Enhanced Minimum Variance Optimisation

A Pragmatic Approach

Author: Lakhoo Lala Bernisha Janti

Supervisor: Professor David Bradfield

Co-supervisor: Mr Tobias Brandt

MINOR DISSERTATION PRESENTED IN PARTIAL FULFILMENT OF THE REQUIREMENTS FOR THE DEGREE OF

MASTERS IN ADVANCED ANALYTICS AND DECISION SCIENCES

M. Sc.

IN THE DEPARTMENT OF STATISTICAL SCIENCES

UNIVERSITY OF CAPE TOWN

AUGUST 2016

Publication

I hereby grant the University free license to publish this dissertation in whole or part in any format the University deems fit.

Plagiarism Declaration

I. I know the meaning of plagiarism and declare that all the work in the project, save for that which is properly acknowledge, is my own.

II. I have used the Harvard referencing guide for citation and referencing. Each significant contribution to, and quotation in, this project from the work, or works, of other people has been acknowledged through citation and reference.

III. I have not allowed, and will not allow anyone to copy my work with the intention of passing it off as his or her own.

Lakhoo Lala Bernisha Janti

August 2016

Abstract

Since the establishment of Markowitz's theory, numerous studies have been carried out over the past six decades or so that cover the benefits, limitations, modifications and enhancements of Mean Variance (MV) optimisation. This study endeavours to extend on this, by means of adding factors to the minimum variance framework, which would increase the likelihood of outperforming both the market and the minimum variance portfolio (MVP). An analysis of the impact of these factor tilts on the MVP is carried out in the South African environment, represented by the FTSE-JSE Shareholder weighted Index as the benchmark portfolio. The main objective is to examine if the systematic and robust methods employed, which involve the incorporation of factor tilts into the multicriteria problem, together with covariance shrinkage – improve the performance of the MVP. The factor tilts examined include Active Distance, Concentration and Volume. Additionally, the constant correlation model is employed in the estimation of the shrinkage intensity, structured covariance target and shrinkage estimator. The results of this study showed that with specific levels of factor tilting, one can generally improve both absolute and risk-adjusted performance and lower concentration levels in comparison to both the MVP and benchmark. Additionally, lower turnover levels were observed across all tilted portfolios, relative to the MVP. Furthermore, covariance shrinkage enhanced all portfolio statistics examined, but significant improvement was noted on drawdown levels, capture ratios and risk. This is in contrast to the results obtained when the standard sample covariance matrix was employed.

Acknowledgement

The path to completing my master's degree began two years ago. This was not a path I walked alone, but included many people that continuously provided a strong support base which aided me every step of the way.

I would like to thank my supervisor, Prof. Bradfield, for his continuous guidance, patience and invaluable contribution to both my coursework and dissertation. Your immense knowledge in the field of quantitative investing and portfolio management is profound. Thank you for always providing insight and feedback to me without delay.

I would also like to thank Tobias Brandt, my co-supervisor who was always on hand to assist me with technical aspects of my thesis; you have a wealth of knowledge and added considerable value always.

To my sister, Nataneya, you have been my guiding light, your encouragement, meticulousness and unwavering involvement in my studies is priceless, I can always depend on you.

To my parents, Janti and Rajna Lala thank you for always being there for me throughout my life, I am because you are. To my sister, Veruschka thank you for coming back and being there for me again.

To Lennie, thank you for being a constant in my life and supporting me in this process.

A special acknowledgement to Yashin Gopi from BNP Paribas Securities for all the assistance you offered on data and technical aspects.

I would also like to thank Associate Prof. Lubbe from the statistics department for understanding that I was a part-time student and always accommodating and assisting me with my countless problems.

Finally to my friends Shawn and Siya, I appreciate all your encouragement along the way.

Table of Contents

List of Figures

List of Tables

1. Introduction

1.1 Background and Objectives

One of the key considerations of Markowitz portfolio theory is diversification with the fundamental objective being that an investor can lessen risk inherent in a portfolio by holding an optimally weighted set of assets which are not perfectly positively correlated. This optimal allocation is determined through Markowitz mean variance optimisation (hereinafter referred to as MV optimisation), which utilises a mathematical formulation where the trade-off between two key parameters – expected risk and return - is considered. Simply put, this implies that an investor should minimise portfolio expected risk while maximising portfolio expected return.

There is a plethora of research available extending on the fundamental theory of MV optimisation. These generalisations and improvements have refined associated theories on the effects of risk on valuation and a key extension in this respect has been minimum variance investing. Studies by Jagannathan and Ma (2003) and Scherer (2011) presented research in support of the superiority of minimum variance portfolios (MVP), which only prioritises risk minimisation in the portfolio construction process. Baker and Haugen (2012) also presented evidence that contradicts the general principle in finance established in the Capital Asset Pricing Model (CAPM): *that greater risk is expected to produce higher returns.*

Academic literature on the persistence of the low volatility anomaly in the South African environment is abundant. An initial study carried out by van Rensburg and Robertson (2003), who found a negative relationship between a stock's beta and return, was followed by research by Khuzwayo (2011) and Panulo (2014), who both presented evidence that indicated portfolios created on the low-volatility premise outperformed higher risk portfolios.

More recently in the South African environment, Oladele and Bradfield (2016) found that using several techniques to construct low risk portfolios including the MVP, resulted in improved performance versus the FTSE/JSE[1] All Share market capitalisation-weighted index.

[1] FTSE-JSE - Financial Times Stock Exchange-Johannesburg Stock Exchange

A primary feature of the MVP is the elimination of expected returns from the formulation. Markowitz theory is postulated on known future expected returns and risk - in practice however these estimates are not always known and must be estimated and are thus subject to estimation error. This issue was discussed by Michaud (1989) who contended that MV optimisers are *estimation error maximisers*, since they have a propensity to maximise the effects of errors in these input parameters. Best and Grauer (1991) also emphasised the sensitivity of MV portfolios to changes in individual asset mean returns. Evidence from the above studies, in addition to research by Bradshaw (2004) and Bonini, Zanetti, Bianchini and Salvi (2010) that highlighted that expected return forecasting is neither superior nor scientific, all support the case for minimum variance investing.

There are investment practitioners who believe in the importance of MVP in portfolio management, but who are dissuaded by its limitations – which are due to several concerns surrounding implementation, concentration, liquidity and unstable non-robust solutions.
The main objective of this study is to investigate if there are pragmatic, robust and theoretically sound methods available that can enhance the performance of the standard MVP and diminish the impact of these limitations in portfolio optimisation. These methods involved the systematic introduction of additional factors into the objective function in the form of *factor tilts (enhancements)*.
Some of primary methods investigated and on which this study builds are discussed in more detail in the paragraphs to follow.

Yanushevsky and Yanushevsky's (2015) postulation around improving MV optimisation is innovative. The main idea introduced in the paper around enhancing MV optimisation was the incorporation of an additional parameter – an *average volume indicator* which measures the average trading volume of shares of a portfolio security for a pre-defined time period as a percentage of its total float number of shares – into the objective function. This technical indicator measures the market demand and liquidity levels for a security.
The *average volume indicator* is introduced into the optimisation framework as an additional parameter - the $J\ Index$:

$$J = (\boldsymbol{X} - \boldsymbol{Z})'(\boldsymbol{X} - \boldsymbol{Z})$$

Where $\boldsymbol{X}$ is the vector of its optimal weights and $\boldsymbol{Z}$ is a vector of proportions of the portfolio's securities based on their average volume traded. The optimal solution, $\boldsymbol{X}$, is sought around minimising the distance between the optimal weights, x_i, for each security i and z_i, the weight of security i in the portfolio based on its average volume traded.

King (2007) in his seminal research tried to address a key concern with MV optimisation - overly concentrated portfolios. His enhancement came in the form of the Herfindahl–Hirshman Index (HHI), which acts as a gauge of the level of concentration of a portfolio. His key concern around concentration was addressed through an adjustment of the objective function, similar to that introduced in Yanushevsky and Yanushevsky (2015). While the aforementioned authors, sought an optimal solution around a measure they believed could enhance portfolio returns, King (2007) sought to mitigate the effect of extremities on the optimal solution, which directly equates to setting the $\boldsymbol{Z}$ vector in the $J\,Index$ to $\boldsymbol{0}$.

Both studies create a platform and methodology that allows for the introduction of multiple factors into the MVP optimisation framework. The average volume indicator described above, tilts the optimal portfolio toward stocks with higher liquidity, the traded prices of which will be more reflective of the fundamental prices, whilst the integration of HHI effectively tilts the portfolio to higher levels of diversification.

Similarly the $\boldsymbol{Z}$ vector can be set up to introduce any *factor* tilt to the MVP, thus increasing the exposure of the portfolio to any economical, fundamental or technical indicator. With the advent of the $J\,Index$, any property can be incorporated to enhance the MV optimisation framework. A factor can thus be employed to create optimal portfolios that incorporate attributes an investment practitioner may believe are important in terms of his or her investment philosophy and portfolio construction process.

The following factor tilts were examined, since the research presented in multiple studies supported the contention of a relationship between the factor and performance:

1. **Concentration** – Although research presented by Kacperczyk, Sialm and Zheng (2005) linked higher levels of fund concentration to improved fund performance and research by Baks, Busse and Green (2007) also found benefit in diversification which was introduced in a two case portfolio, the key goal in integrating this factor in this study is to reduce concentration levels of the MVP, rather than increase it to improve performance. A primary drawback of the MVP is the high concentration in weightings observed (refer to Khuzwayo, 2011). In an examination of the concentration levels across a set of low volatility portfolios (including the MVP and the equally weighted portfolio) in the South African environment, he found that the MVP resulted in the highest concentration across all portfolios investigated.
2. **Volume** - There is research abound which contends that volume is an important factor in determining the level of economic activity and price. Ying (1966), through a series of empirical studies found evidence of a positive relationship between price and volume. Studies by Crouch (1970) and Karpoff (1987) provided theoretical and empirical evidence of this as well.
3. **Active Distance** - which is the Euclidean distance between the fund and benchmark - was introduced by Bradfield, Maritz and Swartz (2005) and is conceptually similar to *Active Share* introduced by Cremers and Petajistoy (2009), who found that funds with a higher active share persistently outperform their benchmarks over time while their counterparts (funds with lower active share) tend to underperform on both a pre- and post-fee basis. Both measures give investors more insight into how actively managed their funds are.

The enhanced optimisation framework established is a suitable platform to examine if funds with exposure to these factors are able to enhance the returns of the MVP in the South African context.

Another aspect which was integrated and examined in this study was covariance shrinkage (Ledoit and Wolf, (2003, 2004)), which addressed problems around estimation error maximisation associated with the utilisation of the sample covariance matrix, $\boldsymbol{S}$, in MV optimisation. These issues were discussed at length in Jobson and Korkie (1981) as well. Some of the implications include extreme coefficients which are inclined to have large amounts of error and are thus assigned the most radical bets within the portfolio. In addition, $\boldsymbol{S}$ imposes too little structure and one of its main disadvantages is that it contains a lot of estimation error when the number of time points (T) is of comparable or smaller order than the number of individual stocks (N). This leads to an ill-conditioned singular covariance matrix with negative eigenvalues, which is not invertible. Given the multiple problems with $\boldsymbol{S}$, Ledoit and Wolf (2003) proposed a shrinkage method which in effect finds a compromise between structure and unbiasedness.

Covariance shrinkage is an alternative technique that attempts to find a more suitable estimator than $\boldsymbol{S}$. The authors introduced, $\boldsymbol{F}$, a highly structured estimator, which is less prone to estimation error, but often misspecified and biased.

Ledoit and Wolf (2003) discussed multiple approaches in determining the structured target, $\boldsymbol{F}$. However, in Ledoit and Wolf (2004), the primary model analysed was the constant correlation model. The authors found that the out-of-sample results were positive and observed increases in the realised information ratio, as well as significant outperformance relative to multiple market value-weighted benchmarks constructed on the US stock market. Munro and Bradfield (2016) and Fletcher (2009) replicated and extended on the Ledoit and Wolf studies in South Africa and the United Kingdom respectively - they have similar findings. This constant correlation model performs best in terms of both risk and return relative to other methods employed.

Given these findings, this paper employs the constant correlation model to determine $\boldsymbol{F}$ and thus $\boldsymbol{\Sigma}_S$ - where $\boldsymbol{\Sigma}_S = \omega\boldsymbol{F} + (1-\omega)\boldsymbol{S}$ and ω is the optimal shrinkage intensity. The method of determining ω is called shrinkage. In essence, shrinkage finds a compromise ($\boldsymbol{\Sigma}_S$) between two extremes and results in an estimator that is generally more robust and performs better than the initial extreme estimators, $\boldsymbol{F}$ and $\boldsymbol{S}$.

The overall objective of this study is to draw on and incorporate the principal ideas introduced in Yanushevsky and Yanushevsky (2015) and King (2007) and Ledoit and Wolf (2004) as well as adapt and extend on these methodologies and examine their impact on the MVP in the South African environment - which is represented by the FTSE-JSE Shareholder Weighted Index (SWIX) universe of stocks. In pursuit of the aforementioned objectives - the study covered the period April 2006 to April 2016.

1.2 Outline of Study

The pertinent areas of focus in this study will be discussed in the sections that follow and are listed below -

1. Chapter 2 – The Literature Review lays the foundation and discusses both international and domestic research available on the key limitations and benefits of MV optimisation. In addition to this, numerous methods on key enhancements of the standard MV framework are contextualised. The conjectures around; expected return, the MVP and more specifically domestic research around the persistence of the low volatility anomaly, and factor tilts and their role in improving returns are examined as well.
2. Chapters 3 and 4 provide a detailed description of the methodology and data used in; the estimation of the shrinkage intensity (ω), the calculation of $\boldsymbol{F}, \boldsymbol{S}$ and $\boldsymbol{\Sigma_S}$, the derivation of the enhanced MVP objective functions and finally the back-testing methodology implemented to determine the levels at which factor tilts should be incorporated (preference parameter estimation).
3. Chapter 5 provides a detailed analysis on the shrinkage parameter results and the out-of-sample performance of the *enhanced portfolios* formulated using the two covariance structures, $\boldsymbol{\Sigma_S}$ and $\boldsymbol{S}$, described above.
4. The summary of findings, as well as limitations and ideas on further research are discussed in Chapter 6.

2. Literature Review

This chapter discusses both international and domestic research available on the key limitations and benefits of MV optimisation. In addition to this, numerous methods on key enhancements of the standard MV framework are contextualised and discussed, which include minimum variance investing, covariance shrinkage, conjectures around expected return, as well as factor tilts and their role in improving returns.

2.1 Limitations of Markowitz Portfolio Theory (MPT)

1. Michaud, R.1989. The Markowitz Optimisation Enigma: Is 'Optimised' Optimal? Financial Analysts Journal, 45(1), 31-42

   A limitation with Markowitz mean variance optimisers is that it results in highly concentrated portfolios. Such portfolios are difficult to accept by the mainstream investment community. The main issue is its propensity to maximise the effects of errors in the input assumptions (risk and return estimates are subject to estimation error). So in effect, securities with large estimated returns, negative correlations and small variances will be over-weighted and favoured in the portfolio et cetera. The author pointed out that MV optimisers magnify the impact of estimation errors and in an unconstrained environment can yield results that are poorer to an equally weighted portfolio as Jobson and Korkie (1981) showed. Michaud termed MV optimisers *estimation-error maximisers*. Hence, where these errors are at extremes, this will result in a maximisation of the impact of the estimation error of the final portfolio weights - these characteristics of portfolio optimisers often prove to be hard for portfolio managers to accept as tools for portfolio construction.

   He also highlighted some other efficiencies and their impact on the MV optimisation:

   - Liquidity is an important factor that MV ignores, which is an important consideration in the investment management profession. Michaud described liquidity as the percentage of a security's market capitalisation that a portfolio holds. Thus if a portfolio holds a considerable portion of the capitalization of a relatively small and /or illiquid company, if a significant proportion of the company shares are traded, the price of the counter will be impacted (market impact). Compared to unconstrained MV frontier, the liquidity constrained portfolio plots below the original MV frontier and provides a lower return and less risk reduction.

- Unstable optimal solutions exist in some cases. For example, any small change in an input assumption can lead to significant shifts in the composition of the portfolio; the MV optimisation problem is not robust to small alterations in input estimators. The author attributes behavioural instability to an ill-conditioned covariance matrix and draws attention to the importance of using meaningful estimates based on adequate historical data. Chopra and Ziemba (1993) and Black and Litterman (1992) also highlighted this.

2. <u>Best, M.J. and Grauer, R. 1991. On the sensitivity of mean variance efficient portfolios to changes in asset means: Some analytical and computational results. Review of Financial Studies 4(2), 315–342</u>

Another important aspect highlighted in this paper is the sensitivity of MV efficient portfolios to changes in the means of individual assets. When budget constraints are introduced into the problem, all results point to the impact that changes in these inputs have on the MV-efficient portfolio's weights, mean, and variance. When non-negativity constraints are also imposed on the problem however, results indicate that positively weighted MV-efficient portfolios' weights are extremely sensitive to changes in asset means, but the portfolios' returns are not. For example, by changing the mean of one asset by a small percentage, rejects nearly half of the assets from the portfolio, but will not have any significant impact on the portfolio return and variance. Britten-Jones (1999) used an OLS regression and found that sampling error in estimates of MV efficient portfolio weights in a global portfolio was significant and large.

2.2 Benefits of MPT

Despite its short comings, mean variance optimisation triumphs other techniques, especially with regard to the manner in which it is able to integrate the objectives of a portfolio with constraints specified by clients. Portfolio Optimisers also make efficient use of information and they are able to process vast quantities of information rapidly. This benefit is particularly useful to large financial institutions that need to see the effect of new information on their portfolios in real time. The MV framework additionally allows for institutions to implement their style objectives and market outlooks by adjusting exposure the benchmark portfolio and the relevant stock universe. The advent of the efficient frontier has been essential in the development of the field of financial economics – it preceded and thus informed the Sharpe-Lintner Capital Asset Pricing model as well as made an important differentiation between systematic and diversifiable risk (Michaud, 1989). With regard to diversification, optimisers allow alterations to risk exposure to be made. Even amidst globalisation amongst markets and economies, during which one would expect that different assets would become increasingly correlated with one another, diversification benefits such as - less portfolio volatility and reduced risk of incurring losses - are still available for the exploitation of investors.

In general, for practitioners, theory suggests that MV efficient portfolios can play an important role in portfolio management, but reported the difficulties in implementing MV analysis. Research also shows that MV analysis is central to many asset pricing theories as well as empirical tests of those theories.

2.3 Improvement of Mean-Variance optimisation

Studies that have addressed this by enhancing the objective function or improving the parameter inputs are listed below.

2.3.1 Enhancement of the objective function formulation

- King, D. 2007. Portfolio optimisation and diversification. Journal of Asset Management, 8(5), 296–307

Again the idea around portfolio optimisation and the resultant overly concentrated portfolios are explored. With an adjustment to the objective function, one can alter the levels of portfolio diversification using the same standard mean-variance optimisation framework. The author alluded to constraints that are often used to mitigate the effect of extremities that result from classical MV optimisation. But this will impose a pre-existing view on the overall process which will inadvertently introduce hindsight or anchoring bias. Employing a nonparametric approach which is based on the classical approach, the author created different levels of diversified portfolios which were linked to a *diversification preference parameter* (similar to the risk parameter). The method to measure diversification is the Herfindahl–Hirshman Index (HHI). The Index is calculated as follows:

$$HHI = \sum_{i=1}^{n} {w_i}^2$$

Where w_i = market capitalisation weight (market share of company i in a particular index or industry)

The HHI ranges between one for an industry with a single constituent and $1/N$ for an industry or index with N participants, each with an equal market share. The author applied the HHI to portfolios and measured the level of portfolio diversification by replacing market share with portfolio weight.

The benefit of diversification was introduced by Baks et al. (2007) in a two case portfolio. For example, with the same level of risk and return, having a measure of diversification can aid in differentiating between the two portfolios. HHI was introduced into the MV-optimiser framework by introducing it to the standard objective function. The results as expected show that at higher levels of the *diversification preference parameter*, the optimiser places a higher penalty on concentrated portfolios - so the resultant portfolios are more diversified in composition.

The introduction of HHI into the standard MV optimisation framework produced results that are very similar to those produced in Resampling Efficiency (RE) portfolio optimisation methodology developed by Michaud (1998). This is a promising finding which may indicate that at an appropriate level of the diversification preference parameter in a single optimisation, there may be no need to perform multiple sampling from a multivariate distribution, that is, RE optimisation.

- <u>Yanushevsky, R. & Yanushevsky, D. 2015. An approach to improve mean - variance portfolio optimisation model. Journal of Asset Management, 16 (3), 209-219</u>

  The authors illustrated that the traditional Markowitz portfolio can be improved by introducing an additional parameter – the average trading volume indicator (which measures the average trading volume of shares of a portfolio security for a pre-defined time period as a percentage of its total float number of shares). This measures the market demand for a security and is used to quantify the potential price increase of a security. Since the authors proposed that the correlation between price and volume is weak, the expected returns of securities cannot be adjusted to incorporate this technical indicator. Rather the standard MV optimisation framework is adapted by incorporating the $J\ Index$ into the objective function -

$$J = (\boldsymbol{X} - \boldsymbol{Z})'(\boldsymbol{X} - \boldsymbol{Z})$$

  Where $\boldsymbol{X}$ is the vector of optimal weights and $\boldsymbol{Z}$ is a vector of proportions of the portfolio's securities based on the expected increase of their price
  In traditional MV optimisation, the portfolios are usually highly concentrated in a few assets. The resulting portfolio weights are extreme and the out-of-sample performance is poor. The means and variances are influenced by outliers and this leads to extreme deviations which are over-weighted and small deviations which are either removed or reduced in the portfolio.

The approach used in the paper reduces to a $l_2\,regularisation$ consequently, but has a sound economic foundation. Regularisation refers to a process of introducing additional information to solve an ill-posed problem. Where the approach differs from previous studies such as Jagannathan and Ma (2003) that used regularisation by imposing constraints on the allocation vector, the inclusion of the volume indicator has economic justification – it adds useful stock market information and is less onerous to use in practice.

The authors also addressed the approach of using historical returns as an input into the MV-optimiser. The main concern here is the soundness of using the past behaviour of a stock as a meaningful predictor of the future performance of a stock, therefore assuming that past trends have information about the future price of the stock. This method fails to take into account any new information that did not exist when the historical data was produced. In order to improve on the above method of computing returns, risk-based models like the Capital Asset Pricing Model (CAPM), the Fama-French three-factor and Carhart four-factor models were introduced. CAPM uses the equity risk premium (the difference between the expected return of the market and the risk-free rate) and the market beta (a measure of risk exposure of the stock to the market) to determine the expected return of a stock. Fama-French (1993) improved on CAPM by introducing factors that measure company size and price-to-book value. Carhart (1997) extended on this and added a momentum factor. With the advent of the internet however, information relating to stock market activity is readily accessible and estimates of company performance are available online. Given the above advancements in forecasts and availability of information, MV portfolios should be modified. The average trading volume indicator (AVT) is an incisive statistic that can be employed to coalesce the information mentioned above as well as investor preferences, to invest in companies that they are familiar with (Massa and Simonov (2006)). Studies conducted by Ying (1966), amongst others, show that the relationship between the volume traded and the potential price increase is weak but positive. The objective function which maximises expected return and minimises risk, is adapted to include the AVT indicator. The portfolio will be tilted toward stocks with greater potential price increase. With the formulation of the $J\ Index$, the optimal portfolio weights x_i, are sought in a region close to the preferable proportions of the portfolio's securities based on the expected increase of their price, namely, z_i. The results indicated that the enhancement does indeed add value over the same subset of stocks which are held in a market index.

- Anagnostopoulos, K. and Mamanis, G. 2010. A portfolio optimization model with three objectives and discrete variables. Computers & Operations Research, 37 (7), 1285-1297

The paper focused on including additional criteria and constraints into the MV optimal framework. The study considered a suitable-portfolio investor, who is concerned with much more than just the return and variance of the portfolio, but with various other aspects such as; the number of assets in his portfolio, the maximum amount allocated to each asset, social responsibility et cetera.
The solution is no longer sought on an efficient line, but rather on an efficient surface in a higher dimension space -$\mathcal{R} > 2$. Solving the multi-criteria model is cumbersome, since objectives may be discrete or non-smooth in nature. The objective function is formulated as a three criteria (returns, variance and the number of assets) model with constraints that restrict the proportion of funds invested in assets or groups of assets. This constraint will limit the number of small holdings in the portfolio and overinvestment in assets with similar characteristics.

The model proposed contains discrete and non-smooth characteristics. In order to solve such a multi-criteria problem, the authors experimented with metaheuristic optimisation techniques like multi-objective evolutionary algorithms (MOEAs). The empirical results indicated that the MOEAs generate surfaces with good diversity characteristics, thus these computational algorithms provide a generalised MV approach with additional portfolios which are not MV efficient, but that have a fewer number of assets.

2.3.2 Enhancement through improved parameter inputs – Covariance Shrinkage

1. Fletcher, J. 2009. Risk Reduction and Mean-Variance Analysis: An Empirical Investigation. Journal of Business Finance & Accounting, 36(7-8), 951–971

The author looked at analysing the performance in the UK market of both the Global minimum variance (GMV) and Tracking Error variance (TEV) portfolios. He employed different forms of the covariance matrix - errors in the covariance matrix do not have as a big impact as those in expected returns (Chopra and Ziemba (1993)), but still may impact results in cases where the number of assets is large in comparison to the time period used. Since the author chose to focus on the estimation errors introduced in the covariance matrix and its impact on the GMV portfolio, he did not make use of expected returns in his analysis.
Various Models of the covariance matrix examined:

Constant correlation matrix (Elton and Gruber, 1973) – this approach assumed an equal correlation for all asset pairs, which is set to the average of sample all pair wise correlations. The covariance matrix was built using the sample standard deviations and sample correlations.

The shrinkage approach proposed by Ledoit and Wolf (2003) applies a shrinkage intensity level, k,in the following equation to determine the shrinkage estimator of the covariance matrix, S_{shrink} :

$$\boldsymbol{S_{shrink}} = \frac{k}{T}\boldsymbol{F} + \left(1 - \frac{k}{T}\right)\boldsymbol{S}$$

T represents the number of time-series observations used to calculate the covariance matrix, $\boldsymbol{S}$, and $\boldsymbol{F}$ is the target covariance matrix. $\boldsymbol{F}$ used in this approach will have a lower estimation error then $\boldsymbol{S}$, but will have biases due to the simplifying assumptions used. Higher $\frac{k}{T}$ ratios, will ensure greater shrinkage of $\boldsymbol{S}$ to $\boldsymbol{F}$. Ledoit and Wolf explained that k will increase as the estimation error of $\boldsymbol{S}$ increases, will decrease as the bias in $\boldsymbol{F}$ increases and will decrease as the covariance between estimation errors between $\boldsymbol{S}$ and $\boldsymbol{F}$ increase. He set $\boldsymbol{F}$ to the constant correlation matrix described above and the **Sharpe single index model** (1963), which uses a variant of the single-index model, the market model, to estimate returns:

$$r_{it} = \alpha_i + \beta_i r_{mt} + e_{it}$$

where r_{it} is the return of stock i in period t, α_i is the unique expected return of security i, β_i, the sensitivity of stock i to market movements, r_{mt}, the return on the market in period t, and e_{it} is the unique risky return of security i in period t and has a mean of zero and variance σ_{ei}^2. The model assumes the correlation between the residuals of any asset pair is zero. The benefit of using the single-index is the reduction in the number of inputs required to calculate the covariance matrix. The covariance matrix only requires N asset variances along the diagonal, the market index variance, N stock betas. The author continued to describe two multi-index models similar to that of Carhart (1997) and Fama and French (1993). He then computed a fourth covariance matrix which combines the different approaches described above. He found that performance and risk relative to both the benchmark and sample covariance matrix is much improved.

2. Ledoit, O. and Wolf, M. 2004. Honey, I Shrunk the Sample Covariance Matrix. Journal of Portfolio Management, 30(4), 110–119.

The paper addressed the various drawbacks around using the sample covariance matrix in MV-optimisation and suggested an alternative technique for estimating the covariance matrix, based on *shrinking* extremes towards the centre. The sample covariance matrix is combined with a target structured estimator, using a convex linear formulation, $\delta\boldsymbol{F} + (1-\delta)\boldsymbol{S}$, where δ, the shrinkage constant, is the weight that is assigned to the structured estimator. In order to determine the *optimal* shrinkage constant, one needs to minimize the expected distance between the shrinkage estimator and the true covariance matrix. This is done using the Frobenius norm defined as follows:

$$\|Z\|^2 = \sum_{i=1}^{N}\sum_{j=1}^{N} z_{ij}^2$$

The norm difference between the shrinkage estimator and true covariance reduces to a quadratic loss function as defined below -

$$L(\delta) = \|\delta \boldsymbol{F} + (1-\delta)\boldsymbol{S} - \boldsymbol{\Sigma}\|^2$$

In Ledoit and Wolf (2003), the best estimate of the optimal shrinkage constant, $\delta^* = \frac{\hat{\kappa}}{T}$, was obtained by minimising the expected value of the loss function, where

$$\hat{\kappa} = \frac{\hat{\pi} - \hat{\rho}}{\hat{\gamma}}$$

$\hat{\pi}$ = sum of the asymptotic variances in the sample covariance,

$\hat{\rho}$ = sum asymptotic covariance's between entries in the sample covariance and shrinkage target,

$\hat{\gamma}$ = squared differences between the terms of the sample covariance matrix and shrinkage target

The authors found that the enhanced covariance matrix determined through shrinkage, significantly improves the information ratio of the portfolio manager.

