			VaR0.05	13.08%	12.60%	0.104%	26.8%	0.0228	0.1667	0.1422	0.2015	0.2474	0.3286	0.1510	0.0278	0.0998	10.3%
			CVaR0.1	6.17%	5.11%	0.104%	27.5%	0.0226	0.1799	0.1435	0.1952	0.2665	0.3670	0.1503	0.0323	0.0979	11.1%
			CVaR0.05	8.47%	7.49%	0.104%	26.7%	0.0237	0.1808	0.1539	0.1953	0.2658	0.3573	0.1541	0.0327	0.1033	11.7%
	Monthly	Raw	V	11.58%	11.34%	0.100%	21.7%	0.0136	0.1386	0.1023	0.1370	0.1847	0.2567	0.1166	0.0192	0.0831	9.5%
			SD	11.44%	10.99%	0.100%	18.7%	0.0172	0.1556	0.1023	0.1684	0.2111	0.2918	0.1311	0.0242	0.0929	11.1%
			DSV	10.85%	10.50%	0.100%	19.1%	0.0145	0.1431	0.1036	0.1722	0.1998	0.2784	0.1205	0.0205	0.0851	9.7%
			DSSD	13.15%	13.07%	0.099%	18.7%	0.0134	0.1409	0.0678	0.1413	0.1830	0.2655	0.1157	0.0199	0.0794	9.1%
			MAD	12.90%	12.61%	0.100%	18.0%	0.0170	0.1513	0.1106	0.1538	0.2046	0.2792	0.1303	0.0229	0.0922	10.5%
			SAD	15.21%	15.18%	0.099%	18.0%	0.0171	0.1513	0.1046	0.1833	0.2024	0.2764	0.1307	0.0229	0.0946	10.6%
			VaR0.1	9.74%	9.27%	0.100%	19.8%	0.0150	0.1446	0.1089	0.1649	0.2025	0.2780	0.1226	0.0209	0.0879	10.0%
			VaR0.05	11.54%	11.14%	0.100%	22.3%	0.0163	0.1501	0.0997	0.1565	0.1980	0.2827	0.1278	0.0225	0.0926	10.6%
			CVaR0.1	11.36%	10.90%	0.100%	19.3%	0.0172	0.1469	0.0997	0.1711	0.2069	0.2924	0.1312	0.0216	0.0939	10.1%
			CVaR0.05	12.18%	11.81%	0.100%	18.9%	0.0171	0.1451	0.0997	0.1424	0.2001	0.2848	0.1307	0.0210	0.0920	9.5%
		EWMA	V	8.14%	7.61%	0.101%	18.4%	0.0140	0.1415	0.1022	0.1686	0.2048	0.2818	0.1185	0.0200	0.0815	9.9%
			SD	11.76%	11.43%	0.100%	19.4%	0.0156	0.1406	0.1032	0.1378	0.1980	0.2791	0.1248	0.0198	0.0881	9.5%
			DSV	8.30%	7.70%	0.101%	22.5%	0.0157	0.1443	0.1133	0.1453	0.2177	0.3124	0.1254	0.0208	0.0864	9.5%
			DSSD	11.13%	10.64%	0.100%	20.3%	0.0173	0.1460	0.1023	0.1329	0.2019	0.2938	0.1316	0.0213	0.0924	9.6%
			MAD	11.82%	11.50%	0.100%	18.9%	0.0153	0.1447	0.0956	0.1442	0.1940	0.2757	0.1238	0.0209	0.0867	9.9%
			SAD	11.91%	11.62%	0.100%	17.8%	0.0152	0.1461	0.0997	0.1340	0.1918	0.2735	0.1231	0.0214	0.0852	9.9%
			VaR0.1	11.28%	10.90%	0.100%	19.4%	0.0157	0.1434	0.1106	0.1416	0.2009	0.2813	0.1255	0.0206	0.0891	9.8%
			VaR0.05	11.60%	11.26%	0.100%	19.2%	0.0156	0.1411	0.1023	0.1391	0.1984	0.2813	0.1251	0.0199	0.0881	9.5%
			CVaR0.1	9.30%	8.89%	0.100%	15.4%	0.0134	0.1371	0.0982	0.1574	0.1993	0.2681	0.1159	0.0188	0.0806	9.4%
			CVaR0.05	12.35%	11.94%	0.100%	21.3%	0.0182	0.1498	0.1228	0.1743	0.2269	0.3165	0.1349	0.0224	0.0938	10.1%

Note. Annualised returns for each strategy based on industry-level indices. V = variance, SD = standard deviation, DSV = downside semivariance, DSSD = downside semideviation, MAD = mean absolute deviation, SAD = semi-absolute deviation, VaR0.1 = 10% value at risk, VaR0.05 = 5% value at risk, CVaR0.1 = 10% conditional value at risk, CVaR0.05 = 5% conditional value at risk.