Munro and Bradfield (2016) conducted a similar study to the aforementioned, but did so in the South African equity market and have similar findings. They found significant differences between the structures of the mean variance portfolio covariance estimators. They also found that the structured target yields better results (lower risk) in an out-of-sample period than the sample covariance. Some of the structured targets they investigate are listed below:

- **Sharpe's (1963) single-index model** estimates stock returns as follows:

$$x_{it} = \alpha_i + \beta_i x_{mt} + \varepsilon_{it}$$

where x_{it} is the stock return i in period t, α_i is the excess return (alpha) for stock i, β_i is the beta for stock i and ε_i is the residual returns of stock i, where $\varepsilon_i \sim N(0, \sigma_i)$. The linear model is generated by regressing the returns of stock i on the market return. Once the slope estimates $\boldsymbol{b}$ (N x 1 vector of stock level estimates) and $\boldsymbol{D}$ (N x N diagonal matrix of residual variance estimates, d_{ii}) are calculated, the covariance matrix estimate can be formulated using the individual stock level beta estimates, the market sample variance s_m^2 and $\boldsymbol{D}$ as follows:

$$\boldsymbol{F} = s_m^2 \boldsymbol{b}\boldsymbol{b}' + \boldsymbol{D}$$

- **Principal Component Analysis** orthogonally transforms the covariance matrix to linearly uncorrelated variables, that is, eigenvectors of the covariance matrix. Each principal component explains a proportion of variance - so far fewer components are required to explain the covariance structure. The number of principal components employed by Munro and Bradfield (2016) explain 80% of the variance.

- **Average covariance matrix** sets the diagonal terms of the sample covariance matrix, $\boldsymbol{S}$, to the average of all stock level variances. The off-diagonal elements are set to the average of all off-diagonal covariances ($\forall (i,j)^{th}$ stock pairs, where $i \neq j$). This is given below:

$$f_{ij} = \begin{cases} \bar{s}_{ii} & if\ i = j \\ \bar{s}_{ij} & if\ i \neq j \end{cases}$$

  Where f_{ij} is the $(i,j)^{th}$ entry of $\boldsymbol{F}$, the structured covariance matrix.

3. Matoti, L. 2009. Building a statistical linear factor model and a global minimum variance portfolio using estimated covariance matrices. Unpublished thesis. University of Cape Town.

Matoti (2009) uses several estimation techniques to determine the non-positive definite covariance matrix (M_{ndp}) applied to emerging market data. He discussed four transformation methods used in the covariance estimation technique, namely, the eigenvector, the $arctan$ shrinkage, the $tanh$ selected shrinkage and the area minimisation methods, were which he applied to M_{ndp} to transform into a positive definite covariance matrix. Two of the methods employed used a shrinkage transformation method ($arctan$ shrinkage and the $tanh$ selected shrinkage) which transformed the non-positive definite correlation (C_{ndp}) matrices into positive definite correlation matrices (C_{dp}). This method shrunk the off-diagonal elements of the C_{ndp} until the resultant matrix was positive definite or all the off-diagonal elements were zero (i.e. I (Identity matrix)). The author employed a function f to transform each off-diagonal element (i.e. shrink to zero). Both $f_1 = arctan$ and $f_2 = tanh$ selected had to satisfy the following conditions:

1) f must be a strictly increasing function
2) must be an odd function
3) $f(0) = 0$

He compared the resultant covariance matrices obtained using the different shrinkage functions against each other using Lindskog's Euclidean distance measure and selected the two functions that resulted in the shortest distance and which were the best estimates of the true covariance matrix.

The results indicated that the $arctan$ transformation method was the best transformation method to use in restoring the positive definiteness property of the sample covariance matrix.

2.3.3 Other techniques employed to improve MPT

1. Fabozzi, F. J., Kolm, P. N., Pachamanova, D. A. and Focardi, S. M. 2007. Robust portfolio optimisation. Journal of Portfolio Management, 33(3), 40-48.

   Robust optimisation incorporates the estimation error directly into the optimiser and is typically used with conventional robust statistical estimation methods. The idea of robust statistics is that it generally promotes the removal (or down-weighting) of outliers, thus, incorporating the idea around robust statistics. The author extended this idea to mainstream finance in the form of robust optimisation. The method tries to minimise the worst-case return for a given confidence region. It requires that the optimal solution remain optimal for all values of the expected returns that are close to the estimates of the true expected returns (μ), that is, all estimates of μ are within a confidence interval around μ. However, Scherer (2007) emphasised that although the estimation error is reduced partly, the methodology introduced by Fabozzi et al. (2007) could be written in terms of ordinary Bayesian shrinkage estimation where the resultant set of optimal portfolios along the efficient frontier stays the same.

2. Jorion, P. 1985. International Portfolio Diversification with Estimation Risk. Journal of business, 58(3), 259-278

   Jorian (1985) also highlighted the importance of international diversification in reducing portfolio risk and enhancing returns, but emphasised that one should take uncertainty of estimates of input parameters into account when forming expectations. He also continued to explain that estimators are less subject to estimation error than the classical sample mean. He questioned alternative estimators of expected returns and their implications for the apparent gains from diversification. By shrinking the sample means toward a common mean. The out-of-sample performance of the optimal portfolio is substantially increased. Findings indicate that the classical method vastly overestimates the possible gains in average returns; instead, benefits from diversification are more likely to accrue from a reduction in risk. Stein (1955) also showed that shrinkage estimators improve out-of-sample performance and have important implications for portfolio selection.

3. Skylogiannis, V. and Xu, J. 2009. Mean Variance Optimization and Beyond. Unpublished thesis. Stanford University.

   The authors addressed the uncertainties introduced in MV optimisation since investors do not know the true values of expected return or risk. Using the historical moments, introduces estimation errors which are maximised by the optimiser. They compared the classical approach to a

Bayesian regularisation. The Bayesian MAP (maximum a posterior) estimator approach introduces Bayesian priors to model the uncertainty in parameter. The authors assumed a multivariate normal distribution and solved the MAP estimators of expected return and risk using Euler equations to derive the closed form solution. They used monthly data for five indices, including the risk- free asset (three-month Treasury bill), to create optimal portfolios under both the maximum likelihood (MLE) approach and the Bayesian MAP. Next, the authors built a risk function to measure out-of-sample performance constructed under the different approaches. They introduced the *certainty equivalent loss* measure, which gauges the loss in portfolio value due to estimation error. The following process was used next: 5000 random samples based on the standard multivariate normal distribution were created and the MV-portfolios for both the MLE and Bayesian MAP approaches for each sample were calculated. The average of certainty equivalent loss across all samples is then calculated. The Bayesian MAP reduces the loss of portfolio value compared to the true optimal portfolio. However, the Sharpe ratio of the tangency portfolio shrinks by a factor of $\frac{T-N-2}{T+1}$ compared to the MLE approach.

4. Bouchaud, J.P., Potters, M. and Aguilar, J.P. 1997. Missing Information and Asset Allocation. http://xxx.lanl.gov/cond-mat/9707042

   Random matrix theory to reduce the error maximisation of Markowitz optimal portfolios and directly control the issue of over concentration is addressed by Bouchaud, Potters and Aguilar (1997).

   They discuss the instability of the covariance matrix and average return over time, which incurs large costs, since the assets selected in the optimisation process change over time.

   In addition they highlight the issues around having partial information and how this impacts the statistical parameters used (e.g. return and covariance matrix), and stress that the optimal portfolio should reflect the lack of information and keep a certain level of diversification.

They introduce 'diversification indicators' (constraints on generalised entropies akin to thermodynamics) which measure the level of concentration in the portfolio and relate these to the information content in the portfolio. A high information content would be representative of a highly concentrated portfolio, while an equally weighted portfolio would have a minimal information content (low concentration). In doing this they generate a "sub-efficient border", based on the effective number of assets, which plots below the unconstrained efficient frontier.

Wilcox and Gebbie (2007) aimed to determine empirical correlations in price fluctuations of daily sampled price data of South African shares in a reliable way where missing data and thin trading were significant. The issues pertained to noise, finiteness of time series, missing data, and thin trading. They examined random matrix theory (RMT) which had been applied to calibrate and reduce the effects of noise in financial time series. Correlation matrices are computed for the data under investigation and quantities associated with these matrices may be compared to those of random matrices. For example, in several studies of shares traded in the S&P 500 it was found that the estimation of covariances was dominated by random noise. In examining the results of the JSE main board, the authors found some agreement between the distributions of RMT predictors (including Wishart distribution for eigenvalues and the Wigner surmise for eigenvalue spacing) and the spectral properties of the calculated correlation matrix estimator.

5. <u>Black, F. and Litterman, R. 1990. Asset Allocation: Combining Investors Views with Market Equilibrium. Fixed Income Research. Goldman, Sachs & Company.</u>

Black and Litterman (1990) tried to highlight the shortcomings of traditional mean-variance optimisation in specifying expected returns as a starting point. They showed that a neutral starting point (used as an expected return input) (provided by Market consensus) yields the benchmark weights under conditions of market equilibrium. The neutral starting point is a set of returns that would "clear the market" – if all investors had identical views. The model uses equilibrium expected returns that are generated from CAPM. All investors have identical views; therefore the resultant weightings should correspond to the "market" or benchmark weightings. If the market portfolio is the optimal portfolio – reverse optimisation shows that the "neutral returns" – are the derived return inputs.

The theory further allows for investors only to specify returns on assets that they have views on. In conventional optimisation, an adjustment to return (of assets involved in the view) can be:

1) Direct - small changes to consensus return inputs, but this results in extreme unrealistic portfolio weightings and significant changes to assets weights that were unaffected by view.
2) Bayesian - weighted between consensus return and direct view, where weighting scheme reflects confidence in view.

Both 1) and 2) result in unrealistic portfolios – that make changes to asset weightings that had no view to begin with.

The Black and Litterman approach – adjusts the view to be covariance consistent, allowing the view to adjust the input returns on other assets (which have no views imposed) according to their covariance with the assets which have the view imposed. Expected returns of all assets are affected by the view imposed, so expected returns are adjusted away from the consensus values – in a way consistent with underlying covariance and view being expressed and as a result the only change is in weights are for those assets that had a view imposed, all other asset weights stay the same.

2.4 Minimum variance and expected return conjectures

2.4.1 The case for minimum variance investing

1. Haugen, R. and Baker, N. 1991. The efficient market inefficiency of capitalization-weighted stock portfolios. The Journal of Portfolio Management 17(3), 35-40.

   The authors purported that capital-weighted stock portfolios are an inefficient way of investing when the following assumptions hold - investors differ with respect to risk and expected return inputs, short-selling is restricted, when investment income is taxed, when some investment alternatives are not included in the target index, or when foreign investors are in the domestic capital market.

   They constructed low-volatility (MVP) portfolios using the MV-optimisation framework, over the period - 1972-1989, using 1000 stocks with the biggest market capitalisation over all U.S exchanges and markets. The performance was then compared to the cap-weighted stock portfolio and a set of random portfolios. The purpose of including this set of random portfolios, which were constructed to have the equivalent structural weights and turnover to the low volatility portfolio, was to prove that the low risk attributes of the low volatility portfolio had not been achieved by chance and that the probability of this was very low in fact. Over five-year rolling periods, the authors showed that this portfolio had lower risk relative to the benchmark and outperformed it as well.

2. Baker, N. & Haugen, R. 2012. Low Risk Stocks Outperform within All Observable Markets of the World.
   Available at: http://ssrn.com/abstract=2055431

   The authors proved comprehensively that low risk stocks have higher expected returns. Their test period covered the period 1990 to 2011 and included 21 developed markets and 12 emerging markets. The low risk anomaly is evident in all jurisdictions and persists over time. In fact, they showed that higher relative risk yields a negative reward across all jurisdictions. Low risk decile stocks outperformed their high risk counterparts on both an absolute basis (where the returns across all markets were higher) and on a risk-adjusted basis (where on average Sharpe ratios were 75 percent higher). In addition, they criticised direct investing in capitalisation-weighted indices, which they believe over-value growth stocks which are in general highly volatile.

3. Oladele, O. & Bradfield, D. 2016. Low Volatility sector-based portfolios: The South African Case. ORiON, 32(1), 55-78

Again, the research in the South African market indicated that portfolios created on the low-volatility premise outperformed higher risk portfolios.
Research on the low volatility anomaly was initially carried out by van Rensburg and Robertson (2003), who found a negative relationship between a stock's beta and return. This was supported by Kruger, Strugnell and Gilbert (2011), who showed the persistence of the low volatility anomaly when they assessed the refined beta estimate. An out-of-sample examination of a series of different low volatility portfolios was carried out by Khuzwayo (2011) and Panulo (2014), who both found conclusive evidence of the South African low volatility anomaly as well.
More recently, Oladele and Bradfield (2016), focused on nine FTSE-JSE sectors where back tests were carried out over the period 2003 – 2013. The portfolios were constructed on various low volatility methodologies which targeted low volatility, low beta, maximum diversification or low correlation, et cetera. Some of the techniques studied included; equally weighting low beta stocks, standard minimum variance optimisation, the low volatility single index model, maximum diversification portfolio optimisation, et cetera. The results indicated that all methodologies outperformed both FTSE-JSE market capitalisation-weighted indices (ALSI[2] and SWIX). The most salient features were higher Sharpe ratios and lower risk. There was a marked improvement in other statistics such as drawdowns and information ratios.

[2] FTSE-JSE All Share (Free Float) market-capitalisation Index

Matoti (2009) and Wilcox and Gebbie (2016) show how the low beta anomaly can be understood from a more rigorous mathematical perspective if one retains the principle of no-arbitrage, then the pricing kernel is necessarily non-linear to accommodate the data.

2.4.2 Expected return conjecture

1. Bradshaw, T. 2004. How Do Analysts Use Their Earnings Forecasts in Generating Stock Recommendations? The Accounting Review 79(1), 25-50.

   The author examined the disparity in analyst forecasts and valuation estimates. He examined four valuation models based on the residual-income model, price-earnings-to-growth (PEG) model and projections of long-term earnings growth. The author found that analyst recommendations have little correlation to the residual income model, but can be explained by the PEG, long-term earnings growth models and heuristic models based on valuation. Although this is the case, there is no evidence of any explanatory power of these models in predicting future excess returns.

2. Bonin, S., Zanetti, L., Bianchini, R. & Salvi, A. 2010. Target Price Accuracy in Equity Research. Journal of Business Finance & Accounting 37(9- 10), 1177 – 1217.

   The paper explored the ability of target prices to predict future stock prices. Target prices measure the potential change in the value of a stock which is a requisite input to the investment decision making process. The expectation is that target prices are accurate estimates of future stock prices, as implied by the efficient markets hypothesis. And that prediction errors should be distributed around zero with known variance. The authors highlighted that there is no comprehensive valuation methodology in place on which to base target prices and thus accuracy in target prices is questionable. Their findings indicate that target prices are consistently biased and that the size of the bias increases as the growth in the value implied increases. This could be seen as a direct and deliberate action on the part of analysts or research providers to bias markets to their advantage.

The authors created a four-fold accuracy metric which measure both intra-period and end-of-period accuracy of the analyst price forecast. This metric was then compared to the actual stock returns over the period. They disproved their hypothesis and found that the frequency of accurate prediction was extremely low and that the size of the prediction error was very large, auto-correlated, non-mean reverting and positive in signs, which suggests the existence of a systematic upward bias.

2.5 A note on Active Distance, Concentration and Volume

2.5.1 Active Distance

1. Cremers, K. and Petajistoy, A. 2009. How Active Is Your Fund Manager? A New Measure That Predicts Performance. Review of Financial Studies, 22(9), 3329-3365

   Active Share, which measures how different a fund's holdings are from its benchmark, was introduced as an active measure that would give investors more insight into how actively managed their funds were. The findings indicate that the measure predicts fund performance, specifically that funds with a higher active share persistently outperform their benchmarks over time while their counterparts (funds with lower active share) tend to underperform on both a pre- and post-fee basis.

2. Schlanger, T., Philips, C.B. & LaBarge, K. P. 2012. The search for outperformance: Evaluating 'Active Share'. Vanguard research report.

   In this study, the authors found no conclusive evidence that active share is a good predictor of fund outperformance. They sub-divided the test period into an evaluation and performance period and calculated certain measures for around 903 long-only active funds selected from the

U.S equity mutual fund category. Of these, 446 funds were identified as having high tracking errors and active shares in excess of 60% (indicative of *concentrated funds*). The measures calculated in the evaluation period included both active share and tracking error and were linked to fund performance measured in the performance period. There was no clear evidence that these concentrated funds delivered superior performance over the performance period. In fact, concentrated funds underperformed over the period and the dispersion in excess returns was much higher relative to other groupings.

This study is supported by Frazzini, Friedman, and Pomorski (2016), who using the same data set as Cremers and Petajisto (2009), found that the correlation between active share and fund outperformance is weak and that active share is driven by its relationship to the benchmark type. They also argued that there is little economic rationale in justifying a preference for Active Share.

2.5.2 Concentration

King (2007) through his seminal work on HHI and its introduction into the standard MV-optimisation framework as a potential diversifier, tried to address a major shortcoming in MV-optimisation - highly concentrated optimal solutions.

Several studies link levels of fund concentration to fund performance. These are discussed briefly below:

1. Kacperczyk, M., Sialm, C. & Zheng, L. 2005. On the industry concentration of actively managed equity mutual funds. Journal of Finance, 60(4), 1983-2011.

   In their study on the relationship between Industry concentration and fund performance of actively managed U.S. mutual funds from 1984 to 1999, the authors found that funds that have high concentration tend to deliver better performance after controlling for risk and style differences.

They believed that active managers, who concentrate their portfolios in industries that they have an informational advantage in, tend to add value. Other studies that have similar findings include Baks et al. (2007), who studied U.S mutual funds over a 25-year period and concluded that managers with highly concentrated funds tend to outperform their more diversified counterparts, but also highlighted using a two portfolio case example, the importance of using diversification in identifying the superior portfolio when the risk and returns are the same.

2. Chen, X. & Lai, Y. 2015. On the concentration of mutual fund portfolio holdings: Evidence from Taiwan. Research in International Business and Finance, 33, 268-286.

   The authors find contradictory results in their survey of equity mutual funds in the Taiwanese market. In bull markets, funds with high levels of concentration are positively correlated with risk-adjusted performance, but in market recessions this relationship reverses, and funds that are broadly diversified tend to outperform.

2.5.3 Volume

Karpoff, J.M. 1987. The relation between price changes and trading volume: a survey. Journal of Financial and Quantitative Analysis, 22(1), 109–126

The author reviewed both the theoretical and empirical relationships between price and volume. Part of his research, involves a review on the previous studies on the relationship between and price and volume. He then established with empirical tests that there is a positive relationship between the two series. These findings correlate and draw on previous evidence by Ying (1966), Crouch (1970) et cetera.

The price-volume relation was motivated with empirical tests of Ying (1966), who contended that volume is an important factor that determines the level of economic activity and price. The relationship was examined using various statistical techniques applied to the Standard and Poor's 500 composite daily price return index (adjusted for dividend rates) and daily volume of stock sales on the NYSE (measured as a ratio of the number of shares traded divided by the total shares outstanding) over the period, 1957-1962. Some of the results of these empirical tests carried out by Ying (1966: 676) are given below:

“ - A small volume is usually accompanied by a fall in price.
- A large volume is usually accompanied by a rise in price.
- A large increase in volume is usually accompanied by either a large rise in price or a large fall in price.
- A large volume is usually followed by a rise in price. ”

Crouch (1970) also found a positive correlation for both market indices and individual stocks between the daily absolute price change and volume data. Research by Westerfield (1977) found the same relation, but in a sample of daily price changes and volumes for 315 common stocks.

Godfrey, Granger and Morgenstern (1964), however, presented new evidence from several data series, including daily and transactional data for individual stocks, but found no correlation between prices or the absolute values of price differences and volume.

This chapter highlighted the plethora of literature available on which this study is based, which in summary addresses both international and domestic literature available on the limitations and benefits of MV optimisation. In addition, it discussed the numerous methods on enhancements of the standard MV framework. The conjectures around; expected returns, the MVP and more specifically domestic research around the persistence of the low volatility anomaly, and factor tilts and their roles in addressing limitations and improving returns are examined as well.

3. Methodology

This chapter provides detail of the methodology used in; the covariance shrinkage estimation process, the derivation of the enhanced MVP objective functions and the back-testing methodology implemented to determine the most appropriate level of factor tilts that should be considered (parameter estimation).

3.1 Covariance Shrinkage estimation

Covariance shrinkage was a technique introduced by Ledoit and Wolf (2003), and the primary objective was to create an alternative to the sample covariance matrix. This approach systematically reduces the estimation error and constricts coefficients to more central values. Another benefit of using the proposed shrinkage technique is that there is no requirement for a semi-positive definite covariance matrix. In several finite sample statistical decision theory studies regarding shrinkage estimators - where the sample size, T is of a smaller order than the number of stocks, N - there is a disintegration of the theory. This is because the loss functions require invertible covariance structures.

The technique involves two primary concepts; the construction of the shrinkage target and thus the shrinkage estimator, as well as the development of the shrinkage intensity parameter. The general form of the shrinkage estimator incorporates two extremes:

1. The *unbiased* sample covariance matrix, $\boldsymbol{S}$ (since the expected value is the true covariance matrix, $\boldsymbol{\Sigma}$), which can be interpreted as a N-factor model, where N is the number of stocks. $\boldsymbol{S}$ is unbiased but prone to estimation error maximisation (Michaud,1989). It also performed poorly out of sample (Jobson and Korkie, 1981). The sample covariance estimator is represented as follows:

$$\boldsymbol{S} = \frac{1}{T-1}\sum_{t=1}^{T}(x_t - \bar{x})(x - \bar{x})'$$

   Where T is the sample size, $\boldsymbol{x_t}$ is an N x 1 vector of stock returns in period t, and $\boldsymbol{\bar{x}}$ is the N x 1 vector of mean of returns of each stock i, $i = 1,2, \dots n$.

2. The shrinkage target, $\boldsymbol{F}$, represents a *biased* but more *structured* K-factor model, where $1 < K < N$. The advantages of using this alternative are clear. It uses a smaller number of parameters and any information on the nature and number of factors is not required. One of the shrinkage target estimators introduced by Ledoit and Wolf (2004) and tested in the South African environment by Munro and Bradfield (2016) was the **constant correlation model**.

The author of this paper (hereinafter referred to as *the author*) made use of the **constant correlation model,** described by Ledoit and Wolf (2004), as the structured covariance target in the study. The reason for this is primarily due to the superior out of sample performance it generated relative to other estimators employed. Munro and Bradfield (2016) found that the model produced minimum variance portfolios with the lowest out-of-sample variance over the period 2006-2009. The model also yielded the best returns using both an information coefficient (IC) of 1 and 0.1. They also found that the enhanced estimator significantly improved the information ratio of the portfolio manager.

The IC introduced by Grinold and Kahn (1999), measures the relationship between a manager's realised and forecasted alphas. The actual forward-returns based on the respective IC were used as inputs into the MV optimisation framework, along with the selected shrinkage estimate. Another benefit of the model is that it is simpler to implement than the Sharpe single-index model described in 2.3.2 above. Ledoit and Wolf (2004) found that the performance of this model was relatively similar to the Shape single-factor matrix.

The construction of the constant correlation model was described as follows:

i. If s_{ij} is defined as the sample covariance of the $(i,j)^{th}$ stock pair then the sample correlation of the $(i,j)^{th}$ stock pair is given by:

$$r_{ij} = \frac{s_{ii}}{\sqrt{s_{ii}\ s_{ij}}}$$

ii. The average sample correlation was calculated by:

$$\bar{r} = \frac{2}{N(N-1)} \sum_{i=1}^{N-1} \sum_{j=i+1}^{N} r_{ij}$$

iii. The $(i,j)^{th}$ element of the $\boldsymbol{F}$, the shrinkage target matrix, was then calculated as follows:

$$f_{ij} = \begin{cases} s_{ii} & if\ i = j \\ \bar{r}\sqrt{s_{ii}\ s_{ij}} & if\ i \neq j \end{cases}$$

Once the shrinkage target, $\boldsymbol{F}$ and the sample covariance matrix, $\boldsymbol{S}$ were determined, an estimate of the shrinkage intensity, δ, could be determined and was always restricted to values between zero and one. A linear combination of $\boldsymbol{F}$ and $\boldsymbol{S}$ could then be formulated to create the shrinkage estimate $\widehat{\boldsymbol{\Sigma}}_{\boldsymbol{s}}$:

$$\widehat{\boldsymbol{\Sigma}}_s = \omega\boldsymbol{F} + (1-\omega)\boldsymbol{S}$$

Where ω denotes the *estimated* shrinkage intensity.

An optimal shrinkage intensity, δ can be obtained by minimising the quadratic loss function (squared distance between the shrinkage estimator and the true covariance matrix) based on the Frobenius norm. There is no requirement for an invertible covariance matrix:

$$\|Z\|^2 = \sum_{i=1}^{N}\sum_{j=1}^{N} z_{ij}^2$$

Where $\boldsymbol{Z}$ is an N x N symmetric matrix of z_{ij}

This results in the quadratic loss function:

$$L(\delta) = \|\delta\boldsymbol{F} + (1-\delta)\boldsymbol{S} - \boldsymbol{\Sigma}\|^2 \qquad (3.1.1)$$

By minimising the loss function (3.1.1), with respect to δ, Ledoit and Wolf (2003) showed that the optimal shrinkage coefficient takes the form:

$$\kappa = \frac{\pi - \rho}{\gamma}$$

Where π is the sum of the asymptotic variances of $\boldsymbol{S}$ scaled by $\sqrt{T}$, ρ is the sum of asymptotic covariances of elements of $\boldsymbol{F}$ and $\boldsymbol{S}$, scaled by $\sqrt{T}$, and γ is the Frobenius norm of the differences between $\boldsymbol{F}$ and $\boldsymbol{S}$. To approximate each of the aforementioned parameters, Ledoit and Wolf (2004) used the following consistent estimators, all of which were proven by them:

1. $\hat{\pi} = \sum_{i=1}^{N}\sum_{i=1}^{N}\hat{\pi}_{ij}$, where $\hat{\pi}_{ij} = \frac{1}{T}\sum_{t=1}^{T}\{(x_{it}-\bar{x}_i)(x_{jt}-\bar{x}_j) - s_{ij}\}^2$ (3.1.2)

Where x_{it} is the stock return i in period t, $\bar{x}_i$ is the mean of returns of each stock i, $i = 1,2, \dots n$ for the period t, s_{ij} is the covariance of the $(i,j)^{th}$ stock pair

2. $\hat{\rho} = \sum_{i=1}^{N} \hat{\pi}_{ii} + \sum_{i=1}^{N} \sum_{\substack{j=1 \\ i \neq j}}^{N} \frac{\bar{r}}{2} \left(\sqrt{\frac{s_{jj}}{s_{ii}}}\, \hat{\vartheta}_{ii,ij} + \sqrt{\frac{s_{ii}}{s_{jj}}}\, \hat{\vartheta}_{jj,ij} \right)$ (3.1.3.1)

   Where $\hat{\vartheta}_{ii,ij} = \frac{1}{T} \sum_{t=1}^{T} \{(x_{it} - \bar{x}_i)^2 - s_{ii}\} \{(x_{it} - \bar{x}_i)(x_{jt} - \bar{x}_j) - s_{ij}\}$ (3.1.3.2)

   $\hat{\vartheta}_{jj,ij} = \frac{1}{T} \sum_{t=1}^{T} \left\{ (x_{jt} - \bar{x}_j)^2 - s_{jj} \right\} \{(x_{it} - \bar{x}_i)(x_{jt} - \bar{x}_j) - s_{ij}\}$ (3.1.3.3)

3. $\hat{\gamma} = \sum_{i=1}^{N} \sum_{j=1}^{N} (f_{ij} - s_{ij})^2$ (3.1.4)

   Where f_{ij} is the $(i,j)^{th}$ of the shrinkage target matrix.

4. Hence $\hat{\kappa} = \frac{\hat{\pi} - \hat{\rho}}{\hat{\gamma}}$, and $\lambda = max\left\{0, min\left\{\frac{\hat{\kappa}}{T}, 1\right\}\right\}$ (3.1.5)

The aforementioned procedure of determining the shrinkage target $(\boldsymbol{F})$, shrinkage intensity estimate (ω) and finally the shrinkage estimator $(\widehat{\boldsymbol{\Sigma}}_{\boldsymbol{s}})$, basically entailed finding an *optimal* combination of an unbiased estimator like the sample covariance matrix, $\boldsymbol{S}$ and the biased but more structured target matrix, $\boldsymbol{F}$**.** This led to a far more robust structured shrinkage estimator that reduced both bias and estimation error. The shrinkage estimator $(\widehat{\boldsymbol{\Sigma}}_{\boldsymbol{s}})$ and the true covariance matrix $(\boldsymbol{\Sigma})$ will hereinafter be represented by the symbols, $\boldsymbol{G}$ and $\boldsymbol{W}$ respectively**.**

3.2 Enhancement Parameter estimation

Markowitz (1952) established the standard MV-optimisation framework which considers two key parameters - namely future expected portfolio return and variance. So for a given set of expected returns, covariances and variances of returns, an investor could identify a suitable portfolio on the efficient frontier that either returned the portfolio with the highest level of return for a given level of variance, or the portfolio with the lowest variance for a given level of return.

If $\boldsymbol{X} = \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{bmatrix}$ and $\boldsymbol{E} = \begin{bmatrix} e_1 \\ e_2 \\ \vdots \\ e_n \end{bmatrix}$ are vectors of the weightings and expected returns of stocks held in a portfolio respectively and $\boldsymbol{W} = \begin{pmatrix} \sigma_{11} & \cdots & \sigma_{1n} \\ \vdots & \ddots & \vdots \\ \sigma_{n1} & \cdots & \sigma_{nn} \end{pmatrix}$ is the expected covariance matrix of stock returns, with σ_{ij} being the covariance of returns of stock i and j, then the portfolio expected return, E_P and portfolio expected variance, W_P can be calculated as follows:

1. $E_P = \boldsymbol{E}'\boldsymbol{X}$
2. $W_P = \boldsymbol{X}'\boldsymbol{W}\boldsymbol{X}$

In practice, one needs to estimate these parameters, since they are unknown.

The objective function in the standard MV-optimisation framework can be formulated as a two-criterion problem, where portfolio return is maximised and portfolio risk is minimised and is represented as follows:

$$\min \left\{-\boldsymbol{E}'\boldsymbol{X} + \frac{\tau}{2}\boldsymbol{X}'\boldsymbol{W}\boldsymbol{X}\right\} \tag{3.2.1}$$

$$\boldsymbol{s.t} \quad x_i \geq 0, i = 1, \ldots, n \tag{3.2.2}$$

$$\boldsymbol{X}'\boldsymbol{1}_n = 1 \tag{3.2.3}$$

Where $\boldsymbol{1}_n$ is a unit vector and τ $(\tau > 0)$ is the conventional risk aversion parameter, which effectively measures the trade-off between risk and return. An investor will select a portfolio that depends on his or her risk aversion. High values of τ penalise variance of the portfolio and so correspond to more risk-averse investors. This is in contrast to smaller values of τ which minimise the impact of portfolio variance within the objective function and thus correspond to less risk-averse investors.

Optimal portfolios were created by varying τ and minimising the objective in 3.2.1, above, at each level of τ, subjected to constraints 3.2.2 and 3.2.3 which ensured that no short selling was allowed and that stock weightings always added up to one, so that the portfolio was not geared in any sense.

Given the above affirmation on the parameter estimation, a key aspect where this analysis differs to prior research such as Yanushevsky and Yanushevsky (2015) is in the estimation of expected returns, where the authors used expert analyst forecast returns to proxy expected return. However there is no conclusive research that indicates that analyst forecast returns are either superior or scientific. Bradshaw (2004) showed that there was no evidence that analyst forecasts are superior estimates of expected return and in fact found that buy-and-hold investing does better than valuation models that include analyst estimates. Bonini et al. (2010) found that analyst ability to forecast returns accurately is very limited and that forecasting errors increase as the growth in the price increases.

Several studies also support the superiority of minimum variance portfolios (MVP). Baker and Haugen (2012) presented evidence that contradicts the general principle in finance which was established in the Capital Asset Pricing Model (CAPM): *that greater risk is expected to produce higher returns*. Over multiple jurisdictions and periods, they found that the low risk stock deciles outperform the high risk stock deciles on both an absolute and risk-adjusted basis. Oladele and Bradfield (2016) tested and extended on this phenomenon in the South African environment and found similar results that are backed by both international studies by Haugen and Baker (1991) and Jagannathan and Ma (2003), as well as domestic studies by van Rensburg and Robertson (2003) and Khuzwayo (2011). They used several techniques to construct low risk portfolios including the MVP and found that all these methodologies have superior performance versus the FTSE/JSE All Share market capitalisation-weighted index. Scherer (2011) also showed that portfolios constructed using MVP principles will implicitly pick up on any risk based pricing anomalies and tend to hold low volatility and low beta stocks. These portfolios tend to have lower realised risk versus the traditional market capitalisation weighted indices.

In view of the above inferences, with respect to the use of expected returns and the MVP, the author, in her estimation of expected returns, chose to assume a constant expected return across stocks. The convex quadratic programming problem in 3.2.1 above is thus simplified to -

$$\min\left\{\frac{\tau}{2}\boldsymbol{X}'\boldsymbol{W}\boldsymbol{X}\right\} \qquad (3.2.4)$$

The MVP holds the least risky combination of stocks and prioritises risk minimisation. This implicitly eliminates the need for expected returns in the construction of the MV-optimisation framework.

3.2.1 Enhancements to Standard MV-Optimisation Framework

3.2.1.1 Enhancement (Tilt) using the Average Volume Indicator

The first enhancement investigated was the methodology introduced by Yanushevsky and Yanushevsky (2015). Here the MV-optimisation framework was enhanced by incorporating the $J\ Index$ into the objective function, given below -

$$J = (\boldsymbol{X} - \boldsymbol{Z})'(\boldsymbol{X} - \boldsymbol{Z}) \qquad \text{(A.1)}$$

In order to calculate the main component of this index, $\boldsymbol{Z} = \begin{bmatrix} z_1 \\ z_2 \\ \vdots \\ z_n \end{bmatrix}$, which is a vector of the weights of portfolio instruments based on the average volume indicator, which one can partly describe as a tilt to more liquid stocks, one requires both the volume and free float shares in issue for all stocks.

The $\boldsymbol{Z}$ vector was calculated as follows:

- For each stock, the average daily volume indicator (av) was calculated using the following steps:
  i. $vd_{it} = \frac{vol_{it}}{ffsiss_{it}}$ (A.2.1)

     Where vol_{it} is the volume for stock i, on day t and

     $ffsiss_{it}$ is the free float shares in issue for stock i, on day t

  ii. $av_i = \frac{\sum_{t=1}^{m} vd_{it}}{m}$ (A.2.2)

     Where av_i is the average daily measure of the proportion of a stock's free float shares that are traded over the preceding m-day period - m was set at 50 business days.
- Once the aforementioned measures were calculated, z_i could be determined:

$$z_i = \frac{av_i}{\sum_{i=1}^{N} av_i} \tag{A.3}$$

The incorporation of the $J\ Index$ into the objective function resulted in

$$\min \left\{ -\boldsymbol{E}'\boldsymbol{X} + \frac{\tau}{2}\boldsymbol{X}'\boldsymbol{W}\boldsymbol{X} + \frac{\lambda}{2}J \right\} \tag{A.4}$$

It can be transformed to an equivalent formulation which is derived below and subject to the same constraints namely, 3.2.2 and 3.2.3 above.

$$\min \left\{ -\boldsymbol{E}'\boldsymbol{X} + \frac{\tau}{2}\boldsymbol{X}'\boldsymbol{W}\boldsymbol{X} + \frac{\lambda}{2}J \right\}$$

$$\Leftrightarrow \min \left\{ -\boldsymbol{E}'\boldsymbol{X} + \frac{\tau}{2}\boldsymbol{X}'\boldsymbol{W}\boldsymbol{X} + \frac{\lambda}{2}(\boldsymbol{X} - \boldsymbol{Z})'(\boldsymbol{X} - \boldsymbol{Z}) \right\} \tag{A.4.1}$$

$$\Leftrightarrow \min \left\{ -\boldsymbol{E}'\boldsymbol{X} + \frac{\tau}{2}\boldsymbol{X}'\boldsymbol{W}\boldsymbol{X} + \frac{\lambda}{2}(\boldsymbol{X}'\boldsymbol{X} - 2\boldsymbol{Z}'\boldsymbol{X} + \boldsymbol{Z}'\boldsymbol{Z}) \right\} \tag{A.4.2}$$

$$\Leftrightarrow \min \left\{ -(\boldsymbol{E}'\boldsymbol{X} + \lambda\boldsymbol{Z}'\boldsymbol{X}) + \frac{1}{2}(\tau\boldsymbol{X}'\boldsymbol{W}\boldsymbol{X} + \lambda\boldsymbol{X}'\boldsymbol{X}) + \frac{\lambda}{2}(\boldsymbol{Z}'\boldsymbol{Z}) \right\} \tag{A.4.3}$$

$$\Leftrightarrow \min \left\{ -(\boldsymbol{E}' + \lambda\boldsymbol{Z}')\boldsymbol{X} + \frac{1}{2}\boldsymbol{X}'(\tau\boldsymbol{W} + \lambda\boldsymbol{I})\boldsymbol{X} + \frac{\lambda}{2}(\boldsymbol{Z}'\boldsymbol{Z}) \right\} \tag{A.4.4}$$

Where $\boldsymbol{I} = \begin{pmatrix} 1 & 0 \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & 1 \end{pmatrix}$, $n\ x\ n$ identity matrix. The preference parameter, λ was in this context determined through experimentation, discussed in 3.2.2.2 below. It can also be set at the discretion of the decision-maker. When the aforementioned derivations and assumptions around expected returns and the MVP (3.2.4) were added, A.4.4 consequently reduced to -

$$\min \left\{ -\lambda\boldsymbol{Z}'\boldsymbol{X} + \frac{1}{2}\boldsymbol{X}'(\tau\boldsymbol{W} + \lambda\boldsymbol{I})\boldsymbol{X} + \frac{\lambda}{2}\boldsymbol{Z}'\boldsymbol{Z} \right\} \tag{A.4.5}$$

The importance of the new formulation is that it is in line with the standard MV-optimisation framework and efficient computation can be executed through existing quadratic optimisers.

Whilst it may be argued that the tilt could be achieved by simply replacing the expected return vector with factor returns, it is worth noting that in adding the $J\ Index$, you are seeking to minimise the distance between the resultant optimal portfolio weights and the weights formulated on the tilt (factor values), in addition to minimising risk. By incorporating and replacing the expected return vector with the factor returns you would be implicitly maximising the factor value (weight), which after transforming the initial objective function (A.4) is actually the outcome, but this would not wholly account for impact the $J\,Index$ has on covariance matrix, which is also affected in the formulation.

3.2.1.2 *Enhancement (Tilt) using Diversification (Herfindahl–Hirshman Index)*

King (2007) introduced diversification into the MV-optimisation framework, through the Herfindahl–Hirshman Index (HHI):

$$HHI = \boldsymbol{X}'\boldsymbol{X}$$

The indicator will always fall in a range between $\frac{1}{N}$ and one, where N denotes the number of stocks in the portfolio. Higher levels of HHI (closer to one) indicate high portfolio concentration, while lower levels (closer to $\frac{1}{N}$) are more representative of an equally weighted portfolio.

The incorporation of the HHI into the objective function resulted in -

$$\min\left\{-\boldsymbol{E}'\boldsymbol{X} + \frac{\tau}{2}\boldsymbol{X}'\boldsymbol{W}\boldsymbol{X} + \frac{\theta}{2}HHI\right\} \tag{B.1}$$

And can be transformed to an equivalent formulation which is derived below and subject to the same constraints, namely 3.2.2 and 3.2.3 above.

$$\min\left\{-\boldsymbol{E}'\boldsymbol{X} + \frac{\tau}{2}\boldsymbol{X}'\boldsymbol{W}\boldsymbol{X} + \frac{\theta}{2}HHI\right\}$$

$$\Leftrightarrow \min\left\{-\boldsymbol{E}'\boldsymbol{X} + \frac{\tau}{2}\boldsymbol{X}'\boldsymbol{W}\boldsymbol{X} + \frac{\theta}{2}\boldsymbol{X}'\boldsymbol{X}\right\} \tag{B.1.1}$$

$$\Leftrightarrow \min\left\{-\boldsymbol{E}'\boldsymbol{X} + \frac{1}{2}(\tau\boldsymbol{X}'\boldsymbol{W}\boldsymbol{X} + \theta\boldsymbol{X}'\boldsymbol{X})\right\} \tag{B.1.2}$$

$$\Leftrightarrow \min\left\{-\boldsymbol{E}'\boldsymbol{X} + \frac{1}{2}\boldsymbol{X}'(\tau\boldsymbol{W} + \theta\boldsymbol{I})\boldsymbol{X}\right\} \tag{B.1.3}$$

The preference parameter, θ, was, in this context, determined through experimentation and is discussed in 3.2.2.3. It can also be set at the discretion of the decision-maker. When the aforementioned derivations and assumptions around expected returns and the MVP (in equation 3.2.4) were added, B.1.3 consequently condensed to -

$$\min\left\{\frac{1}{2}\boldsymbol{X}'(\tau\boldsymbol{W}+\theta\boldsymbol{I})\boldsymbol{X}\right\} \tag{B.1.4}$$

3.2.1.3 Enhancement (Tilt) using Active Distance

Cremers and Petajisto (2009) introduced the concept of active share, which is the aggregate of all the absolute active bets within a portfolio. An active bet is defined as the deviation between the fund and benchmark weighting for each stock held in the portfolio. A related metric – *active distance*, which measures the Euclidean distance between the fund and benchmark - was introduced by Bradfield, Maritz and Swartz (2005). The latter metric given below is basically the *squared version* of active share.

$$AD=(\boldsymbol{X}-\boldsymbol{B})'(\boldsymbol{X}-\boldsymbol{B}) \tag{C.1}$$

Where $\boldsymbol{B}=\begin{bmatrix}b_1\\ b_2\\ \vdots\\ b_n\end{bmatrix}$, represents a vector of SWIX150 weightings.

AD ranges from zero for an index-tracking fund to one for a portfolio with no overlap with its benchmark. A higher AD is representative of a more *active* fund - where the deviation in the manager's stock weightings or stock selection is remarkably different from the benchmark. Lower ADs typify *closet index* funds.

The incorporation of the AD into the objective function resulted in -

$$\min\left\{-\boldsymbol{E}'\boldsymbol{X}+\frac{\tau}{2}\boldsymbol{X}'\boldsymbol{W}\boldsymbol{X}+\frac{\eta}{2}AD\right\} \tag{C.2}$$

This can be transformed to an equivalent formulation which is similar to that derived in A.4.1 – A.4.5 above and subject to the same set of constraints, namely 3.2.2 and 3.2.3 above.

$$\min\left\{-(\boldsymbol{E}'+\eta\boldsymbol{B}')\boldsymbol{X}+\frac{1}{2}\boldsymbol{X}'(\tau\boldsymbol{W}+\eta\boldsymbol{I})\boldsymbol{X}+\frac{\eta}{2}(\boldsymbol{B}'\boldsymbol{B})\right\} \tag{C.2.1}$$

The preference parameter, η, was derived through experimentation similar in process to the estimation of both λ and θ(refer to section 0). It can also be set at the discretion of the decision-maker. When the aforementioned derivations and assumptions around expected returns and the MVP (3.2.4) were added, C.2.1 consequently condensed to -

$$\min\left\{-\eta\boldsymbol{B}'\boldsymbol{X}+\frac{1}{2}\boldsymbol{X}'(\boldsymbol{W}+\eta\boldsymbol{I})\boldsymbol{X}+\frac{\eta}{2}(\boldsymbol{B}'\boldsymbol{B})\right\} \tag{C.2.2}$$

Two additional portfolios were formulated; the MVP and SWIX capitalisation weighted benchmark. In creating the MVP, which serves as a suitable point of comparison, λ, θ and η were set to zero. This eliminated the impact of all the aforementioned enhancements on the optimal solution and reduced the MV-optimisation framework to its standard format.

The second point of comparison used in this analysis was the SWIX benchmark, which was created using a methodology similar to that used by the FTSE-JSE, that is, each stock was rebalanced back to their SWIX free-float weight on a quarterly basis at the end of each of the following months; March, June, September and December.

The following section describes in detail the methodology used to determine the most optimal estimates of, λ , θ and η.

3.2.2 Parameter Estimation

3.2.2.1 Estimation of τ

Many scientific and heuristic methods have been used to estimate the level of τ. Grinold and Kahn (1999) used base estimates of return, E_P and risk, W_P over a long-term period to estimate τ as $\tau=\frac{E_P}{2W_P}$. τ is a representation of the slope that through minimum variance (MVP) investing, one generally seeks to minimise. And so return is in this context, is eliminated from the MV-formulation, since the author in her estimation of expected returns chose to assume a constant return across all stocks.

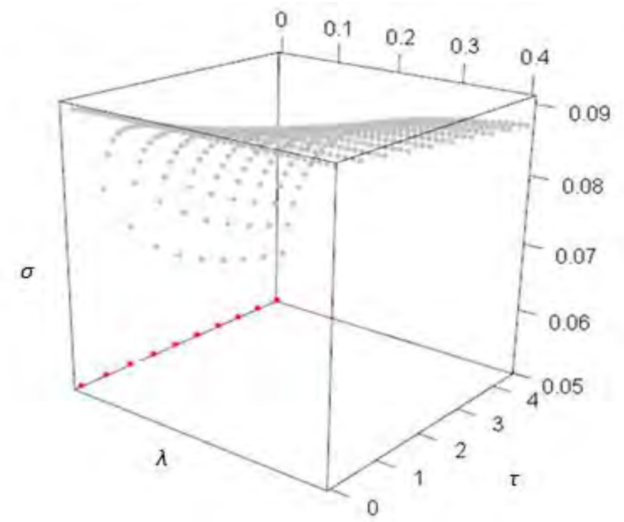

Figure 1 Risk measured at varying levels of λ and τ

In addition to this, if one considers the graphical implications (Figure 1) of using exactly a constant or zero return for each stock – it becomes clear that the efficient frontier becomes flat (at the return selected for each stock), indicated by the red dotted line in the figure above. No matter the value of τ, the MVP is always selected in the standard MV-optimisation framework (λ is set to 0). Thus in this context τ can be set to one.

3.2.2.2 Estimation of λ

Since τ was set to one, the next step was to solve for λ, the preference parameter used in A.4.5 above. A method of experimentation was employed to determine the best estimate of λ and is described in the section below.

Firstly, it was important to identify the investable universe; a *complete set* of stocks which the author defined as stocks that had sufficient data to calculate the weights of portfolio instruments based on their average volume indicator, z_i at each time point and were constituents of the SWIX 150 index. If any stock had not been trading (i.e. had no volume data for the prior 50-business day period) it was excluded from the investable universe. This universe was important to identify, since it was required in the estimation of both the sample covariance matrix, $\boldsymbol{S}$ and of the shrinkage estimator, $\boldsymbol{G}$, which is core to the MV-optimisation process -

$$\boldsymbol{G} = \omega \boldsymbol{F} + (1-\omega)\boldsymbol{S} \quad \text{(E.1)}$$

Where $\boldsymbol{F}$ the shrinkage target is determined from the constant correlation model and ω is an estimate of the shrinkage intensity and set at 0.4 (refer to sections 3.1 and 5.1 respectively). Both $\boldsymbol{S}$ and $\boldsymbol{G}$ were used to estimate the covariance matrix, $\boldsymbol{W}$ respectively. $\boldsymbol{W}$ along with $\boldsymbol{Z}$ were used in the estimation process and substituted into the enhanced objective function below.

The objective function in A.4.5 can be simplified as follows -

$$\min \left\{-\lambda \boldsymbol{Z}'\boldsymbol{X} + \frac{1}{2}\boldsymbol{X}'(\boldsymbol{W} + \lambda \boldsymbol{I})\boldsymbol{X}\right\} \quad \text{(E.2)}$$

$$\boldsymbol{s.t} \quad x_i \geq 0, i = 1, \dots, n \quad \text{(E.3)}$$

$$\boldsymbol{X}'\boldsymbol{1}_n = 1 \quad \text{(E.4)}$$

Since the $J\ Index$ (A.1) is a scalar term that is quadratic in $\boldsymbol{X}$. The $\frac{\lambda}{2}\boldsymbol{Z}'\boldsymbol{Z}$ term in A.4.5 can be ignored since it played no role in the actual optimisation step, because it is effectively a sum of constant terms. These constant terms fall away when the optimal solution is determined by taking partial derivatives with respect to $\boldsymbol{X}$. The first order conditions determine the optimal solution, so the $\frac{\lambda}{2}\boldsymbol{Z}'\boldsymbol{Z}$ term will have no bearing on the outcome. The only difference will be in the absolute value of the two objective functions (A.4.5 and E.2), which will differ by the magnitude of $\frac{\lambda}{2}\boldsymbol{Z}'\boldsymbol{Z}$.

A set of λ levels were selected and spanned a broad range from x_l to x_u, where x_l and x_u represent the lower and upper thresholds of λ. A range of step sizes, δ_s , where $s = 1, \dots n$, were used to create the sequence of λ estimates. The steps below were repeated for each λ level, by iterating through the sequence of λ estimates over the period April 2006 (t_0) to April 2016 (T)-

Table 1 Process Description of λ estimation

STEP 1	Let $\hat{\lambda} = x,\ where\ x_l \leq x \leq x_u.\ x$ is an estimate of λ. For month t, identify the investable universe of stocks that meet the stipulated criteria; the stock forms part of the SWIX 150 and has been trading on the JSE bourse for at least the preceding 50 business days.
STEP 2	Estimate the covariance matrix, $\boldsymbol{W}$, using the covariance shrinkage approach, $\boldsymbol{G}$ or sample covariance matrix, $\boldsymbol{S}$ described in 3.1 above, using three years of historical weekly return data for the investable universe of stocks determined in Step 1 above.
STEP 3	Calculate the $\boldsymbol{Z}$ vector using equations A.2.1 – A.3 above.
STEP 4	Using the above metrics calculated (in steps two to three) as inputs into the MV-optimisation framework given in E.2-E.4 above, solve for the optimal portfolio weightings at month t.
STEP 5	Hold the portfolio until the next month $(t+1)$, this effectively yields an out-of-sample return for month $t+1$.
STEP 6	Repeat steps one through to five for month $t+1$ (for x). This effectively results in the portfolio being rebalanced monthly, yielding an out-of-sample return for each sequential month until the end of the observation period, month T.
STEP 7	Repeat Step six for each sequential x, thus creating an individual portfolio ($portfolio(x, \boldsymbol{W})$) and record the relevant portfolio statistics and optimal weightings for each $portfolio(x, \boldsymbol{W})$.
STEP 8	The results from Step seven above were recorded and analysed in combination to determine the best estimate of λ. A series of absolute and relative statistics were examined in this estimation process, which are illustrated in the 5.2.1 below. Relative statistics were examined against the benchmark.

3.2.2.3 Estimation of θ

The process described in 3.2.2.2, above, was slightly altered to estimate θ. However, there are distinct areas where the process was modified. Firstly, the criteria used to determine the investible universe was simplified to select constituents of the SWIX 150 only. Additionally, since the objective function requires no estimate of $\boldsymbol{Z}$, the formulation was notably condensed as shown in F.1 below -

$$\min \left\{\frac{1}{2}\boldsymbol{X}'(\tau\boldsymbol{W}+\theta\boldsymbol{I})\boldsymbol{X}\right\} \quad \text{(F.1)}$$

$$\boldsymbol{s.t} \quad x_i \geq 0, i = 1, \dots, n \quad \text{(F.2)}$$

$$\boldsymbol{X}'\boldsymbol{1}_n = 1 \quad \text{(F.3)}$$

The identical procedure was used to determine the estimates of the covariance matrix, $\boldsymbol{W}$- specifically $\boldsymbol{S}$ and $\boldsymbol{G}$ were calculated and substituted independently into the enhanced objective function above (F.1). No modifications were made to the objective function to determine the optimal solution.

The procedure regarding the creation of λ levels described in 3.2.2.2, above, was used to determine the sequence of θ levels - the author established an upper bound (y_u) and lower bound (y_l) over which a broad sequence of θ estimates were set. The steps below were repeated for each θ estimate, by iterating through the sequence of θ estimates over the period April 2006 (t_0) to April 2016 (T):

Table 2 Process Description of θ estimation

STEP 1	Let $\hat{\theta} = y,\ where\ y_l \leq y \leq y_u$. y is an estimate of θ. For month t, identify the investable universe of stocks that meet the stipulated criterion; the stock forms part of the SWIX 150.
STEP 2	Calculate the estimate of the covariance matrix, $\boldsymbol{W}$, using the covariance shrinkage estimate, $\boldsymbol{G}$, or sample covariance matrix, $\boldsymbol{S}$, described in 3.1 above, using three years of historical weekly data for the investable universe of stocks determined in step one above.

STEP 3	Calculate the optimal solution by substituting the outputs from step one and two above into the MV-optimisation framework given in F.1-F.3 above, for month t.
STEP 4	Hold the portfolio until the next month $(t+1)$, this effectively yields an out-of-sample return for month $t+1$.
STEP 5	Repeat steps one through to five for month $t+1$ (for y). This results in the portfolio being rebalanced monthly yielding an out-of-sample return for each sequential month until the end of the observation period, month T.
STEP 6	Repeat Step five for each sequential y, thus creating an individual portfolio ($portfolio(y, \boldsymbol{W})$) and record the relevant portfolio statistics and optimal weightings for each $portfolio(y, \boldsymbol{W})$.
STEP 7	The results from Step six above were recorded and analysed in combination to determine the best estimate of θ. A series of absolute and relative statistics were examined in this estimation process, which are illustrated in the 5.2.2 below. Relative statistics were examined against the benchmark.

3.2.2.4 Estimation of η

Again, this process was very similar to that explained in sections 3.2.2.2 and thus, 3.2.2.3 above. The primary difference was that in the estimation of η, the optimal solution was sought around the $\boldsymbol{B}$ vector, not the $\boldsymbol{Z}$ vector. Also, there was no additional filter used to screen for stocks that had been trading for a 50-day period. This criterion was ignored and the investible universe was only reliant on whether the stock was a constituent of the SWIX 150. The method described in 3.2.2.2 can be followed to determine the best estimate of η.

This chapter covered in detail the methodologies used in the estimation of the following; risk inputs determined through covariance shrinkage, as well as the back-testing methodology employed to determine the preference parameter estimates. In addition, the derivation of the enhanced objective functions was described.

4. Data

This chapter provides detail of the data used in; the covariance shrinkage estimation process and the back-testing methodology implemented to determine the most appropriate level of factor tilts (parameter estimation).

4.1 Covariance Shrinkage estimation

In order to determine both estimates for the shrinkage target and intensity, the author used weekly total returns dating back to May 2003. The author's dataset ended in April 2016. To avoid introducing a bias due to the asymmetry of the return distribution, the author transformed to log returns and converted back to simple geometric returns at the end. The following procedures were carried out to ensure that the data had integrity and that all calculations were based on accurate and sensible data:

1. For a share without enough history: the data was back-filled using the underlying sector return. The lowest available sub-sector return was used as a proxy; however, if it was not available, the subsequent sector return could be used. The same sequence was followed in that if the sector return was not available, then the industry and finally the major sector return was used. The issue in some instances, however, was that sectors may have changed over time due to data availability for the indices. This process of back-filling data ensured that the exact number of data points were available when the covariance matrix is calculated.
2. If the return was more than a certain number of standard deviations away from the median return, it was trimmed, that is, it was reduced to a maximum of three median absolute deviations (MAD) away from the median. This ensured that the data was more robust to outliers. Using the median and MAD is suitable, since they are asymptotically equal to the mean and standard deviation for a normal distribution and are more robust to outliers when the assumption of normality does not hold.
3. This screening was applied to all returns per given time period.

The universe of stocks was determined by selecting the largest 150 companies from the FTSE-JSE SWIX index. This index represents the same constituents as the JSE All share, but represents the proportion of a constituent's market capitalisation that is dematerialised and held on the South African share register only and therefore down-weights dual-listed companies[3].

In order to determine a good estimate of the shrinkage intensity, it is important to calculate the parameter over multiple time periods. The sample covariance matrix was thus, estimated using weekly returns for the top 150 stocks over rolling three-year window periods starting in May 2003 and ending in April 2016. Using weekly returns over a three-year period ensured that the number of time points (T) was always greater than the number of stocks (N). Additionally, Munro and Bradfield (2016) showed that shorter estimation periods increased the sampling error and noise as well as increased the shrinkage intensity significantly. The universe of stocks was determined at week t and the sample covariance matrix, $\boldsymbol{S}$, was estimated over the preceding 156-week period [$(t–155)$, t]. Similarly other parameters; $\bar{r}$, $\boldsymbol{F}$, $\hat{\pi}$, $\hat{\rho}$, $\hat{\gamma}$ and $\hat{\kappa}$ required in determining the shrinkage intensity, ω were calculated on the same basis.

Whilst it may be argued that replacing missing share return data with the relevant sector data would tend to increase diversification and lower overall portfolio volatility, it is worth noting that missing returns were only substituted in determining the weekly returns which were used to estimate the sample covariance matrix. These returns were backfilled to prevent a sample covariance matrix which would more likely be ill-conditioned (this method usually leads to a covariance matrix with negative eigenvalues which is something one wants to avoid). This, together with the requirement for 36 data points for 150 shares, would leave the covariance matrix with a significant number of missing data points.

This method could lower portfolio volatility, but both the MVP and enhanced (tilted) portfolios were formulated using the same covariance matrix, so this resulted in robust and unbiased comparisons across all portfolios. Also the way in which returns were backfilled attempted to use the most granular sector level data before using the subsequent upper sector returns; the process used more granular data unless it was not available. This should decrease the likelihood of obtaining even lower

[3] When a company's securities are listed on more than one exchange, this increases liquidity and gives investors flexibility around where they trade shares

volatilities (if only the major (top level) sector return was employed). This is an acceptable method in proxying returns and is used in the investment industry.