Source: Author’s calculations

Figure 9.1

Weekly raw risk measure-optimised strategy returns

Note. Graph of the cumulative return of each weekly raw risk measure-optimised strategy. V = variance, SD = standard deviation, DSV = downside semivariance, DSSD = downside semideviation, MAD = mean absolute deviation, SAD = semi-absolute deviation, VaR0.1 = 10% value at risk, VaR0.05 = 5% value at risk, CVaR0.1 = 10% conditional value at risk, CVaR0.05 = 5% conditional value at risk.

Source: Author's calculations

Figure 9.2

Weekly EWMA-optimised strategy returns

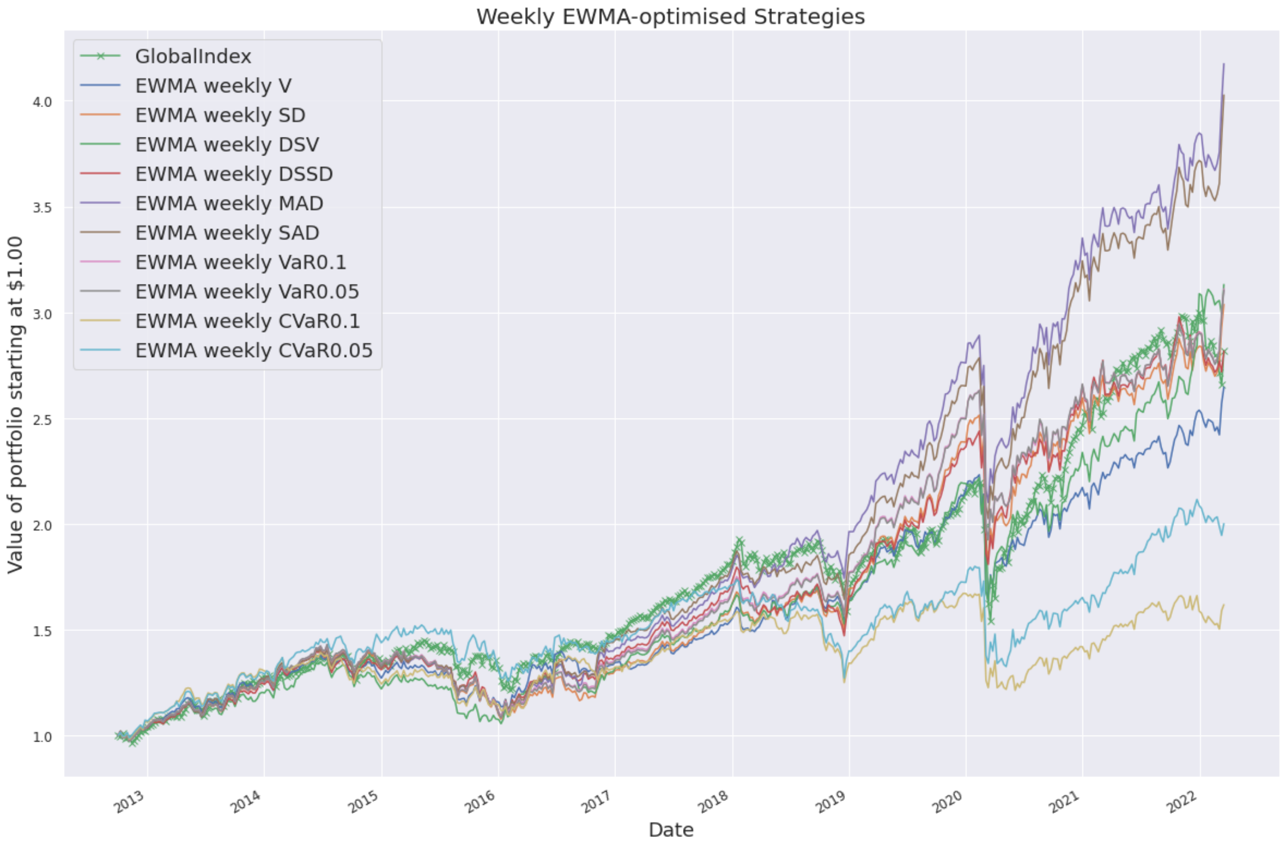

Note. Graph of the cumulative return of each weekly EWMA-optimised strategy. V = variance, SD = standard deviation, DSV = downside semivariance, DSSD = downside semideviation, MAD = mean absolute deviation, SAD = semi-absolute deviation, VaR0.1 = 10% value at risk, VaR0.05 = 5% value at risk, CVaR0.1 = 10% conditional value at risk, CVaR0.05 = 5% conditional value at risk.

Source: Author’s calculations

Figure 9.3

Monthly raw risk measure-optimised strategy returns

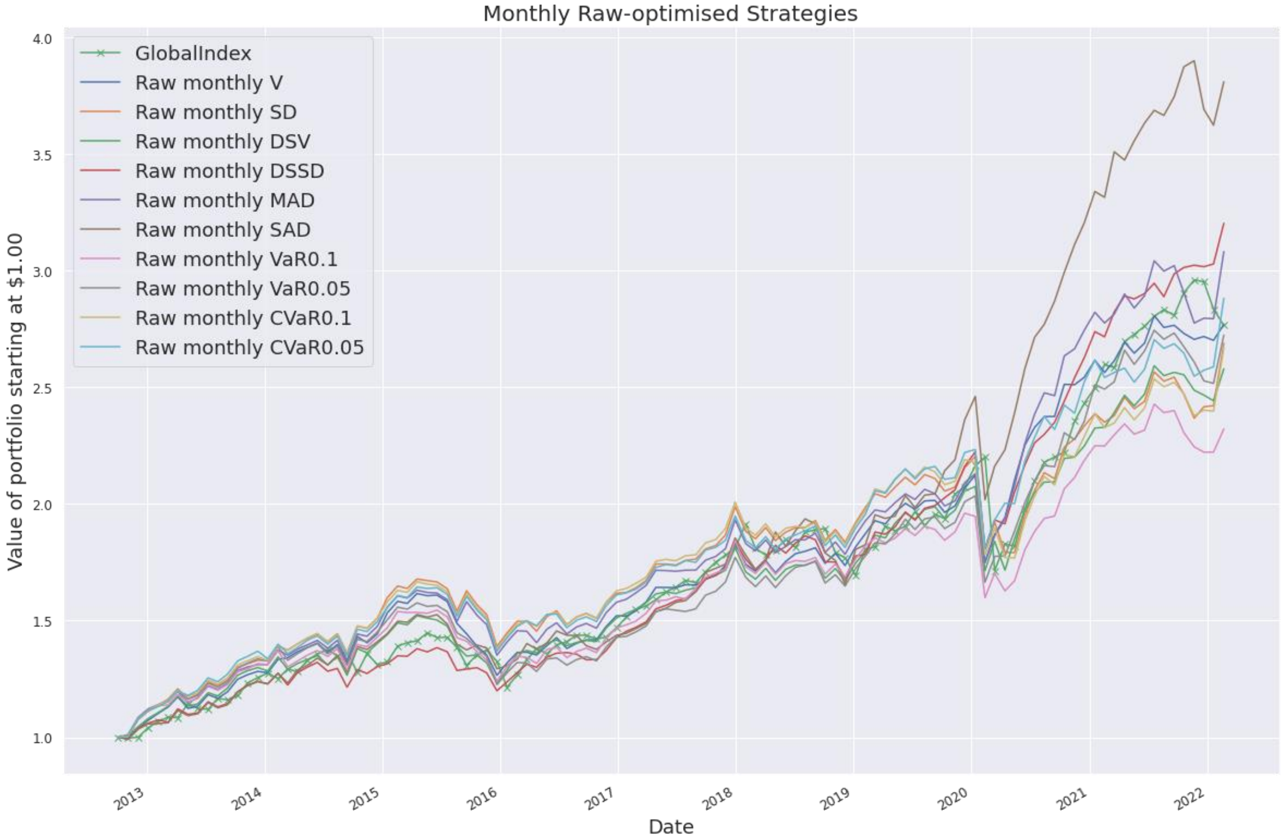

Note. Graph of the cumulative return of each monthly raw risk measure-optimised strategy. V = variance, SD = standard deviation, DSV = downside semivariance, DSSD = downside semideviation, MAD = mean absolute deviation, SAD = semi-absolute deviation, VaR0.1 = 10% value at risk, VaR0.05 = 5% value at risk, CVaR0.1 = 10% conditional value at risk, CVaR0.05 = 5% conditional value at risk.