Whilst it also may be argued that winsorising extreme returns removes the effect of volatility that the optimisation process would minimise, it is worth noting that although extreme or influential observations are actual data points and are not always considered to be problematic, they may have an effect on statistical inferences. This was an observation made by Belsley, Kuh and Welsch (1980). They can be a result of erroneous data or unusual events. Erroneous data can be as a result of incorrect measurement, for example see Kraft (2006), who considers the company Smith Corona, which filed for bankruptcy and delisted in May 1996 when the stock closed at $0.375. After reorganisation, its new shares started trading, but in 1997 the incorrect delisting price of $3.12 was used and this resulted in a 700% delisting return instead of -100%. In the instance of unusual events, Kraft (2006) also describes how in 1998 Triton Energy Ltd. had 203 percent buy and hold abnormal return. There are conflicting views on dealing with influential observations, in fact Leone, Minutti-Meza, Wasley (2015) document that in reviewing 157 studies using stock return data, 53% of studies winsorised or truncated extreme stock returns, versus 47% who used data in the raw format. So the literature does not necessarily reject the idea of windsorising data to create more robust statistical inferences.

4.2 Estimation of preference parameters

The following section describes the data sources required for the estimation of λ, θ and η, which are parameters that were required in the minimisation of the following objective functions. These objective functions are repeated below for ease of reference:

i. Enhancement through the average volume indicator

$$\min\left\{-\lambda \boldsymbol{Z}'\boldsymbol{X} + \frac{1}{2}\boldsymbol{X}'(\boldsymbol{W} + \lambda \boldsymbol{I})\boldsymbol{X}\right\} \tag{4.2.1}$$

ii. Enhancement through diversification (Herfindahl–Hirshman Index)

$$\min\left\{\frac{1}{2}\boldsymbol{X}'(\tau \boldsymbol{W} + \theta \boldsymbol{I})\boldsymbol{X}\right\} \tag{4.2.2}$$

iii. Enhancement through Active Distance

$$\min\left\{-\eta\boldsymbol{B}'\boldsymbol{X}+\frac{1}{2}\boldsymbol{X}'(\boldsymbol{W}+\eta\boldsymbol{I})\boldsymbol{X}\right\} \tag{4.2.3}$$

Both weekly and monthly returns were required for each of the aforementioned parameter estimations. The procedures carried out to ensure the return data had integrity and that all calculations were based on accurate and sensible data are described at length in 4.1 above. The universe of stocks was determined by selecting the largest 150 companies from FTSE-JSE SWIX index. Both the sample covariance matrix, $\boldsymbol{S}$, and the shrinkage estimate, $\boldsymbol{G}$, were estimated using weekly returns. The back-tests were carried out using monthly returns.

In estimating λ and η however, additional data sources were required, these are summarised in the paragraph below:

1. The process of λ estimation required the collation of two additional data sources to calculate the average daily volume indicator required in the calculation of the $\boldsymbol{Z}$ vector. Daily volume and free float shares in issue were collected for around 290 counters dating back to January 2006. These stocks would have formed part of the SWIX 150 at some point in time historically.
2. η estimation required benchmark weightings, which were based on the SWIX market capitalisation weighted index.
3. There were no additional data sources required in the estimation of θ.

This chapter provided detail of the data used in; the covariance shrinkage estimation process and the back-testing methodology implemented to determine the most appropriate level of factor tilts (parameter estimation).

5. Empirical Results

This section provides an in depth analysis of the results from the shrinkage parameter estimation process. Furthermore, it outlines the out-of-sample performance of the enhanced portfolios formulated using the two covariance structures. This informs the parameter estimation process or level at which factor tilts could best be incorporated to improve performance relative to the .

5.1 Covariance Shrinkage estimation

The shrinkage intensity estimate, ω, which is based on the constant correlation model, is fairly stable over the measurement period as shown in Figure 2 below.

Figure 2 Shrinkage intensity calculated over rolling three-year periods, for the period 2006-2016

The shrinkage intensity is calculated using formulae (3.1.1) – (3.1.5) described in section 3.1 above. There is a clear shift in the average shrinkage intensity level around October 2011. Over the period, May 2006 to September 2011, the average value of the shrinkage intensity is approximately 0.48. This indicates that the shrinkage target, $\boldsymbol{F}$ will be allocated around a 48% weight over this period. Over the entire period, the overall level reduces to 0.40, since in the latter half of the period (Oct 2011 to April 2016) there is a significant decline in the shrinkage intensity to an average level of 0.32. Since the universe is determined at time t, and the parameters estimated on the previous three-year period of weekly data, the drop coincides with and would include the financial crisis that occurred in the latter half of 2008. One would expect a decline in ω in as a consequence of a market crash. Volatilities and correlations tend to increase significantly during market crashes like the 2008 financial crisis, when markets trend however, they tend to decline.

In *Table 3* below, some descriptive statistics of the shrinkage intensity are reported. The standard deviation would be far less if measured over the separate periods as defined above. This holds true for the first moment as well. On closer inspection of the underlying parameters; $\hat{\pi}$, $\hat{\rho}$, $\hat{\gamma}$, all tend be highly correlated with each other.

Table 3 Shrinkage intensity (ω) - Descriptive Statistics

Min.	1st Qu.	Median	Mean	3rd Qu.	Max.	Std dev.
0.289	0.341	0.456	0.413	0.475	0.546	0.074

5.2 Parameter Estimation

Both $\boldsymbol{G}$ (shrinkage estimator) and $\boldsymbol{S}$ (sample covariance) were employed in the estimation process. In the sections, below, however the author outlines in detail the results of the experimentation which only incorporated $\boldsymbol{G}$. A condensed summation of results obtained using both $\boldsymbol{G}$ and $\boldsymbol{S}$ are provided in the sections; 5.2.1.3, 5.2.2.3 and 5.2.3.3.

5.2.1 Estimation of λ (Average Volume Indicator Tilt)

A range of metrics were calculated to determine the most optimal levels of λ. It was imperative to assess the performance of each portfolio over the entire period with respect to both absolute and relative measures, including return, risk, and risk-adjusted returns.

5.2.1.1 Broad λ Range

In initial tests, the estimation interval was set over a broader range, with a greater step size. Since the λ level would be determined from experimentation, the author chose to select values within a range from 0 to 0.5. The outcome of the experimentation would inform any subsequent steps or improvements in the process that would be required.

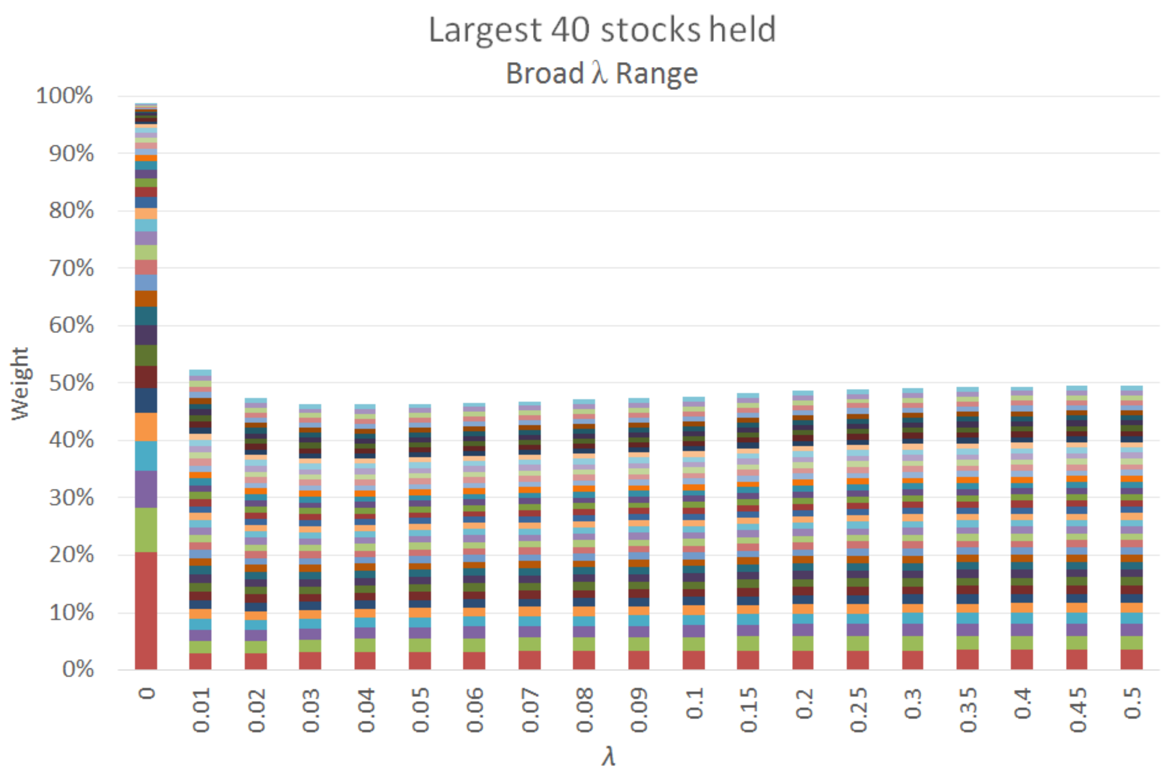

Figure 3 Top 40 portfolio holdings held for broad λ range

Figure 3 above depicts the average weights of the largest 40 stocks; this was determined by ranking stock weightings per month from largest to smallest and then calculating the average weight over time per ranking. It is evident that as λ increases beyond 0.15, the portfolio quickly approaches the *volume-weighted* portfolio, that is, the portfolio based on their average volume indicator weighting, z_i. As $\lambda \rightarrow \infty$, the impact the risk parameter ($\boldsymbol{W}$ which was estimated with the shrinkage estimator, $\boldsymbol{G}$) has on the optimal solution is minimised. The minimum variance portfolio (MVP) is generated when λ is set to zero. For $\lambda = 0$, the random portfolio is highly concentrated and holds in excess of 20% in the largest stock on average over time. The next largest stock is approximately eight percent. From an investible universe of 150 stocks, the MVP holds on average 20% holding in the largest stock and the next 39 stocks account for around 78% of the portfolio. In comparison, the introduction of the enhancement through the average volume indicator ($\lambda > 0$), adds a level of diversification to the portfolio. There was a substantial drop in the aggregate weighting of the largest 40 holdings in the portfolio ($\lambda > 0$) in comparison to the MVP which was fully invested at this stage. On average the largest 40 holdings account for around 52% of the total portfolio.

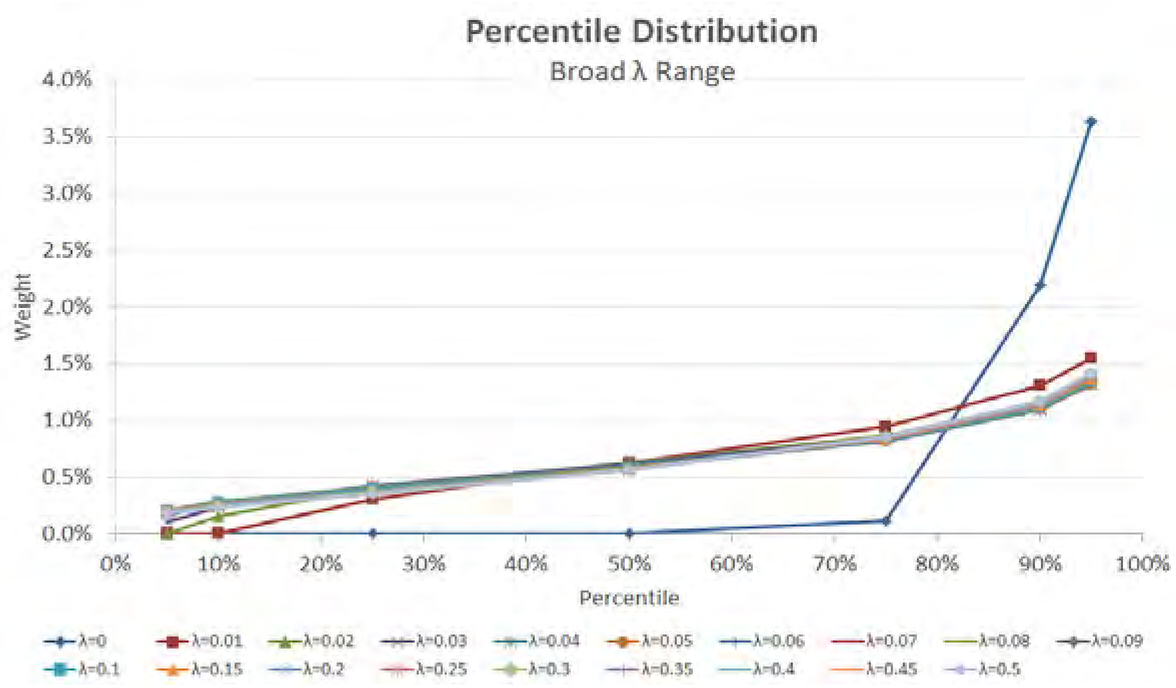

Figure 4 Percentile distribution for multiple values of λ

The above theme is mirrored in the percentile distribution in Figure 4 above, where the percentile weights are calculated in a similar way to that described above, are calculated over the entire period. Two interesting observations can be made:

i. The MVP never holds the complete investible universe at any particular point. It assigns a zero weighting to more than 50% of the portfolio. Only five percent of the portfolio holds a weight in excess of three-and-a half percent. The average maximum weight held over the measurement period is around 20%.

ii. In comparison, the volume-weighted portfolios have exposure to around 90% of the investible universe over time. Around half the stocks have a weighting less than 0.5%. The 90th-percentile is approximately 1.15%. As λ increases above 0.15, the portfolios seem to approach the volume-weighted portfolio. The most benefit of introducing the $\boldsymbol{Z}$ vector is evident in the lower percentiles, where average exposure to stocks is increased.

iii. A key aspect which required further analysis, was the volatility of weights in each portfolio. This was completed by calculating the standard deviation of weightings for each stock over the complete period and then determining the average of these values across all stocks. Table 4 below, records these results. It is interesting to note that the MVP has a standard deviation around five times that of the average λ portfolio (where $\lambda > 0$). This could be indicative of higher levels of turnover and thus trading costs. Conversely the lower volatility observed in other portfolios (where $\lambda > 0$) could suggest *strategies akin to buy-and-hold*.

Table 4 Standard deviation of weights measured over the period (April 2006 - April 2016)

λ=0	λ=0.01	λ=0.02	λ=0.03	λ=0.04	λ=0.15
1.00%	0.32%	0.29%	0.28%	0.27%	0.26%

Figure 5 Cumulative Return - Broad λ range

An important aspect that needed consideration was the performance of the λ portfolios. It was clear that for portfolios with a $\lambda > 0.15$, there was convergence to the volume-weighted portfolio. There was no significant change in return as λ increased and approached 0.5.

The cumulative returns of a subset of λ portfolios are illustrated in Figure 5 above. These portfolio returns are plotted along with both the MVP and SWIX 150 (benchmark) portfolios. The plot includes λ levels that were significantly different from each other in terms of their return profiles. The selection was based on the sequential change in the return as the λ level increased. If the change in the monthly mean return was greater than 0.05%, the portfolio was selected. If this was not true however, the next permissible portfolio would be required to have a cumulative difference (from the last selected portfolio) of no less than 0.05%. The benchmark and MVP are the most distinctly different portfolios in comparison to the other portfolios. From inception to around April 2012, both the MVP and benchmark underperformed the λ portfolios. The benchmark does underperform for a considerable amount of time. The MVP outperformed significantly from around May 2012 and was the best performing portfolio for the entire period – as measured at the end of April 2016. There is very little that differentiates the λ portfolios from each other, especially in the initial period from April 2006 to April 2012. Post this period however, there is a fair amount of dispersion in returns, and portfolios with higher λ levels seemed to be quite heavily impacted and experienced larger drawdowns in the latter part of 2015.

Figure 6 Cumulative Relative returns vs. SWIX 150 benchmark

The cumulative relative returns are plotted in the Figure 6 above. This plot highlights the differences between the portfolios more clearly. The MVP underperformed in the period prior to the 2008 financial crisis, this is expected since in the lead up to the crash, exposure to high beta stocks (with higher risk) in the MVP would be very limited. These stocks benefitted from the strong bull run in commodities and exhibited inherent momentum and would thus be favoured in portfolios that are tilted to stocks which had higher average volume indicators, *that is*, higher z_i weightings. The MVP gained ground post the crisis, but underperformed extensively again in the beginning of 2009 before it rebounded.

The λ portfolios also suffered drawdowns over the aforementioned periods; 2008 and 2009, however the relative performance versus the MVP is less profound. Post these periods, the relative performance of the MVP was better than all λ portfolios; in fact it returned the highest relative return over the whole measurement period. Portfolios with a higher λ level started to underperform significantly in the latter half of the measurement period, delivering a negative relative return over the period 2012-2015. The performance in the λ portfolios and MVP diverged considerably in the periods May 2006 – June 2007 and April 2015 – April 2016. These were periods where serious sector dislocations occured. In the latter period, the resource sector (high risk) underperformed while consumer defensive stocks outperformed - these are low risk stocks and so would have a higher weighting and be favoured in a low risk portoflio like the MVP.

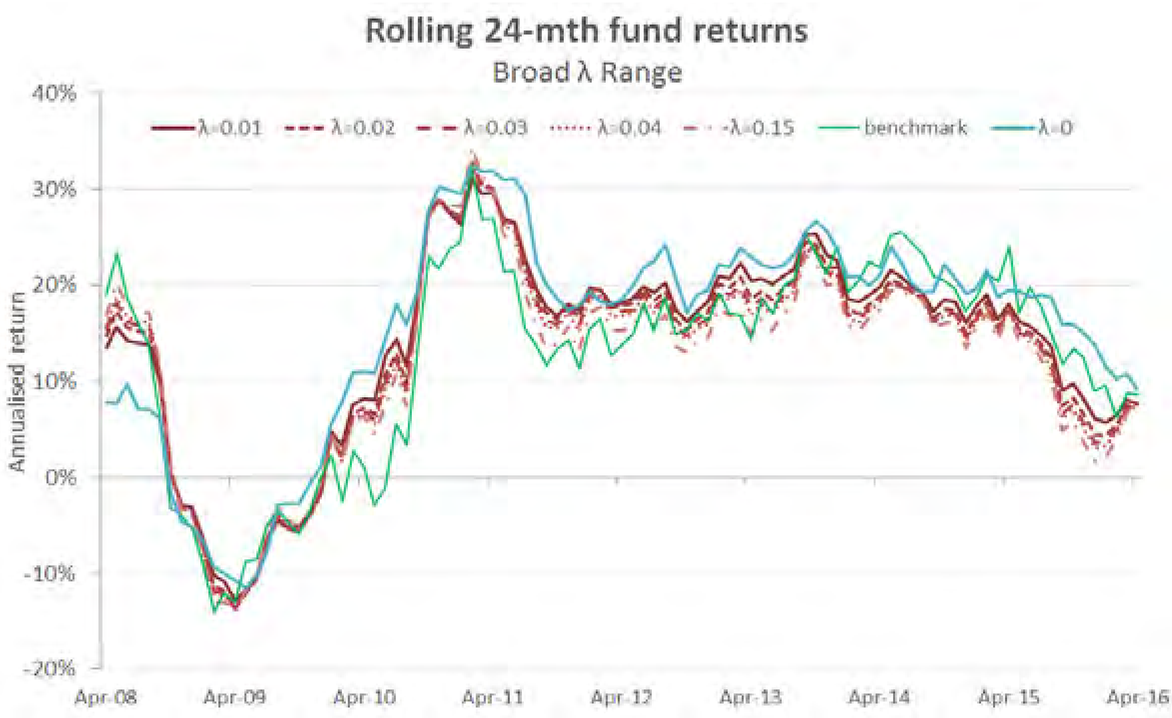

Figure 7 Rolling 24-month returns – Broad λ range

The rolling 24-month returns in Figure 7 above indicate that the MVP performance was superior over most of the period- it only underperformed all portolios for one year from inception. The λ portfolios on the other hand, underperformed the benchmark from around August 2013 until the end of April 2016. Another aspect which was analysed was the realised risk of the portfolios.

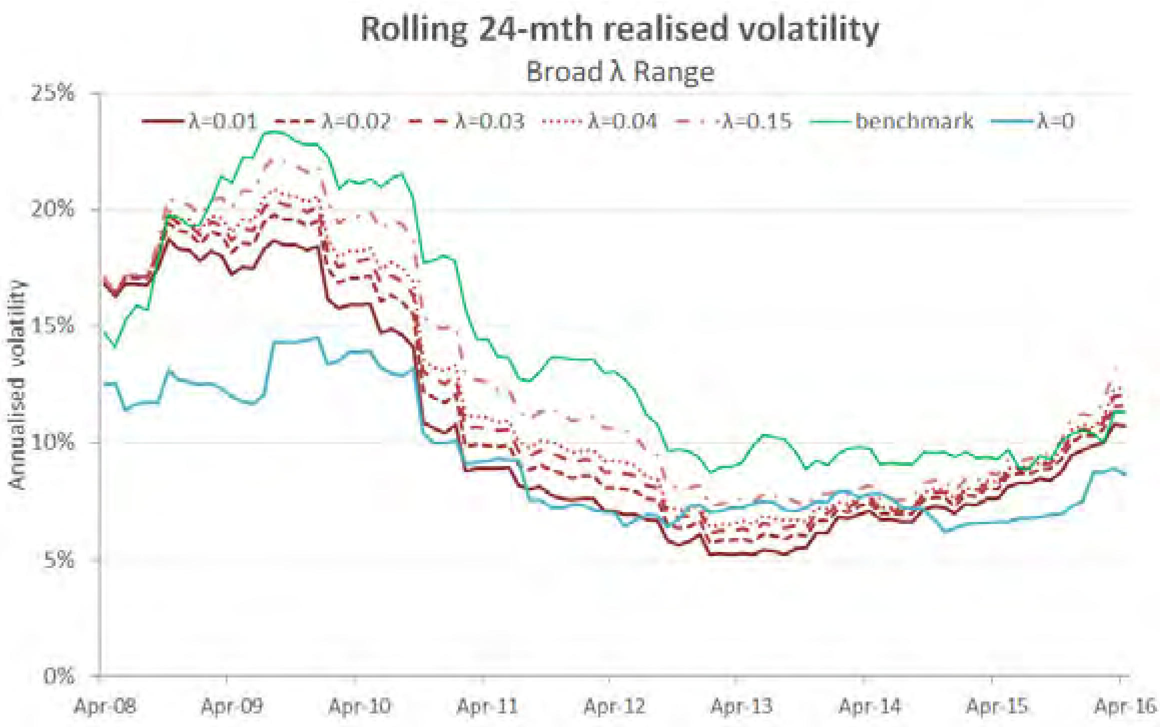

Figure 8 Rolling 24-month annualised risk– Broad λ range

From Figure 8 above, it is evident that the benchmark has generally been riskier than the other portfolios. As expected, the MVP had the lowest risk over most rolling measurement periods. For portfolios with $0.01 \leq \lambda \leq 0.04$, over the period June 2012 to June 2014, the realised volatility was lower than that of the MVP.

The period represented a strong bull market, and since markets trend higher in bull markets, realised volatilities tend to be lower. The λ portfolios (where $0.01 \leq \lambda \leq 0.04$), seem to have benefited most from this and delivered the lowest overall risk over the period. For portfolios with $0.01 \leq \lambda \leq 0.04$, over the period June 2012 to June 2014, the realised volatility was lower than that of the MVP.

The tracking error is shown in Figure 9 below; the MVP was inherently riskier on a relative basis in comparison to the other portfolios. The risk measured over the entire period was approximately 1.5 times higher than the λ portfolios.

Interestingly at a λ level of 0.01, there is a considerable reduction in the level of risk, which may indicate the optimal solution's sensitivity to small changes in the level of λ.

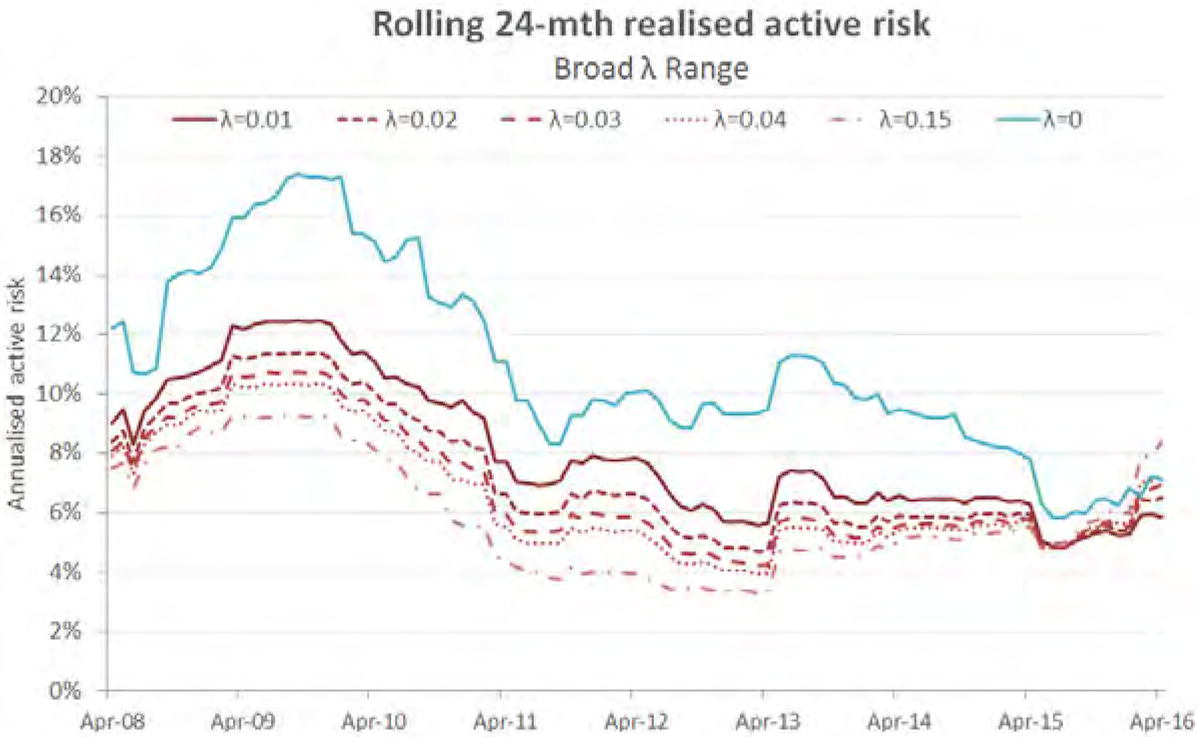

Figure 9 Rolling 24-month annualised active risk– Broad λ range

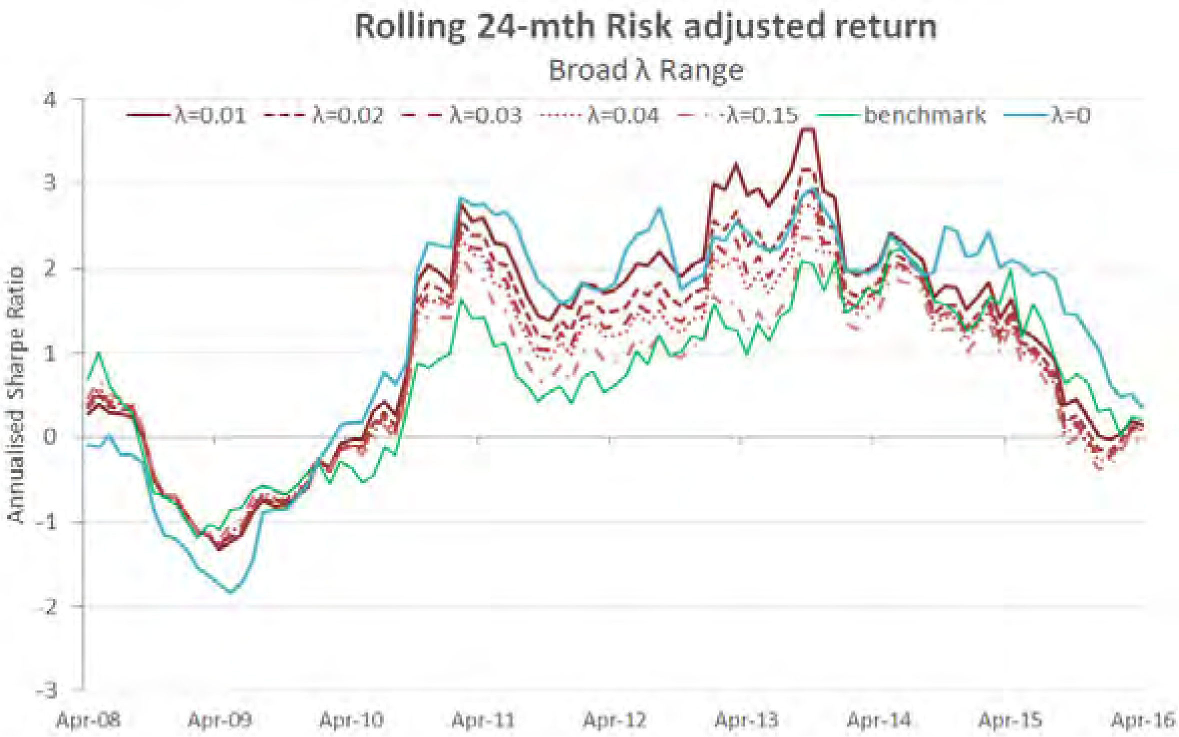

Figure 10 Rolling 24-month annualised risk-adjusted returns – Broad λ range

One can draw similar conclusions from the rolling risk adjusted return (Sharpe ratios) in Figure 10 above. The MVP benefitted from lowest volatility over the entire period and although the Sharpe ratio was affected in the initial months (April 2008 to June 2009); it was the only portfolio that was increasing in volatility whilst the others which were at higher levels, seemed to oscillate around a constant. The benchmark was clearly the most volatile and this resulted in the lowest Sharpe ratio overall. It seems that the volatility is a bigger differentiator then the actual return across both the benchmark and the funds.

Given the above results of the initial broad λ range, the author believed that there was merit in narrowing the range over which to search for the most optimal level of λ. The results are discussed in the section below.

5.2.1.2 Narrow λ Range

In determining the most optimal level of λ, one needs to distinguish from the range of λ's (where $\lambda > 0$), the portfolio that exhibits the following properties (ranked in order of importance):

1. Highest return
2. Highest risk adjusted return
3. Lowest risk

From initial tests described in the section 5.2.1.1, the portfolio identified with the above properties had a λ value of 0.01. Given this outcome, the λ range was narrowed and the identical absolute and relative measures described above, were recalculated on the narrower range. Narrowing the λ range did not guarantee finding a portfolio that would deliver superior performance compared to the portfolio determined at a λ level of 0.01 (if measurement was carried out on the three criteria listed above). This portfolio formed the base off which all comparisons were made. The MVP and benchmark were included since they were important for comparative analysis and formed the basis required for the relative performance metric calculation.

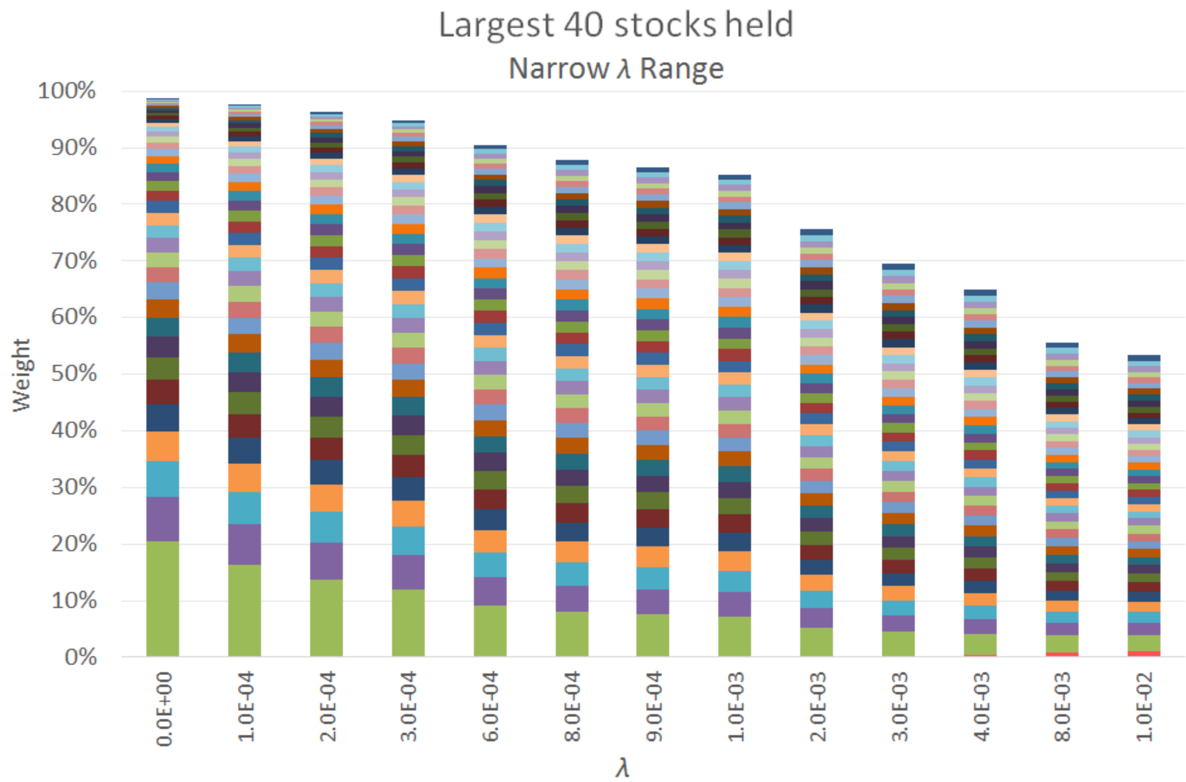

Figure 11 Largest 40 constituents held in the narrow λ range

The Top 40 holdings of the narrower range are displayed in Figure 11 above and were determined on the same principles discussed in 5.2.1.1. An aspect that needed to be examined was if there was any inherent change in the composition of the portfolios as the λ level was varied. The subset of portfolios selected was representative of the λ range and were determined on the same basis described in the preceding section 5.2.1.1. Again, the sensitivity of λ is quite apparent, a small change in the value results in a significant change in the portfolio. There is a substantial change in the distribution of weights and increase in diversification as λ increases.

In comparison to the previous analysis, where the Top 40 holdings accounted for approximately 52% of the portfolio on average over time (refer to Figure 3 in 5.2.1.1), this average percentage increased to 95% at lower levels of λ ($\leq$ 6.0E-4). When λ increased beyond 1.0E-3, this percentage dropped to 52% again. The distribution of weights in all λ portfolios were far more diversified in comparison to the MVP, which held 21% on average in the largest stock over time. This is shown in Figure 11 above.

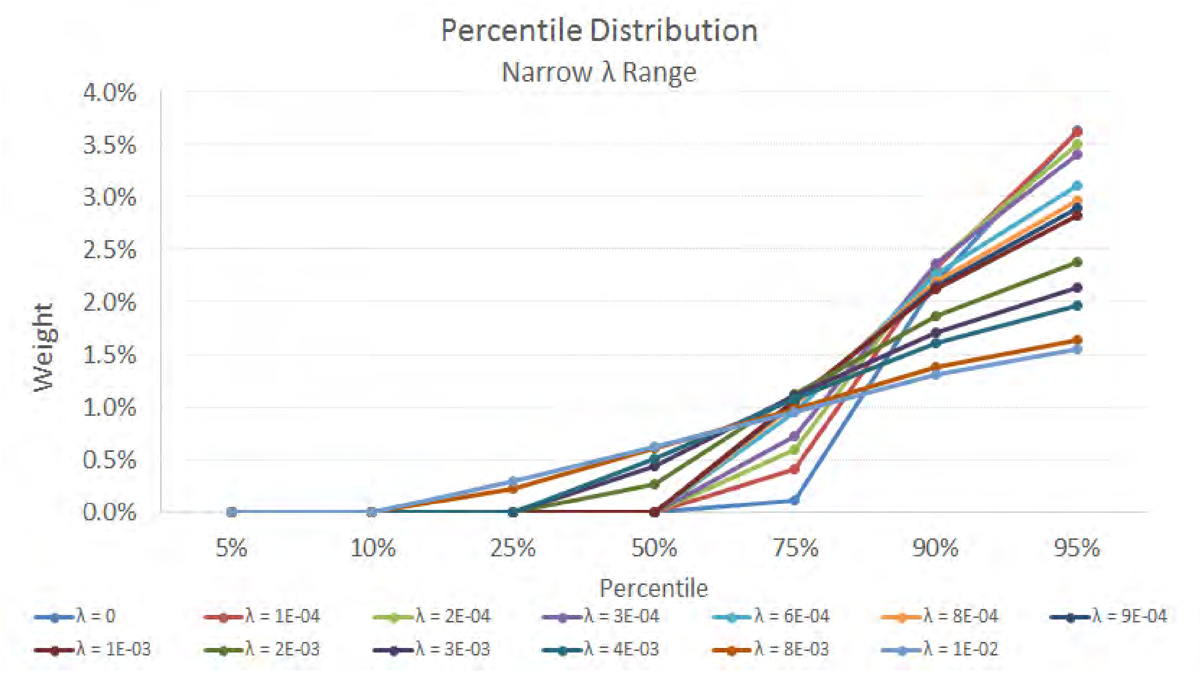

Figure 12 Percentile Distribution of Narrow λ Range

Figure 12 above indicates that at lower levels of λ ($\lambda \leq$ 1.0E-3), around 50% of the portfolio was assigned a zero weight. Between the 50th and 75th percentiles, λ portfolios became more distinctive, allocating more significant weightings to portfolio constituents. There was uniform increase in the portfolio weightings for λ levels in excess of 1.0E-3. From Table 5 below it is clear that the relationship between the standard deviation of weights, σ_w and λ still holds; so as λ increases, σ_w decreases. This was observed over the broader λ range as well.

Table 5 Standard deviation of weights measured over the period (April 2006 - April 2016)

λ	0	1.0E-4	2.0E-4	3.0E-4	6.0E-4	8.0E-4	9.0E-4	1.0E-3	2.0E-3	3.0E-3	4.0E-3	8.0E-3	1.0E-2
$\sigma_{w\,(\%)}$	1	0.88	0.77	0.72	0.61	0.55	0.53	0.51	0.43	0.39	0.37	0.33	0.32

The cumulative returns shown in Figure 13 below, can distinctly be grouped into two categories; portfolios with $\lambda \leq 0.001$ and $\lambda \geq 0.001$. As the λ level gets finer, the portfolios started to outperform the benchmark and MVP. Increasing λ reduced the performance gap however and as λ approached 0.01, the portfolios started to underperform the MVP toward the end of the period. In the beginning of period, the performance reversed and one can note the portfolios with higher λ levels outperforming.

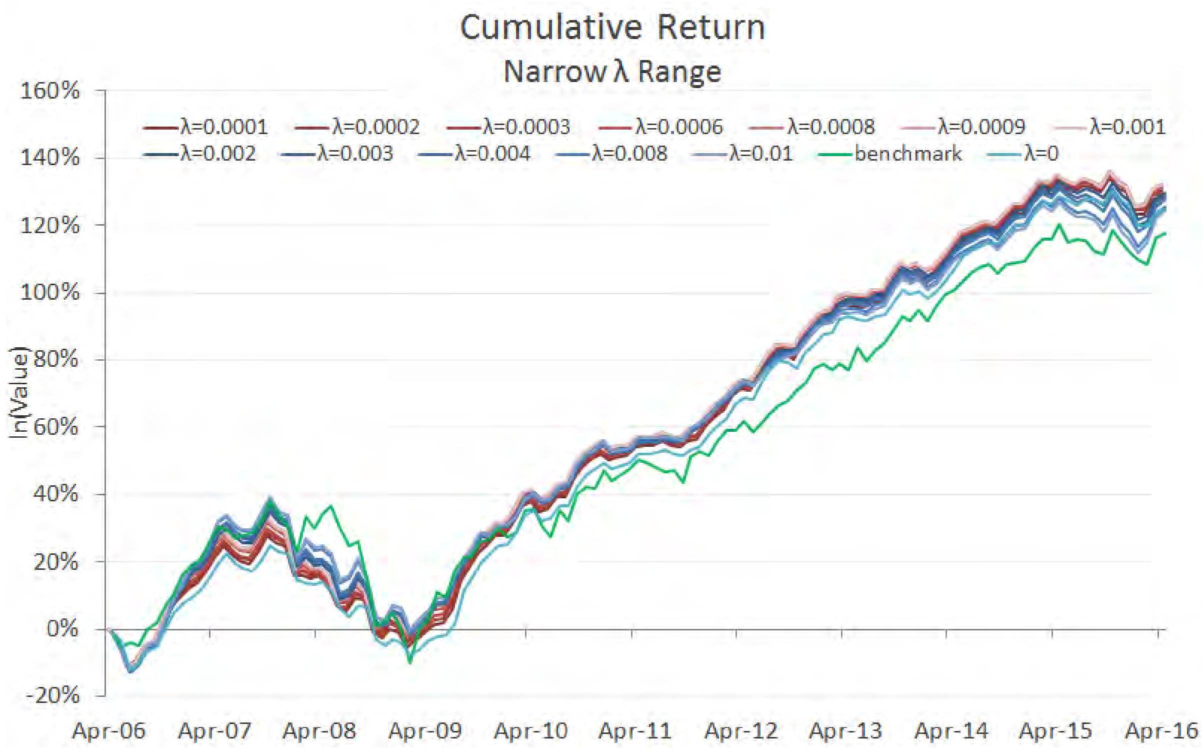

Figure 13 Cumulative Returns for Narrow λ Range

In the Figure 14 below, the portfolio generated at a λ level of 0.01 was set as the base portfolio off which all relative performances for all other portfolios were determined.

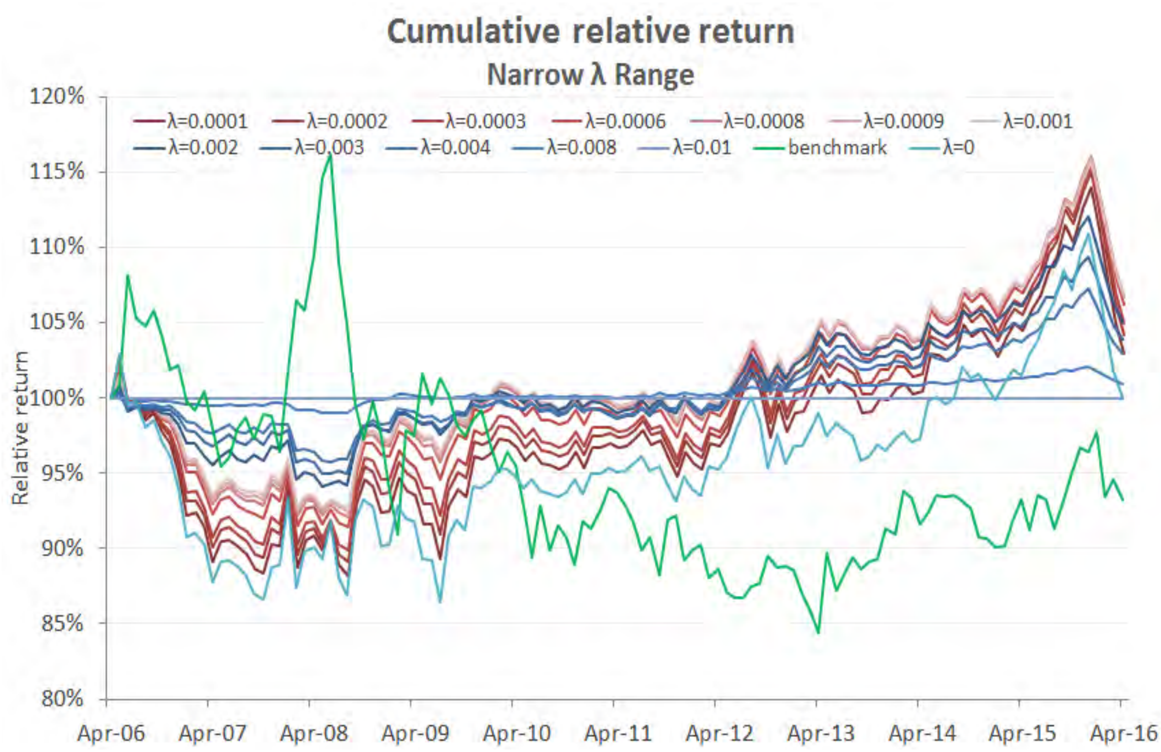

Figure 14 Relative return vs. portfolio (where $\lambda = 0.01$) – Narrow λ range

Prior to June 2009, the portfolios underperformed, however the differences in the performances are spread, in comparison to the period from June-2009 to April 2012 where the performance is flat and similar across most portfolios.

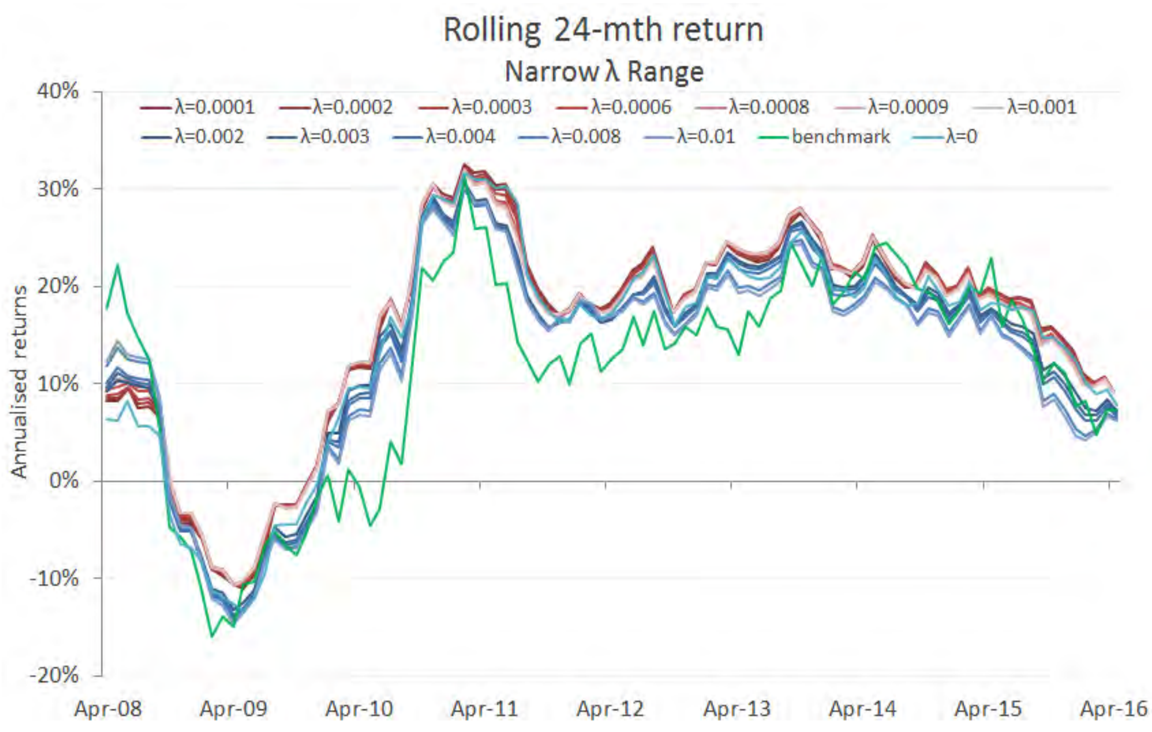

Figure 15 Rolling 24-month annualised returns – Narrow λ range

From Figure 15 above, the rolling performance is more distinct, in that prior to 2010; the λ portfolios seemed to perform differently from the benchmark. The finer λ portfolios ($\lambda \leq 0.001$) benefitted from having exposure to both the z_i weighting tilt and exposure to the MVP and consistently delivered positive performance over the majority of the period.

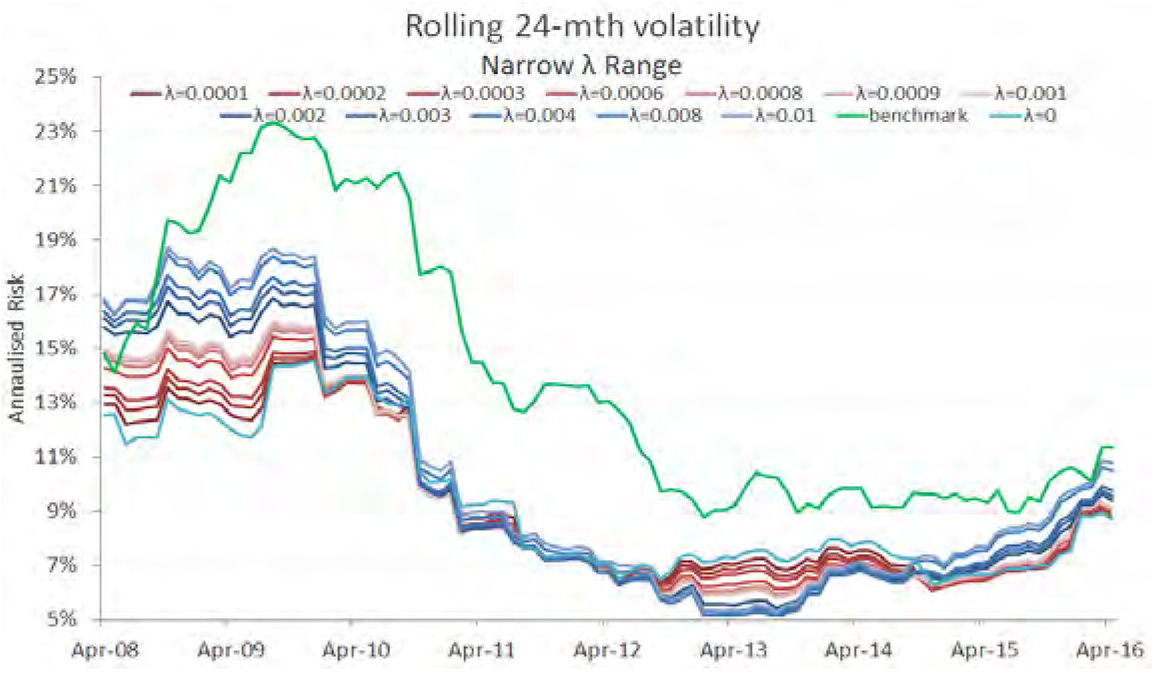

Figure 16 Rolling 24-month annualised risk – Narrow λ range

From Figure 16, one can note a similar pattern, in that as the λ level increases, there is an increase in risk in similar periods to that observed in broad λ range analysis and a decrease in risk in others. As λ approaches zero; the influence of the risk parameter increases and this results in the convergence of λ portfolios to the MVP.

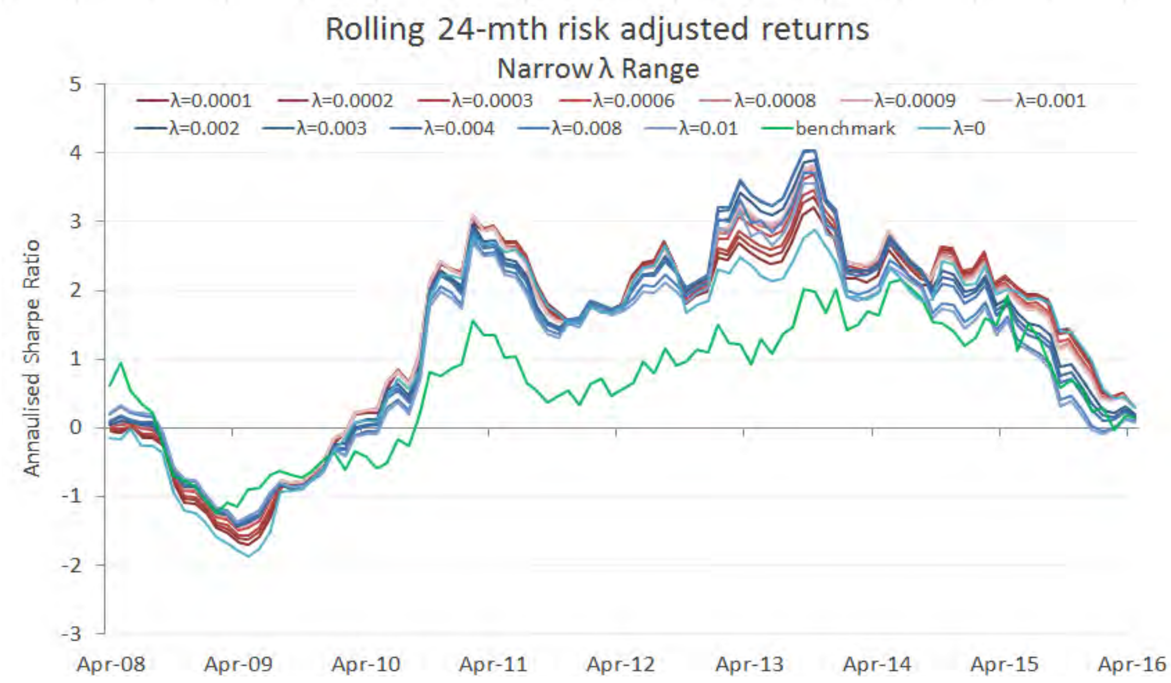

Figure 17 Rolling 24-mth annualised risk adjusted returns – Narrow λ range

In Figure 17, the convergence is more defined; however it is important to note that there is benefit in holding a portfolio with a level of exposure to the z_i weighting tilt.

In general the rolling performance is more distinct, in that prior to 2008, the λ tilted portfolios seemed to perform differently from both the MVP and benchmark. But during the crash of 2008, the portfolios faired much worse and all portfolios seem to converge to the same point. The second underperformance relative to the benchmark in around 2009-2010 - the λ tilted portfolios again performed fairly better than the MVP, which had a big drawdown. Post this period, the performance seemed to be in line.

Using finer levels of λ, the resultant portfolios have characteristics closer to the MVP, as expected; one interesting aspect is that there are periods where the returns do not mirror the returns of the MVP.

5.2.1.3 Outline of Results - λ Level

The findings of the sections in the above sections - 5.2.1.1 and 5.2.1.2 - are consolidated below and more importantly informed the parameter estimation of λ used in the study. The results obtained using $\boldsymbol{S}$ were examined in combination with the findings highlighted in the sections above. The main properties which were considered in this estimation process are highlighted below (and ranked in order of importance):

1. Highest return
2. Highest risk-adjusted return
3. Lowest risk

The analysis revealed that there was no precise λ level that delivered the best overall return, risk-adjusted return and volatility. In fact if one targets any individual property when selecting the λ level, this could result in fairly different results. Examining the first property, return, one can note from Figure 18 below that over the total measurement period all portfolios formulated with a λ level lower than 0.01, outperformed the MVP, irrespective of the covariance parameter ($\boldsymbol{S}$ or $\boldsymbol{G}$) used. It is clear however that $\boldsymbol{G}$ enhances portfolio returns across the range of λ estimates in comparison to $\boldsymbol{S}$. As λ increases and emphasises the role of $\boldsymbol{Z}$ (i.e. prioritising the minimisation of the distance between the optimal weightings ($\boldsymbol{X}$) and $\boldsymbol{Z}$) on the objective function, this directly reduces the impact of risk on the formulation and results in very similar portfolios formulated at higher levels of λ, regardless of the covariance parameter utilised. As λ increases beyond 0.01, the return reduces significantly and at higher λ levels around 0.15, the portfolio return is approximately on par with the benchmark return.