Source: Author's calculations

Figure 9.4

Monthly EWMA-optimised strategy returns

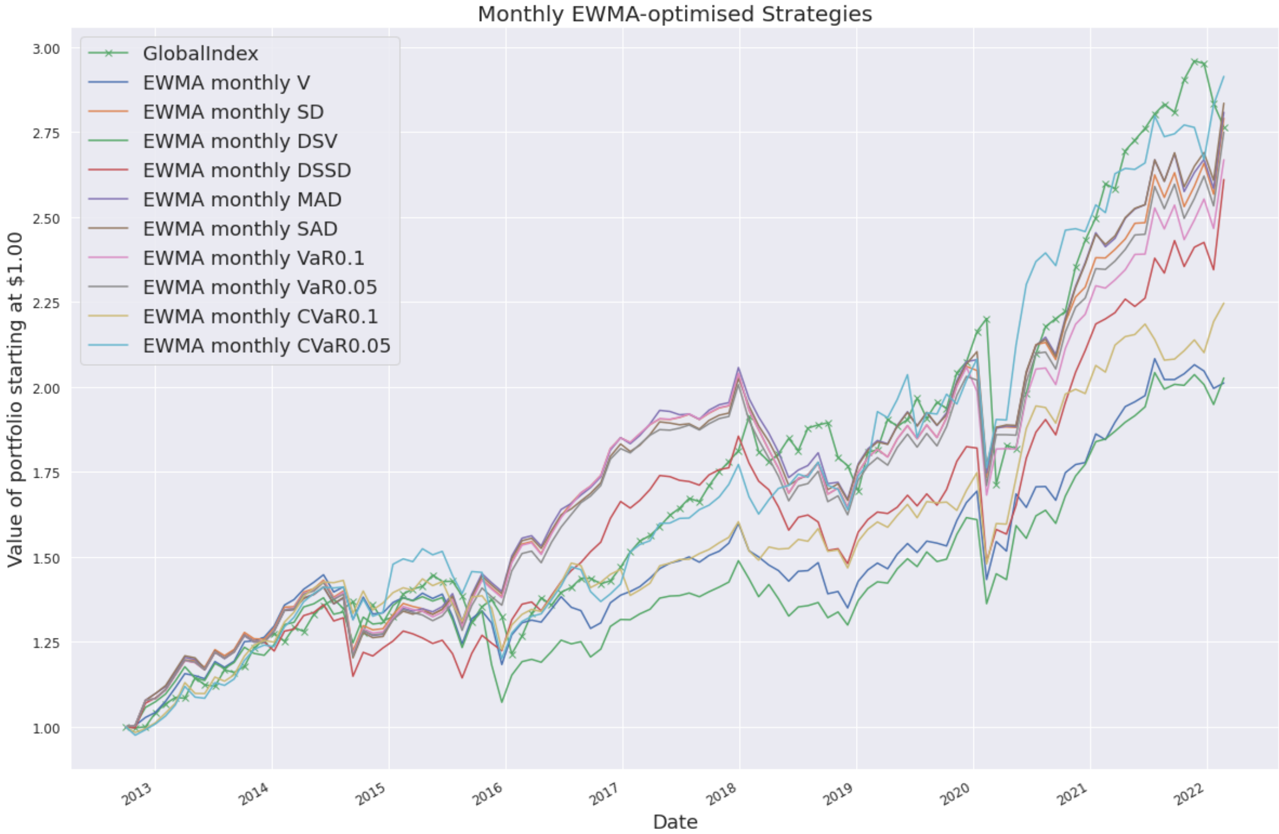

Note. Graph of the cumulative return of each monthly EWMA-optimised strategy. V = variance, SD = standard deviation, DSV = downside semivariance, DSSD = downside semideviation, MAD = mean absolute deviation, SAD = semi-absolute deviation, VaR0.1 = 10% value at risk, VaR0.05 = 5% value at risk, CVaR0.1 = 10% conditional value at risk, CVaR0.05 = 5% conditional value at risk.

Source: Author's calculations