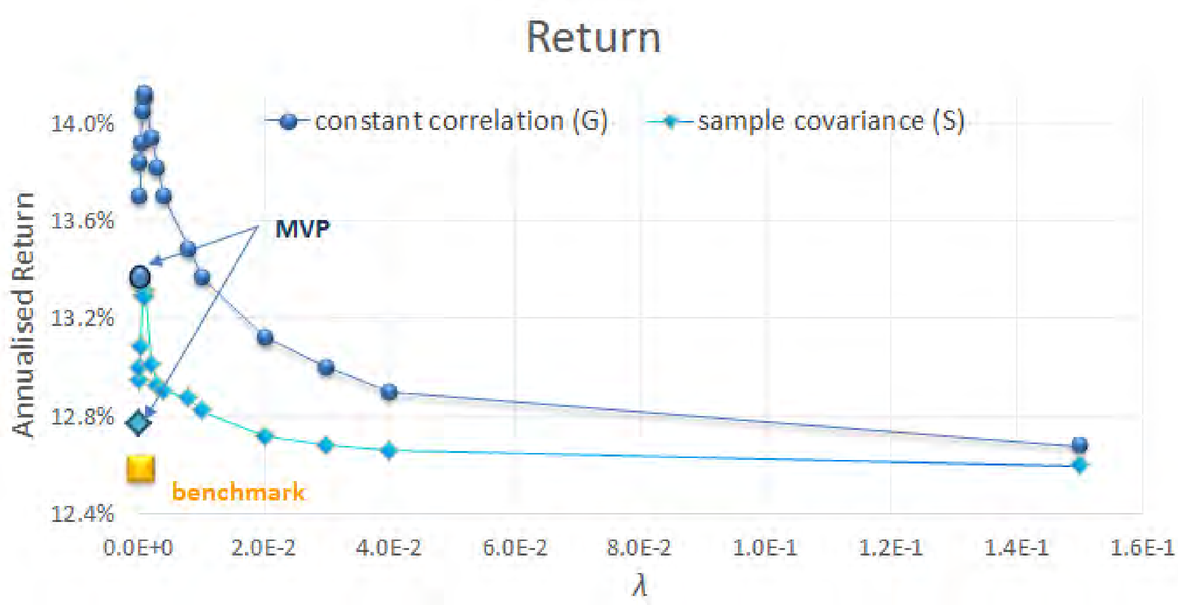

Figure 18 Annualised return measured over combined λ range

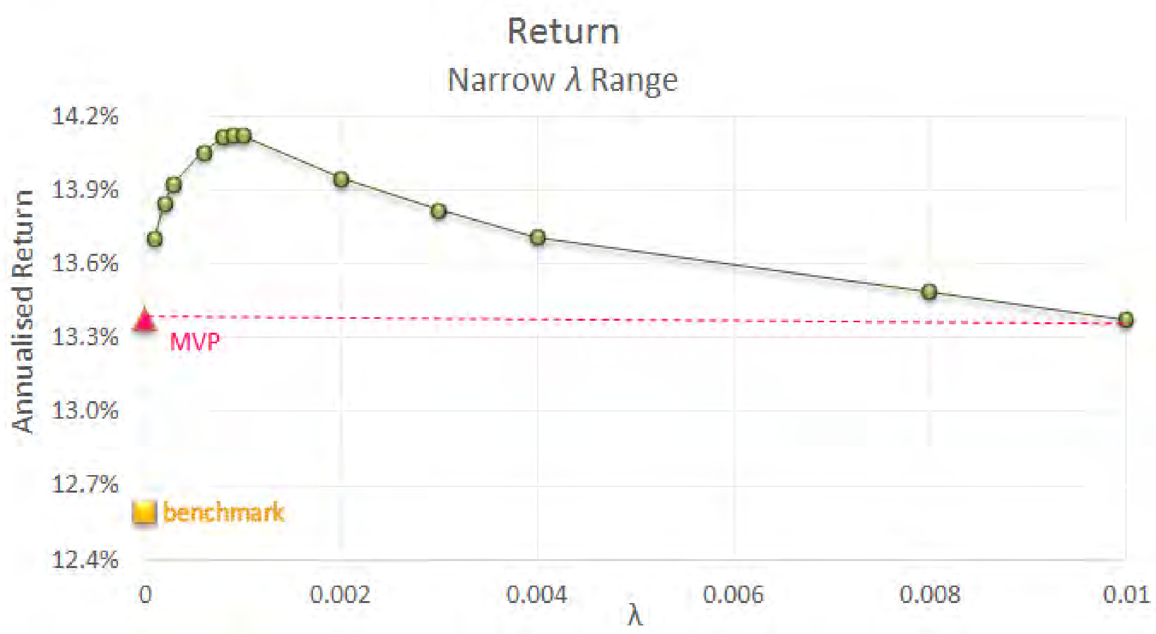

Figure 19 Annualised return measured over narrowed λ range using shrinkage (G)

If shrinkage is utilised and the period is narrowed to λ levels less than 0.01, as shown in Figure 19 above, one can note that returns increase monotonically and reach a maximum in the range: $0.0008 < \lambda < 0.001$. The return is quite sensitive to small changes in λ and as λ increases beyond the 0.001, the return starts to decay at a rapid rate. If $return_i = f(\lambda_i)$, for $i = 1, 2 \dots n$, defines the relationship between λ and return, then $f(\lambda)$ is not strictly monotonic for all values of λ, but can be represented by -

$$return_i = \begin{cases} f(\lambda_{i+1}) > f(\lambda_i) & 0 < \lambda_i < 0.001 \\ f(\lambda_{i+1}) < f(\lambda_i) & \lambda_i > 0.001 \end{cases} \qquad 5.1$$

A similar relationship to 5.1 above was observed using $\boldsymbol{S}$, specifically that the relationship between λ and return was non-monotonic.

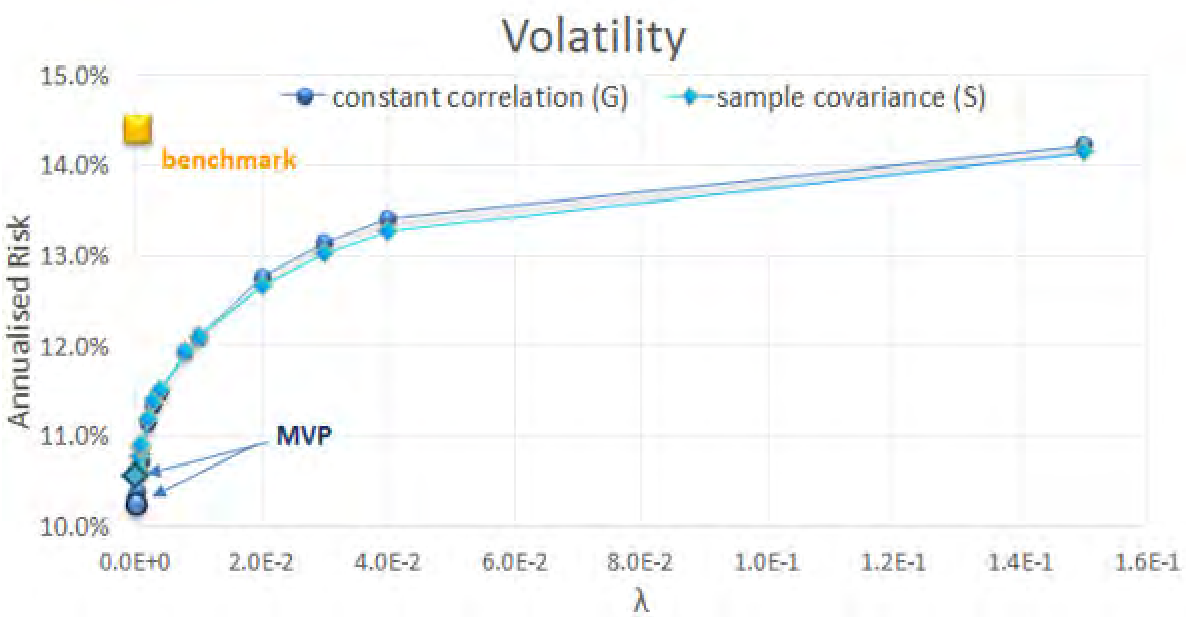

Figure 20 Annualised risk measured over combined λ range

The third attribute, risk, is more representative of a monotonically increasing sequence since for all $\lambda > 0$, the realised risk increases as λ increases. This is highlighted in Figure 20 above. All portfolios' realised risk plots between the MVP, which has the lowest risk and the benchmark which serves an upper bound and has the highest risk over the measurement period. Therefore portfolios formulated with λ values close to zero, tend to *mimic* the MVP and have lower risk. An interesting observation is that although $MVP(\lambda, \boldsymbol{G})$ (MVP formulated under the shrinkage method ($\boldsymbol{G}$)) measured risk is lower than $MVP(\lambda, \boldsymbol{S})$ (MVP formulated using the sample covariance ($\boldsymbol{S}$), for λ values > 0.01 the measured risk of

$portfolio(\lambda, \boldsymbol{S})$ (portfolios formulated using the sample covariance ($\boldsymbol{S}$)) is lower than $portfolio(\lambda, \boldsymbol{G})$ (portfolios formulated under the shrinkage method ($\boldsymbol{G}$)).

Finally, the relationship between $portfolio(\lambda, \boldsymbol{G})'s$ risk-adjusted return (measured by the Sharpe ratio) and λ can be represented similarly to 5.1 above, but adjusted as follows -

$$Sharpe_i = \begin{cases} f(\lambda_{i+1}) > f(\lambda_i) & 0 < \lambda_i < 0.0003 \\ f(\lambda_{i+1}) < f(\lambda_i) & \lambda_i > 0.0003 \end{cases} \qquad 5.2$$

Where $Sharpe_i$ represents the Sharpe ratio for λ_i where $i = 1, 2, \dots n$.
Since realised risk is monotonically increasing and rises at a faster rate than return, the maximum Sharpe ratio is attained at a lower λ threshold. The relationship is presented in Figure 21 below, for both covariance parameters ($\boldsymbol{G}$ and $\boldsymbol{S}$). As the λ increases above 0.001, the Sharpe ratio starts to decrease and approach the benchmark level (at high levels of λ).

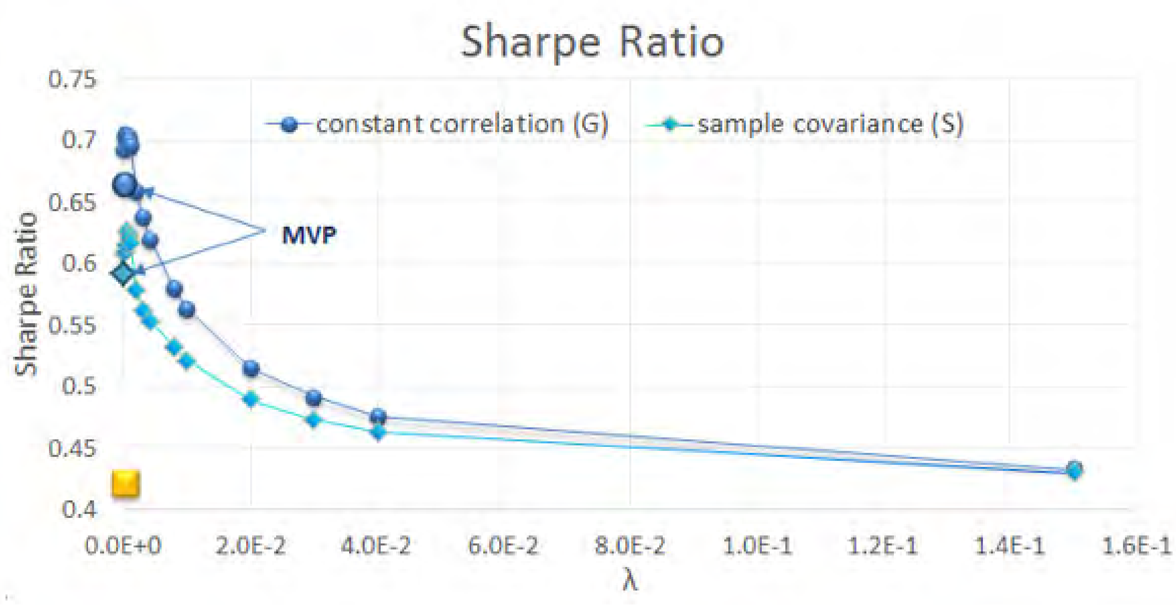

Figure 21 Annualised Sharpe Ratio measured over combined λ range

The maximum Sharpe Ratio for $portfolio(\lambda, \boldsymbol{G})$ is attained at a λ level close to 0.0003. The function above (equation 5.2) is represented over a narrower range $(0 \leq \lambda \leq 0.01)$ in Figure 22 below.

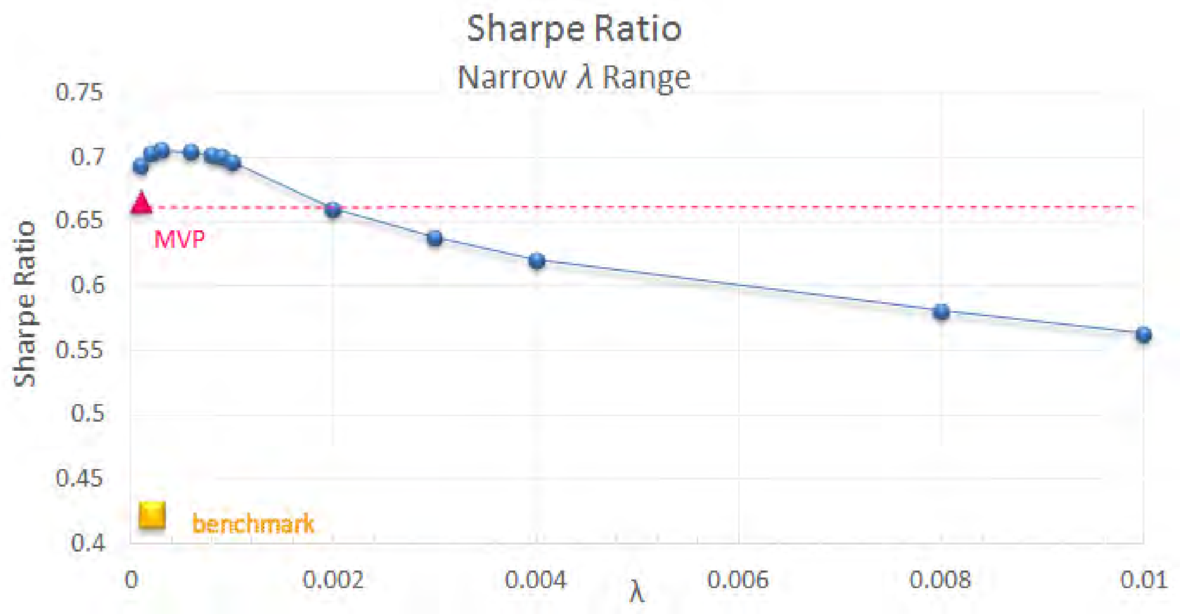

Figure 22 Annualised Sharpe Ratio measured over narrow λ range using shrinkage

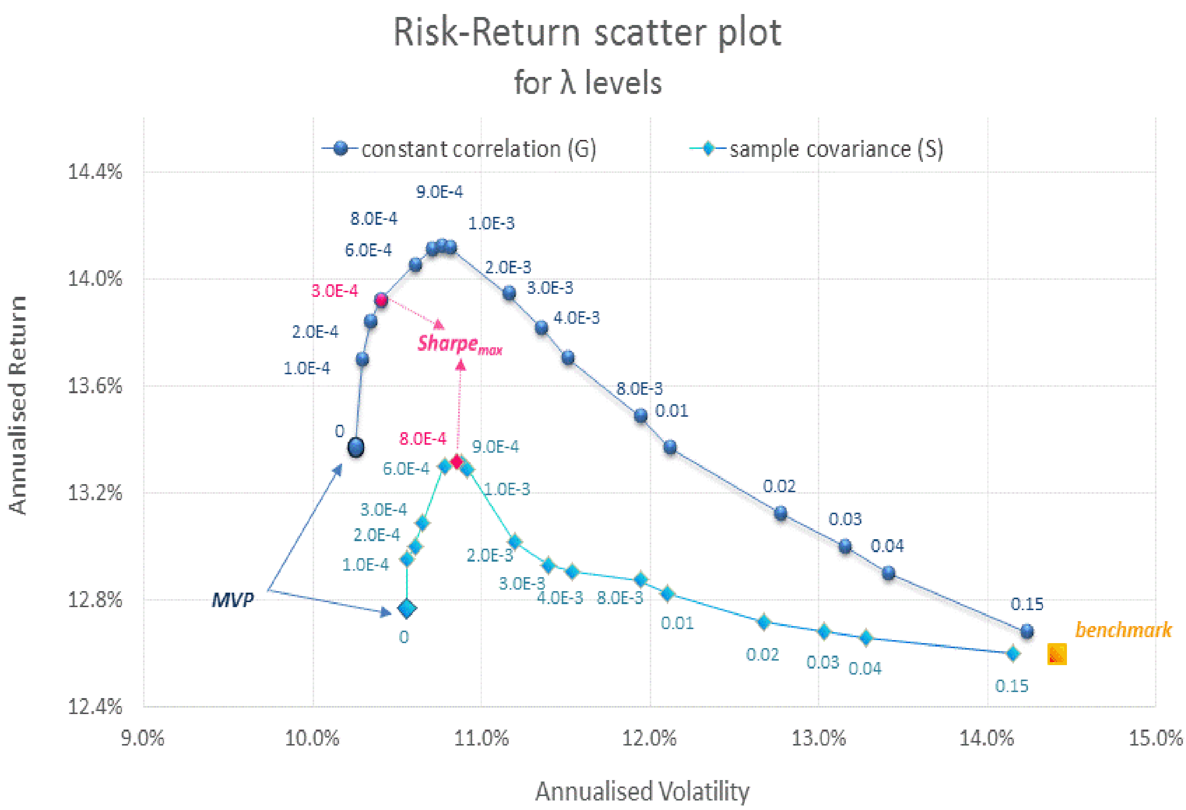

Figure 23 Risk-return scatter plot for various λ levels

The risk-return plot in Figure 23 above summarises the above relationships concisely. There is a clear advantage in employing a shrinkage method in determining the estimate of covariance used in MV-optimisation framework. One of the highlights is the reduced risk and elevated returns at lower levels of λ.

The efficient frontier formulated under $\boldsymbol{G}$ dominates the corresponding efficient frontier formulated under S. The maximum Sharpe Ratios are attained at different λ levels. There is a benefit in using the average volume indicator $(\boldsymbol{Z})$ as an enhancement to the standard MV-optimisation framework, this is evident mostly in terms of the enhanced return obtained relative to the MVP.

The average weighted volume was calculated per portfolio over time, and indicated (refer to Figure 24 below) a monotonically increasing relationship between λ and the average weighted volume. This indicated that as the portfolios approached the volume-tilted portfolio, stocks with a higher average volume indicator $(\boldsymbol{Z})$ were favoured in the selection process.

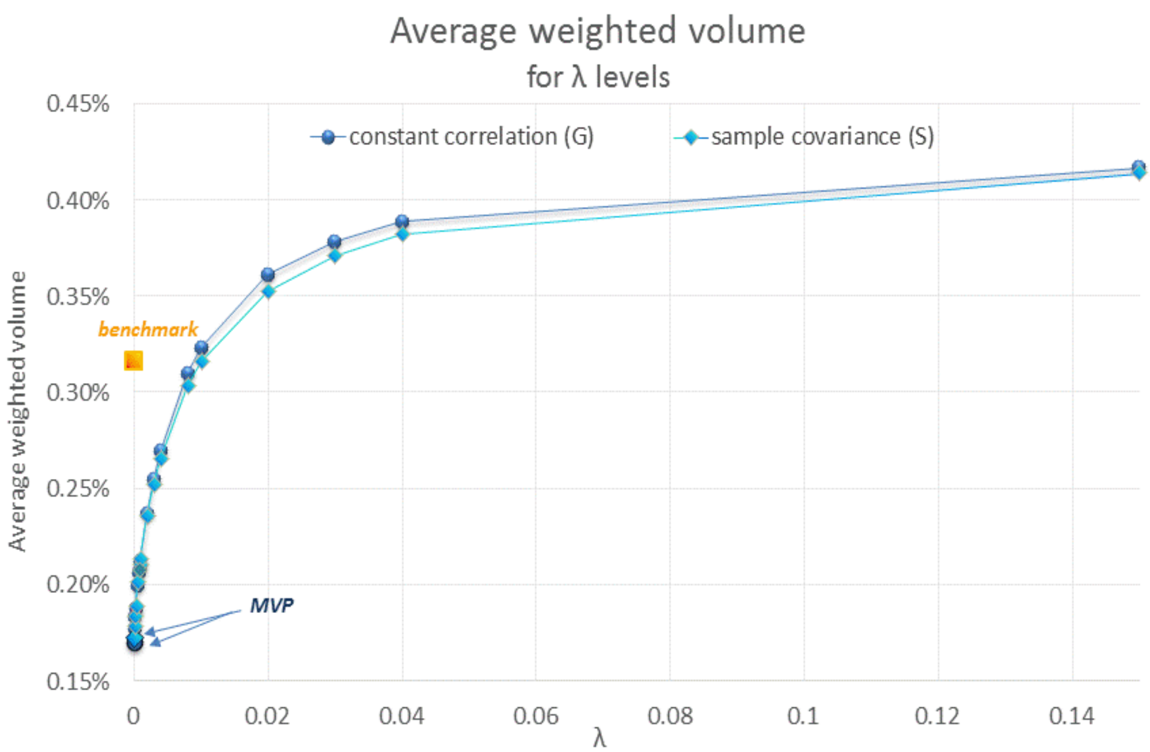

Figure 24 Average weighted volume for over combined λ range

An important consideration in any strategy is its performance in different market states, like bull or bear markets. The Capture ratio (refer to Bacon, 2004) encapsulates the portfolio's compound return over periods of directional change in the benchmark. The *upside capture ratio* measures fund performance relative to the benchmark over periods when the benchmark performance is positive, while the *downside capture* ratio measures relative performance in downturns.

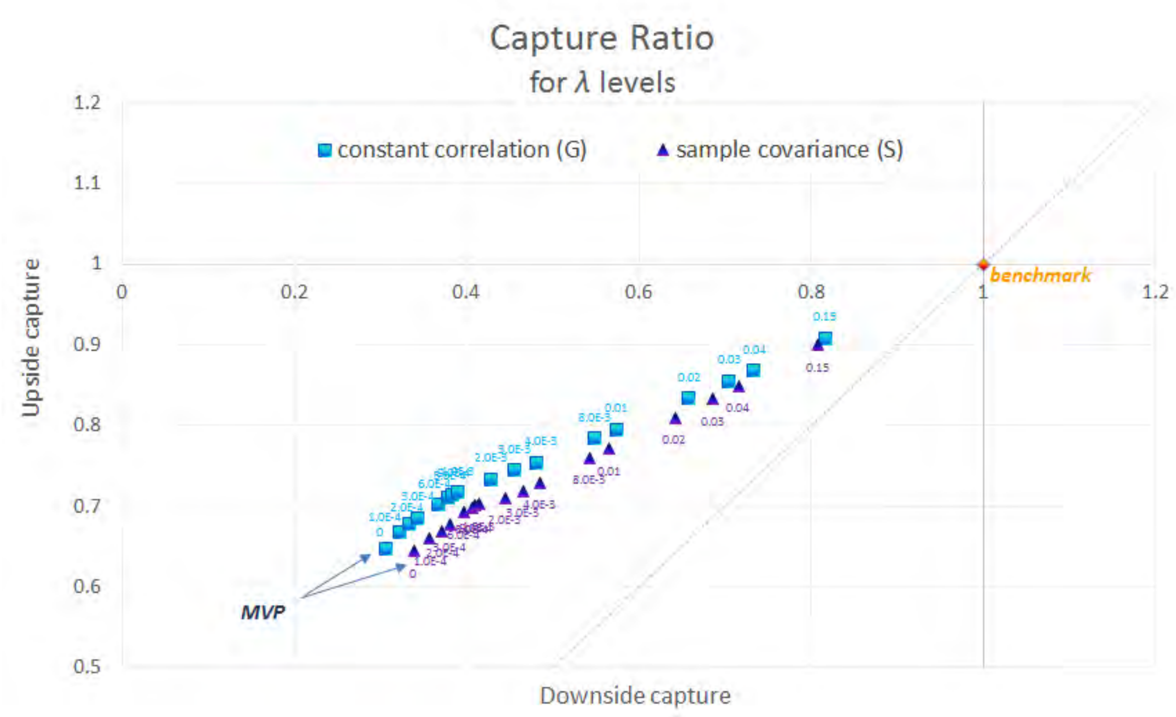

Figure 25 Capture Ratio - λ Levels

From Figure 25 above, it is evident that in periods when the benchmark returns positive performance, all portfolios regardless of λ level or covariance structure used, do not fully participate in the *upturn phase*. Both MVP_G and MVP_S benefit least in these periods and only *capture* around 64% of the benchmark return. As the λ tilt increases, there is an improvement in participation - at high levels of λ, this number increases to 90%.

A benefit however is the resultant performance in downturns, all portfolios formulated at $\lambda \geq 0$, fare better than the benchmark. The downside capture ratio ranges between 30% (for the MVP_G) and 81% (for λ =0.15). Portfolios formulated under $\boldsymbol{G}$, generally perform better then portfolios formulated under $\boldsymbol{S}$, capturing more performance in upturns and less of the negative performance associated with market drawdowns.

Refer to 8.1 Appendix A, Table 12 and Table 13, for all capture ratio statistics.

In addition, turnover levels were also examined - refer to 8.2 Appendix B, Table 18. There is a general reduction in turnover across all portfolios (where λ > 0) in comparison to the MVP (where λ=0) regardless of the covariance structure employed. Turnover directly affects the costs associated with trading and higher turnover levels will result in lower overall net-of-cost returns. Employing shrinkage in addition is advantageous, since the recorded turnover is lower for all portfolios.

There is a marked improvement in the maximum drawdowns observed in comparison to the benchmark. The MVP results in the lowest drawdowns for both covariance structures. As λ increases, the maximum drawdown observed increases as well. The benchmark recorded the worst drawdown over time. The most significant observation however was the lower drawdown statistics obtained when $\boldsymbol{G}$ (shrinkage) was employed in comparison to $\boldsymbol{S}$. Refer to 8.3 Appendix C, Table 21.

A useful tool which examines the robustness of a portfolio strategy includes the computation of the probability that an investor will outperform the benchmark over a given time horizon regardless of entry point.

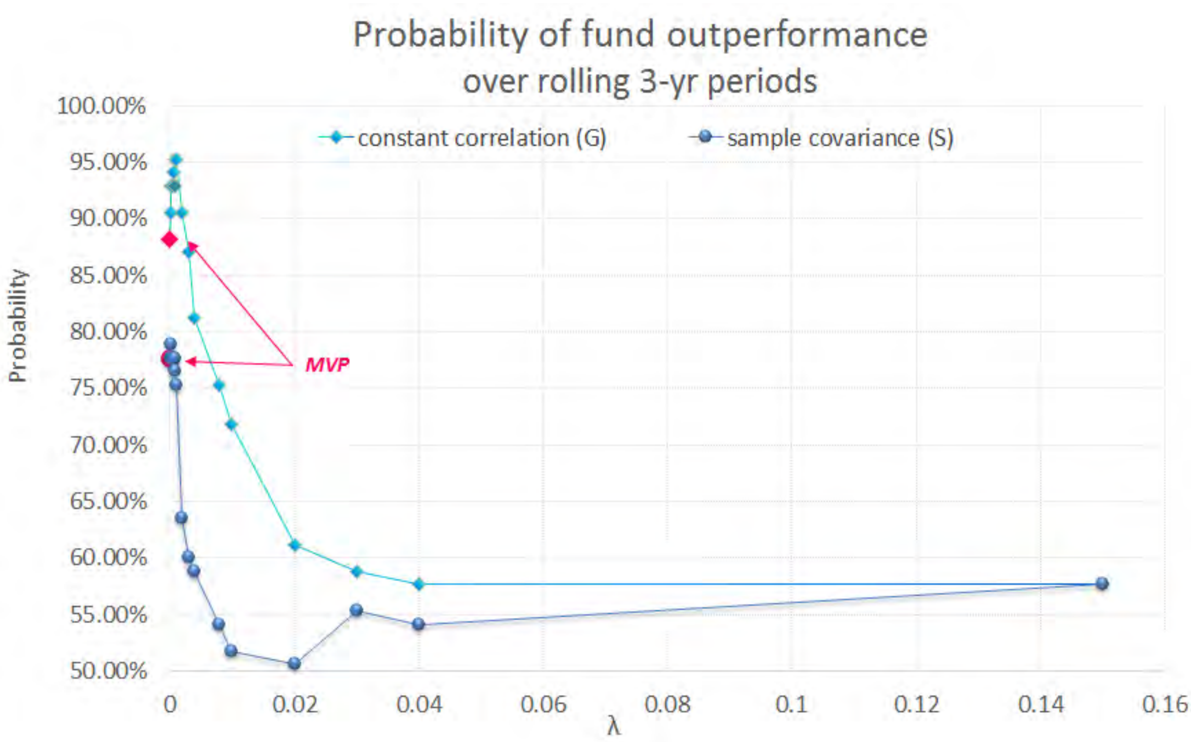

Figure 26 Probability of outperformance for λ various levels

This statistic was calculated over trailing three-year periods, highlighted in Figure 26 above and is quite compelling for portfolios generated at lower levels of λ. Again the benefit of employing shrinkage is evident, in that there is a formidable increase in the probability of outperformance across the majority of portfolios.

5.2.2 Estimation of θ (Diversification (Herfindahl–Hirshman Index) Tilt)

The same technique described in the section 5.2.1 above, was applied to determine efficient levels of θ.

5.2.2.1 Broad θ Range

The distribution of the largest 40 holdings indicate that with the introduction of the concentration parameter, θ, the effect on the weighting distribution was quite evident even at very low levels of θ, as indicated in Figure 27 below.

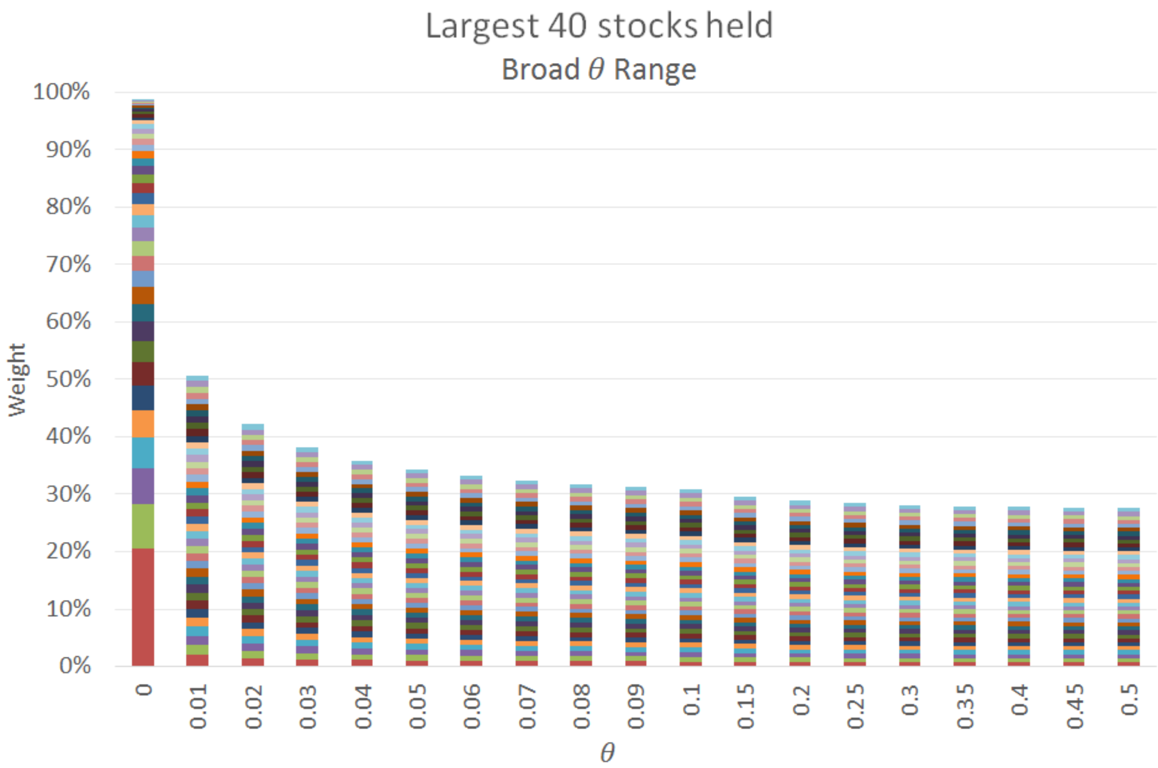

Figure 27 Weight distribution of largest 40 holdings for $\theta \geq 0$

The top 40 holdings account on average for close to 100% of the MVP (where $\theta = 0$), indicating a high level of concentration, which is amplified by the fact that one stock makes up 21% of the portfolio. This exposure declined substantially however as the tilt toward the equally weighted portfolio was increased ($\theta > 0$). As θ increased, one effectively increases diversification in the portfolio. The aggregate exposure of the top 40 holdings decreased substantially to around

40% at lower levels of θ. This number moderated further at higher levels of θ ($\theta \geq 0.25$) to 28%. It is clear that as θ increases, the portfolios approach the equally weighted portfolio ($HHI = \frac{1}{N} = 0.67\%$).

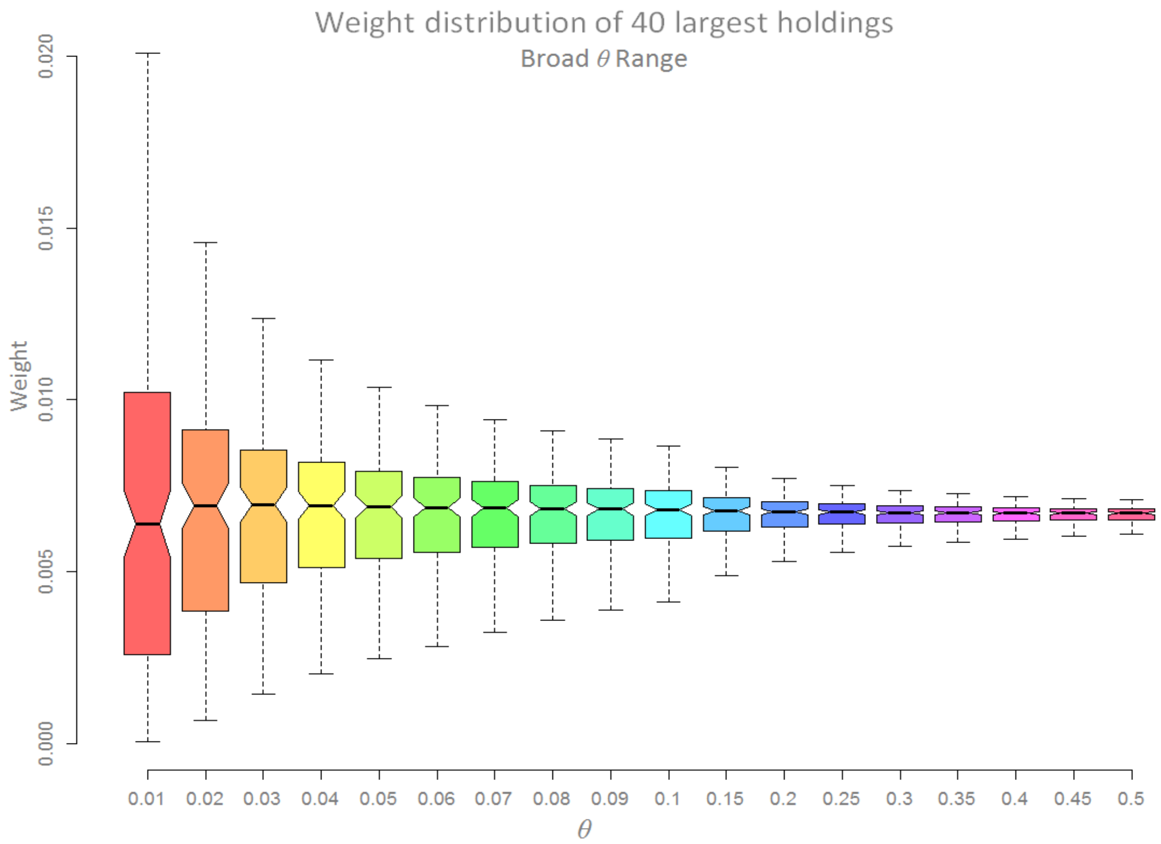

Figure 28 Box plot of largest 40 weights for portfolios formulated at $\theta > 0$

Taking a closer look at the distribution of the largest average holdings across time, the box plot shown in Figure 28 above, indicates that at lower levels of θ, the maximum weight which is in excess of 2.0%, is fairly different from the corresponding maximum weight in the MVP, which is approximately 21%. The inter-quartile range is broader and most distinct at lower levels of θ. This range becomes very compressed at θ levels greater than 0.15, as portfolios approach the equally weighted portfolio.

Table 6 Standard deviation of weights measured over the period (April 2006 - April 2016)

θ	0	0.01	0.02	0.05	0.09	0.25	0.35
σ	1.00%	0.19%	0.14%	0.08%	0.05%	0.02%	0.02%

Table 6 above represents the volatility of weights across a subset of portfolios. The portfolios were selected on a similar basis as described above in 5.2.1, *that is*, when there was at least a 0.05% change in the average monthly return, the portfolio was added.

There is a monotonically decreasing relationship between the value of θ and volatility of weights, which declines substantially at higher levels of θ. When θ=0.35, the measured volatility is 0.03%. This is very low, which is expected, since the month-to-month change in the portfolios formulated at θ levels greater than 0.35 characterise the equally weighted portfolio, which by its very nature will have this property.

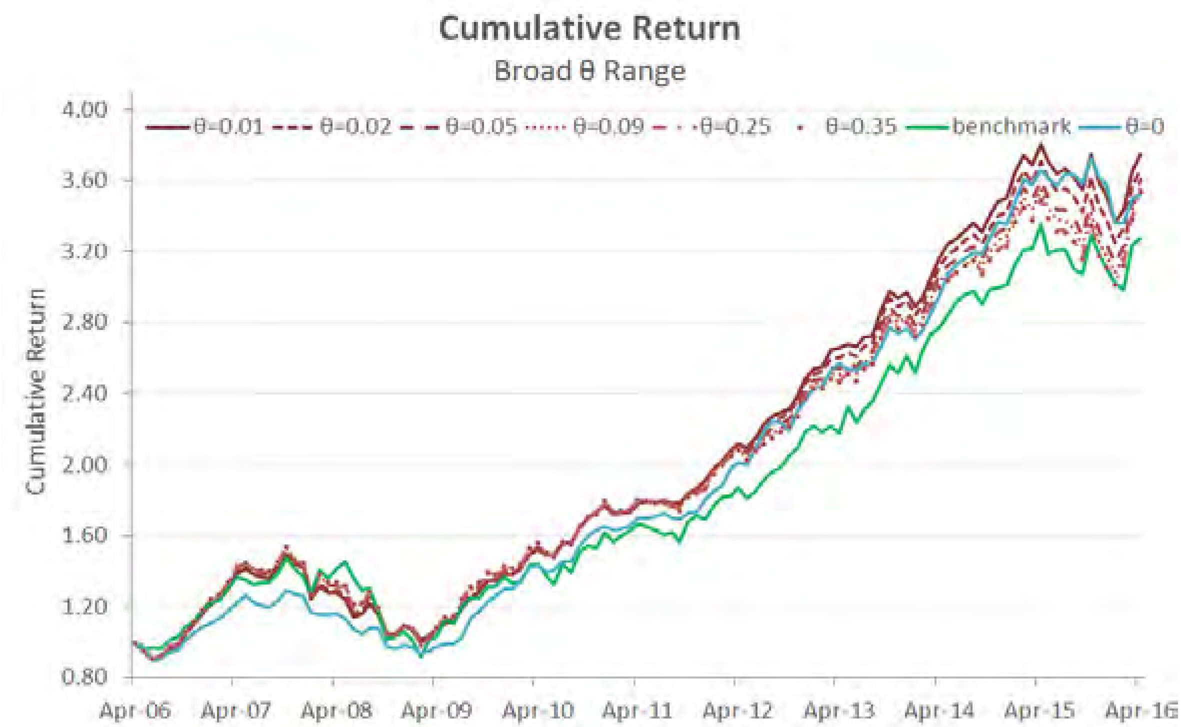

Figure 29 Cumulative Return - Broad θ range

The cumulative return and relative return measured over the entire period is shown in Figure 29 above and Figure 30 below. All θ portfolios performed better than the MVP ($\theta = 0$) and benchmark. The benchmark and MVP are the most distinctly different portfolios in comparison to the other portfolios. From inception to around November 2007, both the MVP and benchmark underperformed the θ portfolios. All portfolios including the MVP underperformed the benchmark over the financial crisis during November 2007 and June 2008. Post the crisis both the

MVP and θ portfolios recovered sharply and outperformed until February 2009. The drawdown over the next period which ended in July 2009 saw the MVP lose 15% of its value from its high, while the θ portfolios fared much better losing only 9% of value against the benchmark.

The benchmark underperformed for a considerable amount of time thereafter. The MVP outperformed significantly from around May 2012 – and was the best performing portfolio until December 2015 - very little differentiated the θ portfolios from each other, especially in the initial period from April 2006 to April 2012.

Post this period however, there was a fair amount of dispersion in returns and portfolios with higher θ levels seemed to be quite heavily impacted and experience larger drawdowns in the latter part of 2015, before sharply rebounding in the first four months of 2016.

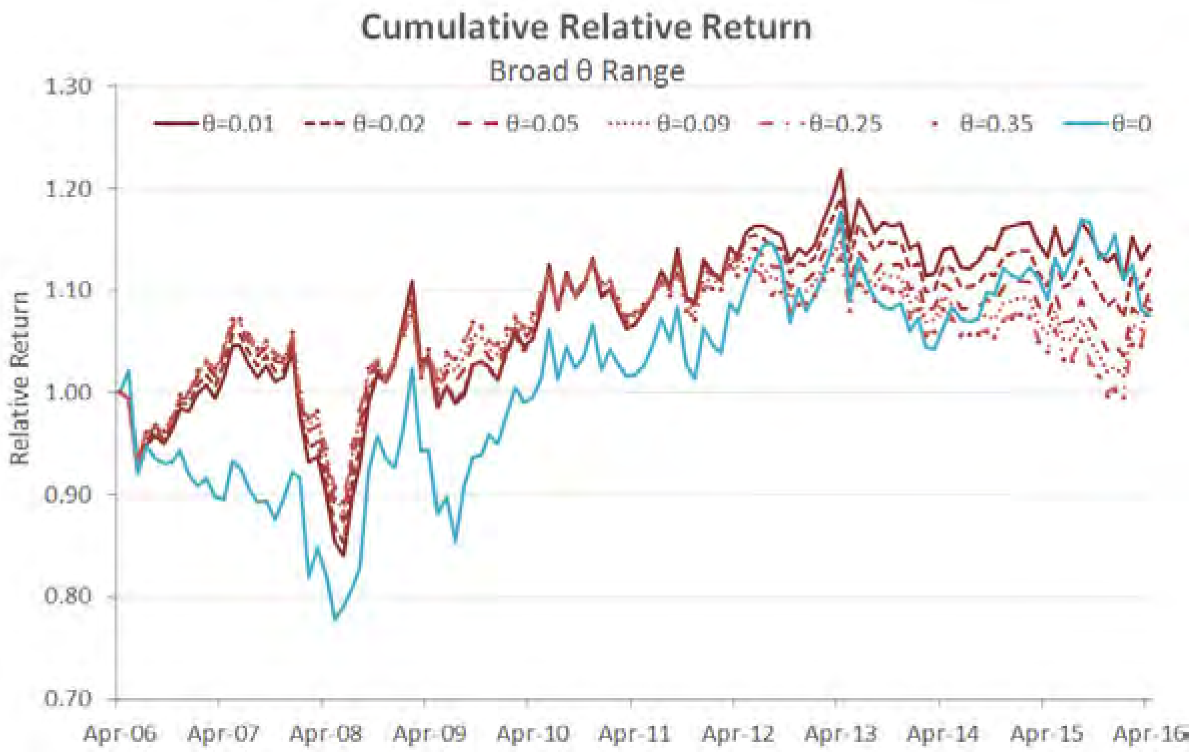

Figure 30 Cumulative relative Return - Broad θ range

The MVP post April 2009, over most 24-month rolling periods, fared better then all θ portfolios as well as the benchmark (see Figure 31 below). There was a convergence of returns from December 2013 to around December 2014. Pre- and post this period however, the MVP performance was quite distinct from the

θ portfolios and benchmark. Another attribute that was investigated was the volatility of the θ portfolios. It was important to include the benchmark and MVP, for comparative purposes.

Figure 31 Rolling 24-month returns – Broad θ range

Figure 32 Rolling 24-month volatility – Broad θ range

In Figure 32 above, the risk is shown over 24-month rolling periods. As expected the MVP returned the lowest realised risk over the majority of rolling measurement periods. Although post June 2010, there was a convergence in realised risk across both the set of θ portfolios and the MVP. In fact, from around May 2012 to June 2014, the θ portfolios' realised risk was lower than the MVP. The benchmark returned the highest risk across the majority of the period.

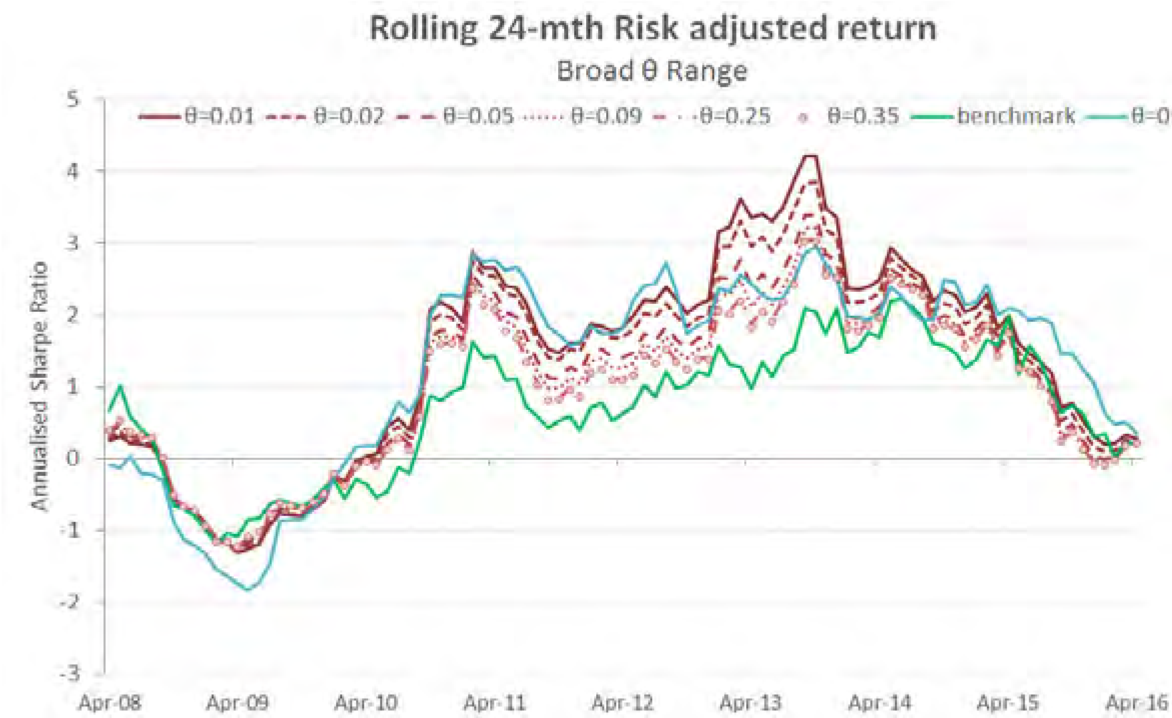

Figure 33 Rolling 24-month risk adjusted return - Broad θ range

The rolling 24-month Sharpe ratios shown in Figure 33 indicate that at lower levels of θ, portfolios tend to yield higher risk-adjusted returns on average than the benchmark and MVP.

Table 7 Average 24-month Sharpe Ratio for Broad θ Range

Portfolio	Sharpe Ratio
θ=0.01	1.39
θ=0.02	1.27
θ=0.05	1.10
θ=0.09	1.02
θ=0.25	0.95
θ=0.35	0.94
benchmark	0.69
MVP	1.25

Table 7 above shows the average 24-month rolling Sharpe ratio. For $\theta \leq 0.02$, the Sharpe ratios are higher than the MVP. The benchmark is an outlier and the low value can be attributed in the most part to higher levels of realised risk over the measurement period.

The fact that higher returns, higher risk-adjusted returns and lower volatilities were observed at lower levels of θ, resulted in further examination of θ levels within a narrower range $(0 < \theta \leq 0.01)$. The results of this analysis are discussed further in the section below.

5.2.2.2 Narrow θ Range

The range was narrowed and portfolios were selected based on the same criteria described above. The θ portfolios are given in Table 8 below.

Table 8 θ levels selected from the narrow θ range

θ	1.0E-4	2.0E-4	3.0E-4	8.0E-4	1.2E-3	1.8E-3	2.5E-3	3.0E-3	6.0E-3	8.0E-3	1.0E-2

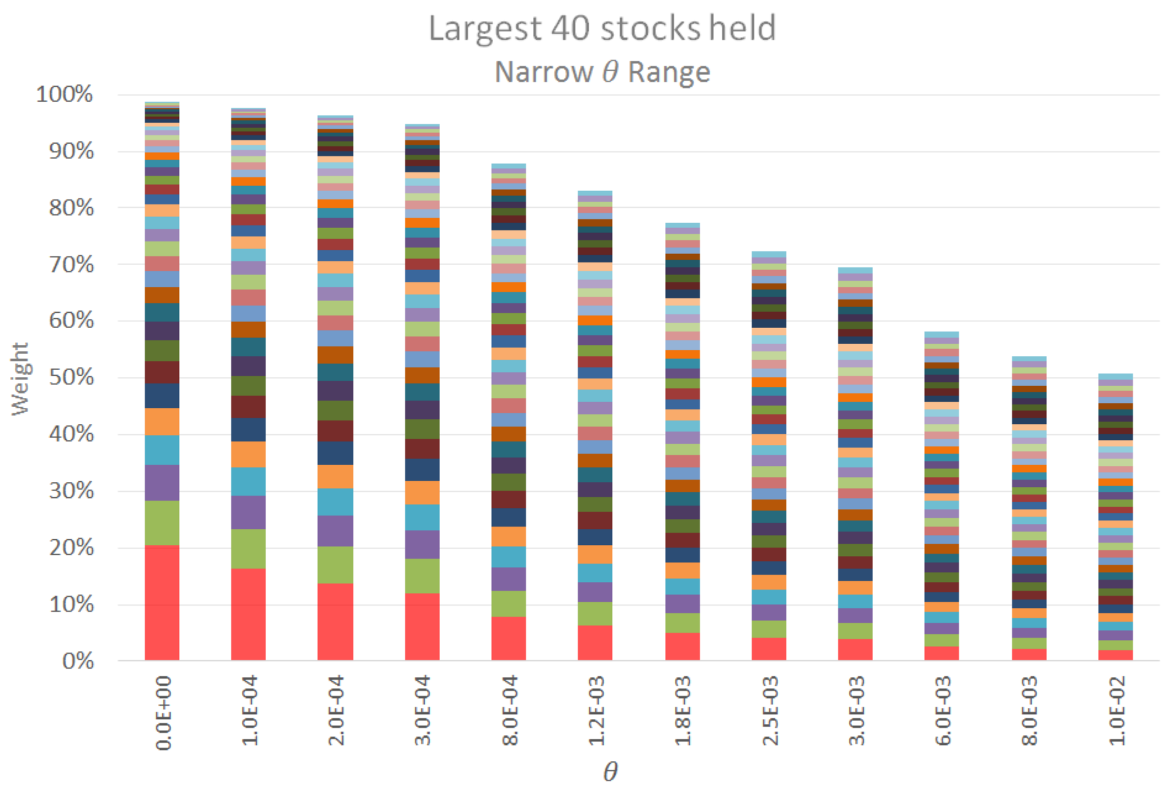

Figure 34 Largest 40 holdings for the Narrow θ range

Figure 34 above, indicates that as the θ range is narrowed, the diversity decreases. The top holdings now account for close to 100% of the total portfolio exposure. Although this is the case, there is a far more diversified holding within this grouping then the MVP. Exposure to tail stocks in the portfolio will be increased as θ increases.

The volatility of weights and maximum HHI observed over time are given in Table 9 below. There is a reduction in the HHI as θ increases. The MVP returns the highest HHI of 31.9%, indicating that the MVP portfolios generated tend to be fairly concentrated. With a slight introduction or tilt to the *concentration* preference parameter (θ), the HHI decreases and increases the diversification effect. The benchmark is fairly diverse in comparison to the MVP, but still more concentrated than portfolios with greater tilts, that is, with θ levels ≥ 1.2E-3.

Table 9 Standard deviation of weights and maximum HHI measured over the period (April 2006 - April 2016)

θ	0	1.0E-4	2.0E-4	3.0E-4	8.0E-4	1.2E-3	1.8E-3	2.5E-3	3.0E-3	6.0E-3	8.0E-3	1.0E-2	benchmark
HHI	31.9	21.1	15.4	12.1	6	4.4	3.4	2.8	2.5	1.7	1.5	1.4	4.7
σ	1.00%	0.88%	0.77%	0.72%	0.54%	0.46%	0.40%	0.35%	0.32%	0.25%	0.22%	0.19%	0.12%

The volatility of weights tends to give an indication of how the average individual stock weight is changing over time. Higher volatility will incur far more rebalancing month-to-month, which will directly incur higher turnover and trading costs. It would seem that the benefit of adding a tilt reduces this volatility of weights over time. The benchmark has the lowest volatility, it is rebalanced on a less frequent basis and stock weightings will generally float with price intra-quarter between rebalances.

The cumulative absolute and relative returns are plotted below in Figure 35 and Figure 36 respectively. There is a convergence in the returns as the level of θ is narrowed, especially post 2008. The interesting aspect is that there seems to be an interchange in performance of the θ portfolios. Prior to the credit crisis of 2008, the higher θ levels fared better and do not underperform as much as portfolios which have lower exposure to the tilt, *that is*, portfolios that have lower θ levels, will indirectly increase their exposure to the standard MV risk parameter ($\boldsymbol{W}$) and reduce exposure to the equally weighted portfolio. However, post 2008, these portfolios catch up and performance is less different. The highest return was achieved at a θ level close to 0.0018; this portfolio seems to have an optimal combination of exposure to both the risk and concentration preference parameters.

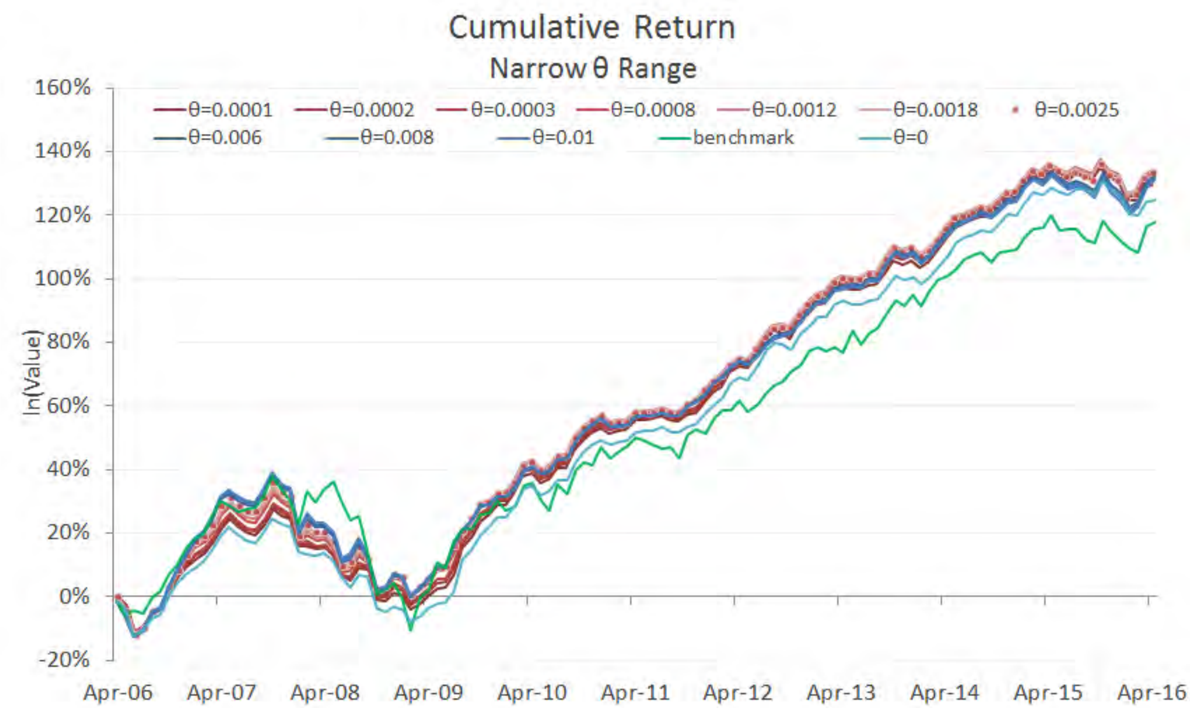

Figure 35 Cumulative Return - Narrow θ range

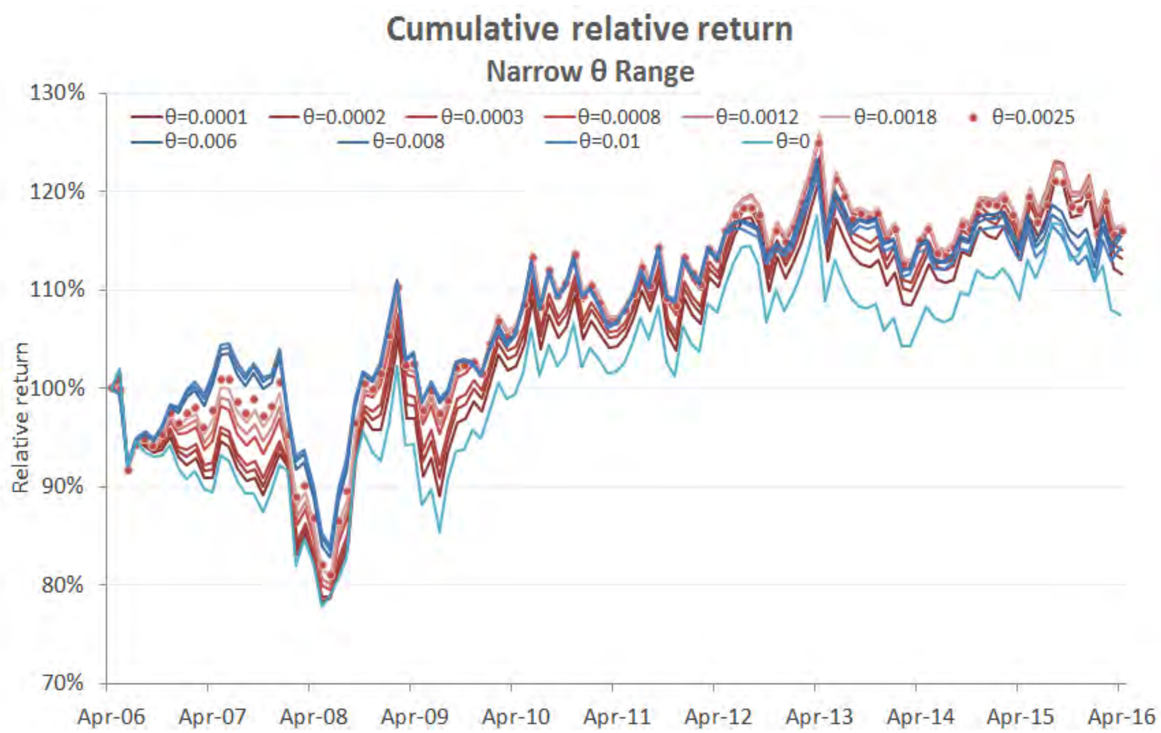

Figure 36 Cumulative relative Return - Narrow θ range

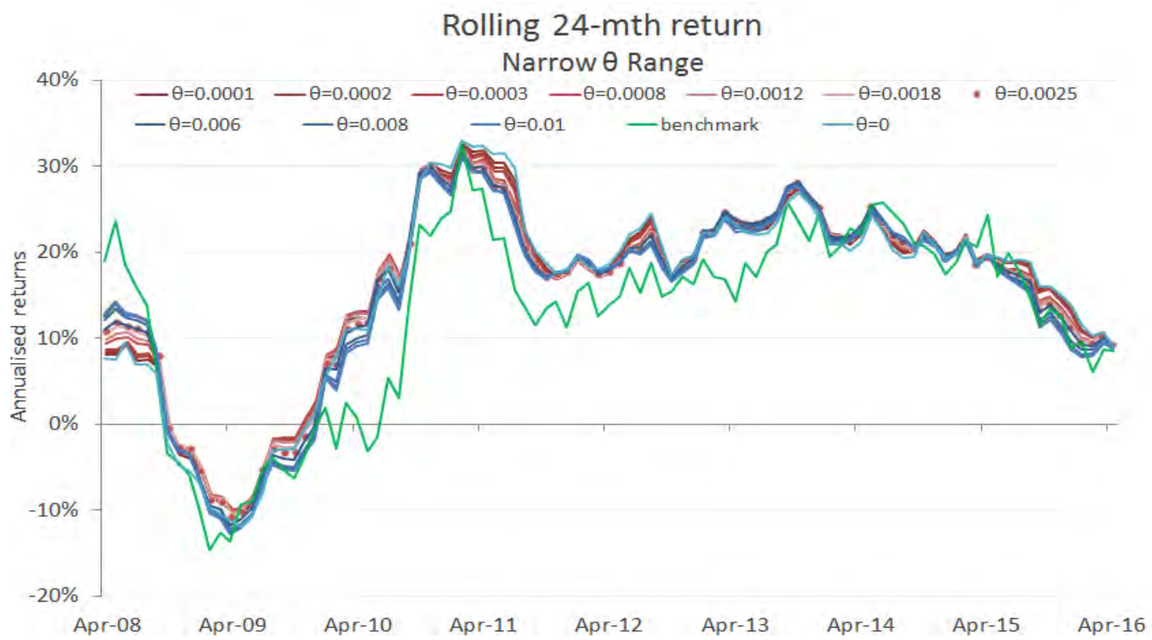

Figure 37 Rolling 24-month returns – Narrow θ range

Again, from Figure 37 above, there is no clear difference in return between the portfolios formulated at finer θ levels.

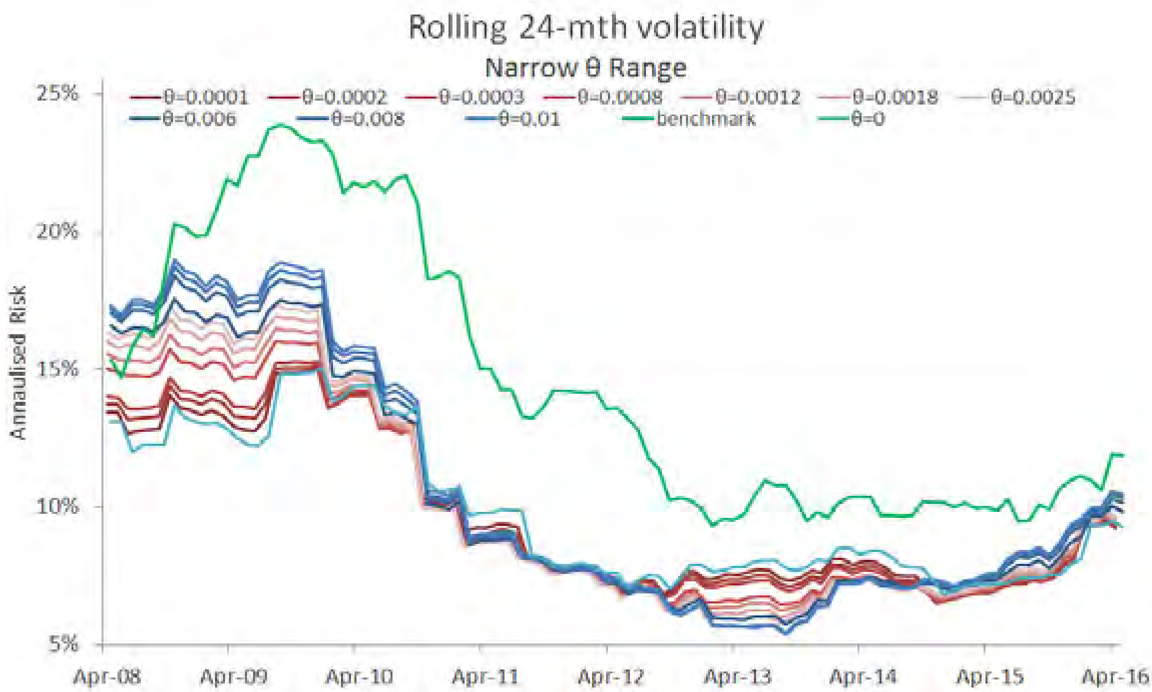

Figure 38 Rolling 24-month volatility – Narrow θ range

In Figure 38, only periods where there was a noticeable difference in risk is between April 2008 and April 2010 and June 2012 and January 2014. Over the remaining periods the rolling 24-month risk was very similar across all the sets of θ portfolios and the MVP.

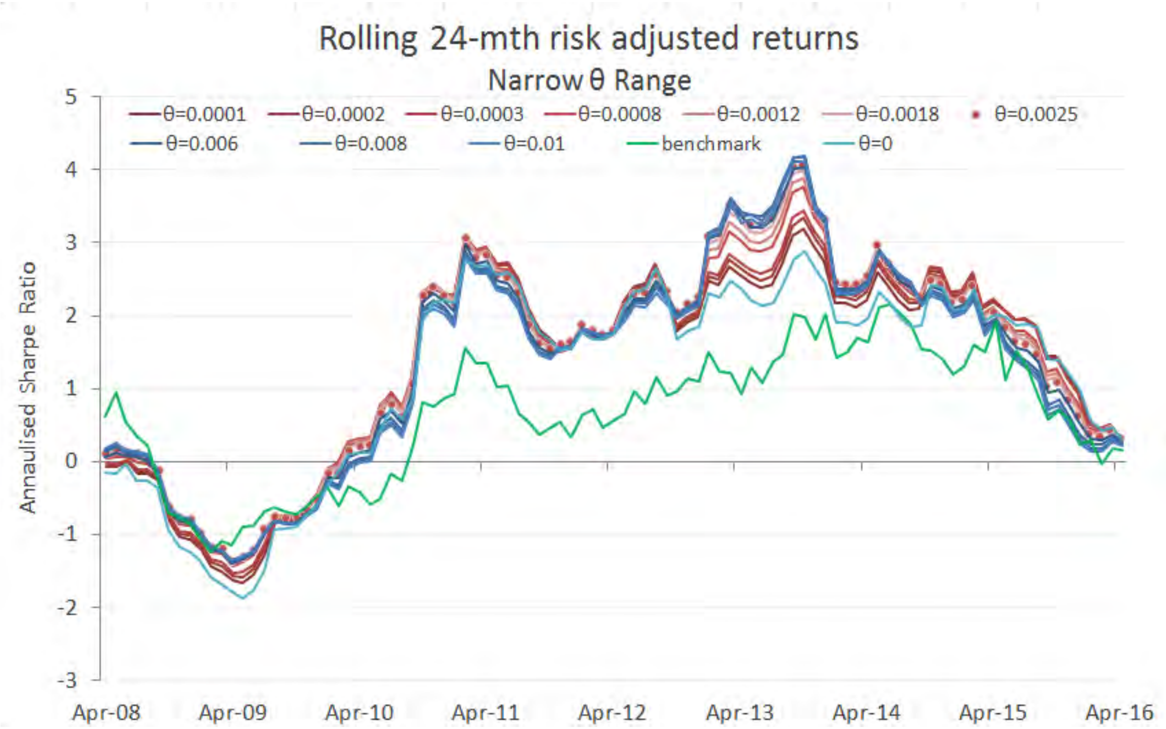

Figure 39 Rolling 24-month risk-adjusted return – Narrow θ range

Table 10 Average 24-month Portfolio Statistics for Narrow θ Range

Portfolio	Return	Risk	Sharpe Ratio
θ=1.0E-4	15.88%	9.17%	1.33
θ=2.0E-4	15.96%	9.16%	1.37
θ=8.0E-4	15.86%	9.30%	1.43
θ=1.2E-3	15.79%	9.39%	1.44
θ=1.8E-3	15.68%	9.49%	1.45
θ=2.5E-3	15.50%	9.57%	1.45
θ=3.0E-3	15.88%	9.61%	1.45
θ=6.0E-3	15.09%	9.82%	1.44
θ=8.0E-3	14.94%	9.94%	1.42
θ=1.0E-2	14.78%	10.05%	1.39
benchmark	12.79%	13.96%	0.69
MVP	15.53%	9.22%	1.25

The rolling 24-month Sharpe ratios shown in Figure 39 indicate that at lower levels of θ, portfolios tend to yield higher risk-adjusted returns on average than the benchmark and MVP. Table 10 shows the average 24-month rolling Sharpe ratios, volatilities and returns for $\theta \leq 0.01$. All portfolios performed better in terms of their Sharpe ratios and returns than the MVP. The benchmark is an outlier and the low Sharpe ratio can be attributed in the most part to higher levels of realised risk over the measurement period.

5.2.2.3 Outline of Results - θ Level

The results of the sections 5.2.2.1 and 5.2.2.2, above were consolidated below and inform the parameter estimation of θ used in the study. The results obtained using the sample covariance, $\boldsymbol{S}$, were examined in combination with the findings highlighted in the aforementioned sections. Over the course of the this process, which at inception was to identify the most optimal level of θ, the objective evolved into trying to find levels that were suitable across multiple facets of measurement.

The main properties which were considered in this estimation process are highlighted below (and ranked in order of importance):

1. Highest return
2. Highest risk-adjusted return
3. Lowest risk

The returns were measured over the total period and are highlighted in Figure 40 below.

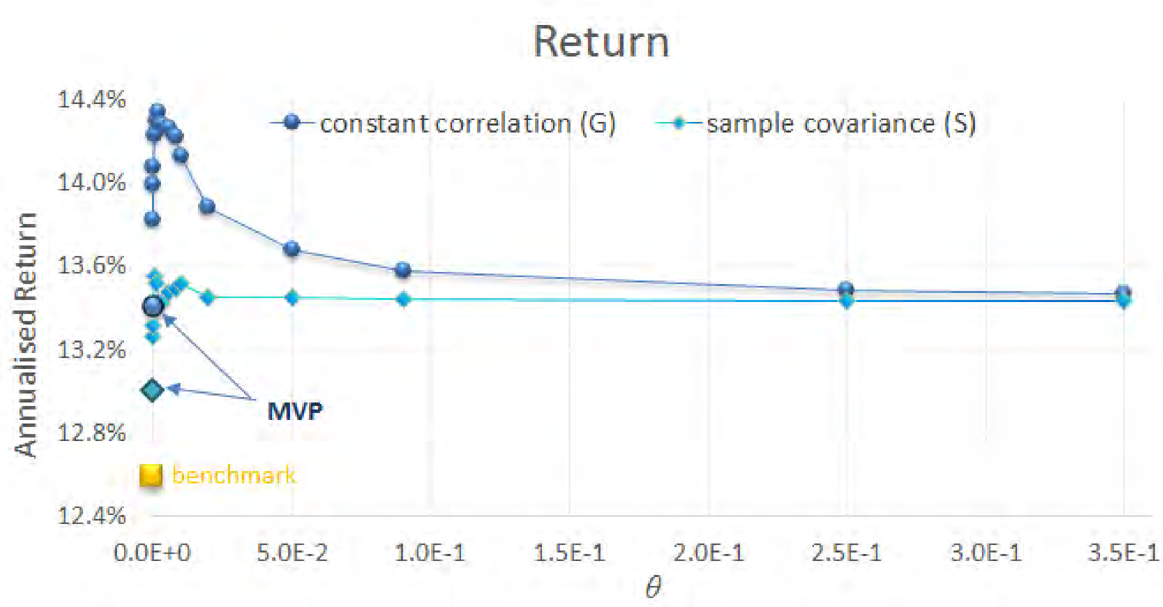

Figure 40 Annualised return measured over combined θ range

It is clear that introducing the *diversification tilt* has the advantage of lowering overall concentration; but this has resulted in outperformance of both the MVP and benchmark over the total measurement period. Portfolios formulated using $\boldsymbol{G}$ produced a return series that was not monotonic, since for θ levels below 0.002, the returns increased and approached a maximum at θ levels around 0.0018 before there was a sharp decline as θ increased beyond 0.002.

As θ increased and portfolios approached the equally weighted portfolio, the returns declined but still outperformed the MVP and benchmark. Employing $\boldsymbol{S}$ produces portfolios that have similar properties to those described above (using $\boldsymbol{G}$), the effect however is far less pronounced. The return series plots below that of $\boldsymbol{G}$ for values of $\theta \geq 0$. But as θ increases beyond 0.25, there is a general convergence in returns of both $portfolio(\theta, \boldsymbol{S})$ and $portfolio(\theta, \boldsymbol{G})$, as the impact of the risk parameter is minimised and the tilt to diversification amplified. Returns generated for $portfolio(\theta, \boldsymbol{S})$ increase and reach a maximum when $\theta \approx 0.0008$, there is a slight decrease in the return however as θ increases beyond this level and it *flat lines* as θ increases beyond 0.0012.

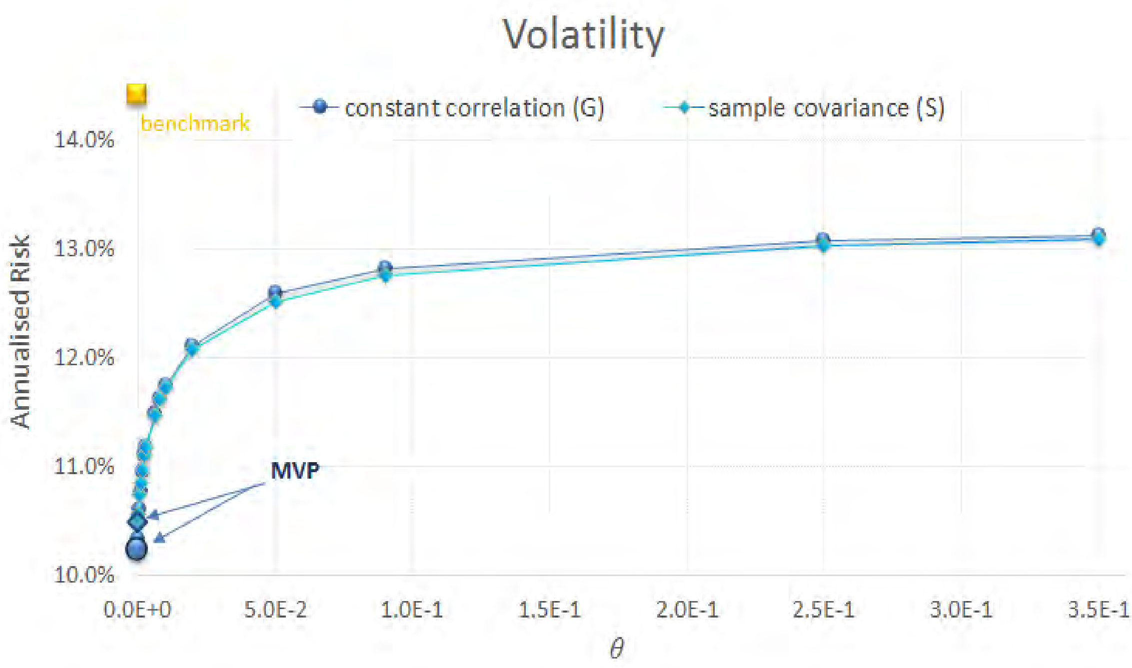

Figure 41 Annualised risk measured over combined θ range

The risk measured for all θ portfolios ranged between the MVP and benchmark, as indicated in Figure 41 above. No level of θ offers a risk reduction lower than the MVP. The benchmark has a risk around one percent higher than the portfolios constructed at higher θ levels. As θ increases the realised risk seems to approach 13.2%, which is the measured risk of the equally weighted portfolio. The series increases monotonically but at higher levels of θ, portfolios quickly approach the equally weighted portfolio (which has the highest risk of all θ portfolios, since the tilt completely reduces the effect the risk parameter has on the optimal solution for all levels of θ. The $MVP(\theta, \boldsymbol{G})$ measured risk is lower then $MVP(\theta, \boldsymbol{S})$, but as θ increases there is convergence between the risk of portfolios constructed under both $\boldsymbol{G}$ and $\boldsymbol{S}$. In fact for $\theta > 0.003$, $portfolio(\theta, \boldsymbol{S})$ risk is lower than $portfolio(\theta, \boldsymbol{G})$.

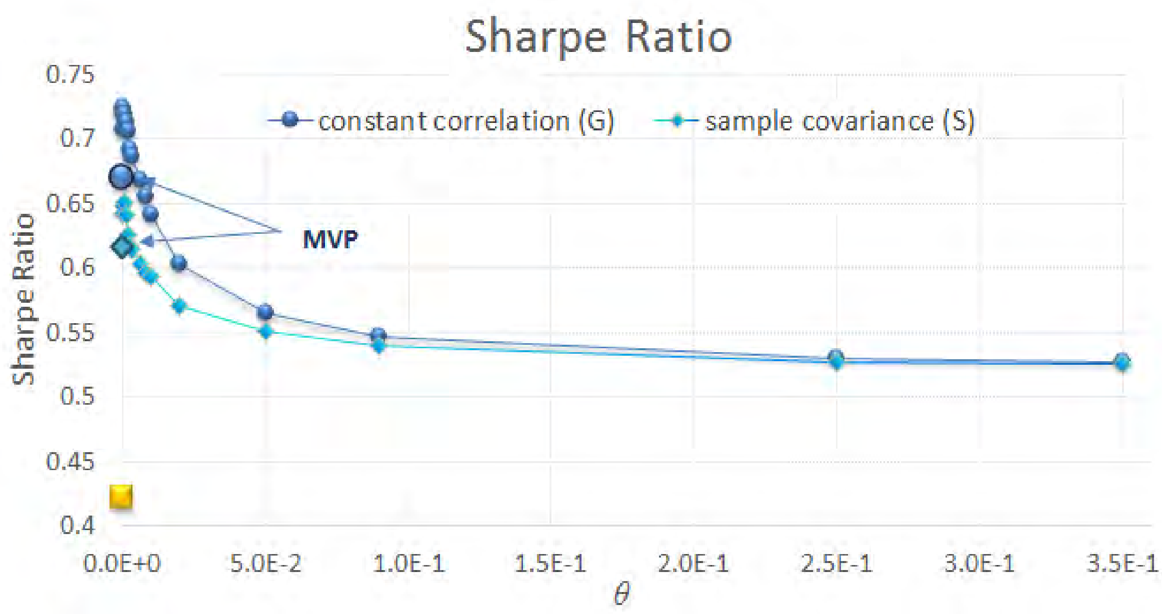

Figure 42 Annualised risk-adjusted return measured over combined θ range

Finally, the relationship between risk-adjusted return (measured by the Sharpe ratio) and θ can be represented as follows for $portfolio(\theta, \boldsymbol{G})$-

$$Sharpe_i = \begin{cases} f(\theta_{i+1}) > f(\theta_i) & 0 < \theta_i < 0.0003 \\ f(\theta_{i+1}) < f(\theta_i) & \theta_i > 0.0003 \end{cases} \qquad 5.3$$

Where $Sharpe_i$ represents the Sharpe ratio of portfolios constructed using $\boldsymbol{G}$ for θ_i where $i = 1, 2, \dots n$. The threshold (0.0003 in equation 5.3 above) for $portfolio(\theta, \boldsymbol{S})$-portfolios constructed using $\boldsymbol{S}$ - was approximately 0.0008.

Since realised risk is monotonically increasing and rises at a faster rate than return, the maximum Sharpe ratio is attained at a lower θ threshold. The relationship is presented in Figure 42, above. As θ increased above 0.006 for $portfolio(\theta, \boldsymbol{G})$, the Sharpe ratio started to decrease and dropped below the $MVP(\theta, \boldsymbol{G})$ level (at high levels of θ). Refer to Figure 43, below.

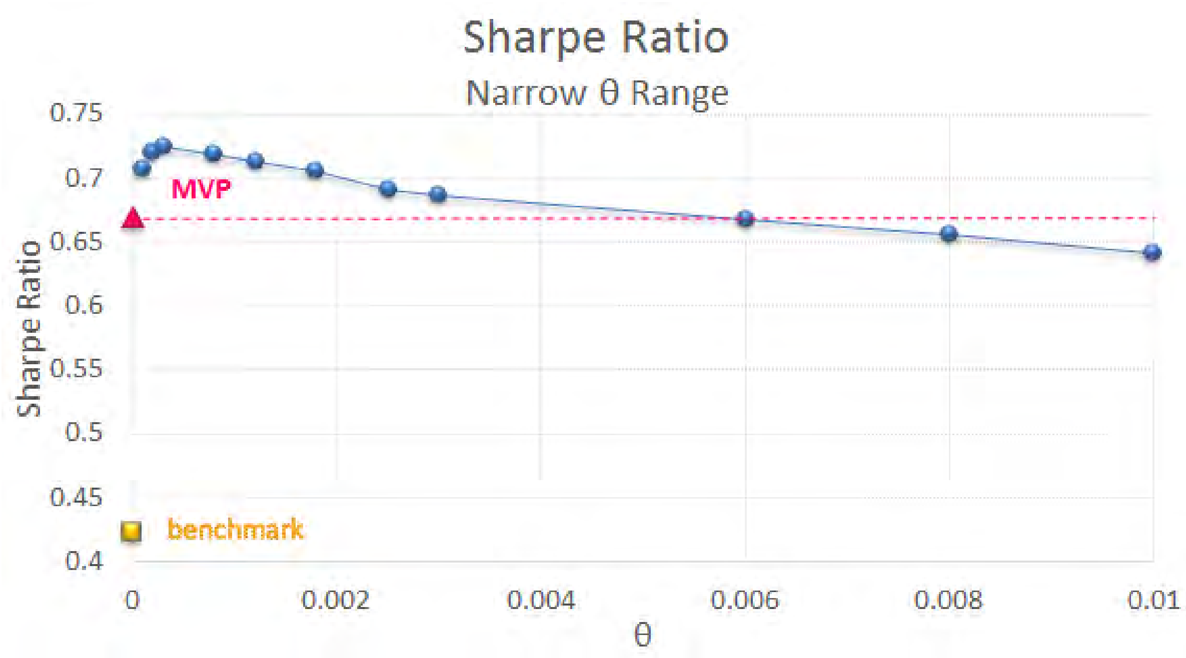

Figure 43 Annualised risk-adjusted return measured over narrowed θ range for $portfolio(\theta, \boldsymbol{G})$

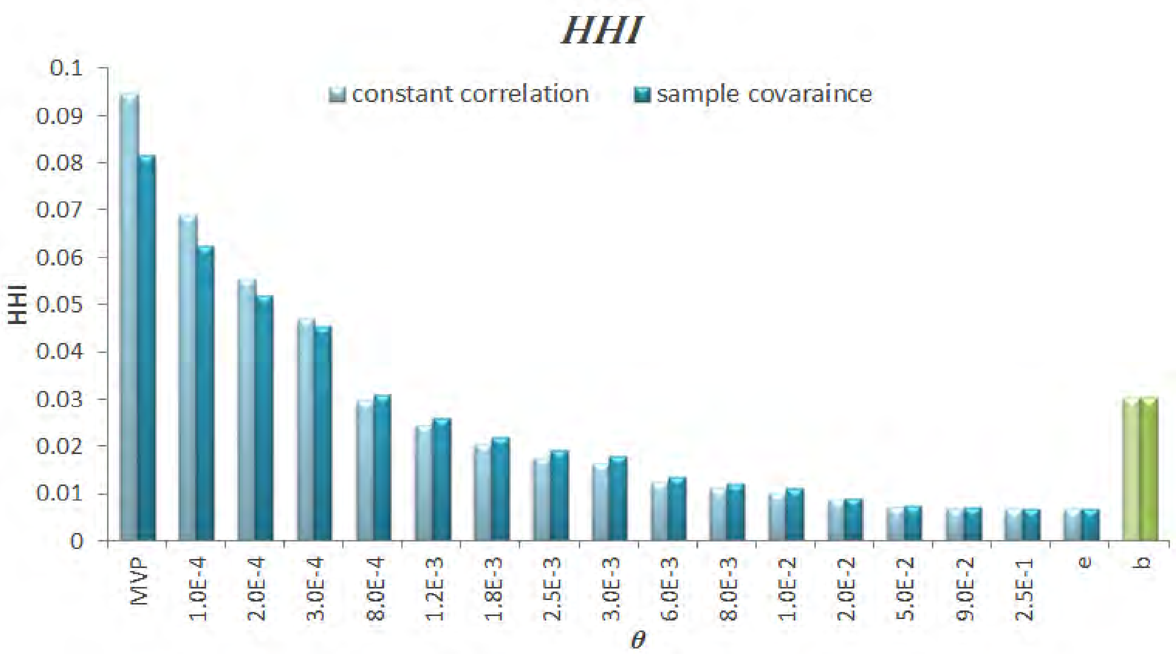

Figure 44 HHI for combined θ range

The average HHI was measured for the subset of θ portfolios over time and is shown in Figure 44, above. At higher levels of θ, the portfolios quickly reduce their overall concentration and approach that of the equally weighted portfolio (e). The benchmark (b) in comparison to MVP is fairly diverse, since its HHI measurement is about six times lower. At low levels of θ (θ < 3.0E-4), there is a significant reduction in concentration levels in comparison to the MVP. So although the portfolios generated at very fine levels of θ were fully equitised holding just 40 stocks, there was more diversity in the composition. The HHI for $portfolio(\theta, \boldsymbol{G})$ is higher than $portfolio(\theta, \boldsymbol{S})$, where $\theta \leq 0.0008$, indicating that at lower levels of θ, portfolios formulated under $\boldsymbol{G}$ are generally more concentrated then under $\boldsymbol{S}$. This reverses as θ increases beyond 0.0008; however this can be attributed to the risk parameter making less of an impact on the optimal solution.

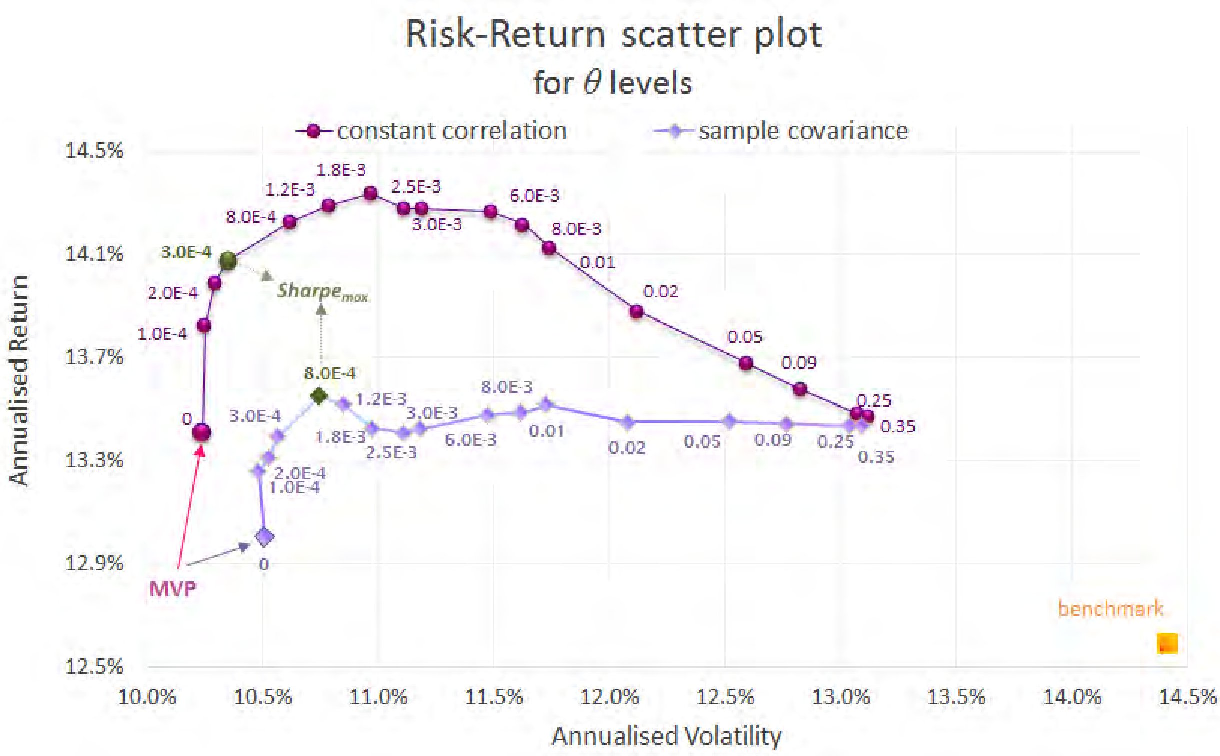

Figure 45 Risk-return scatter plot for various θ levels

The risk-return plot in Figure 45 above, clearly highlights the benefit of employing a shrinkage method ($\boldsymbol{G}$) over utilising a simple covariance matrix structure ($\boldsymbol{S}$). But the general pattern in the risk-return plot is similar for both covariance structures. At lower levels of θ, portfolios formulated benefit from both higher returns and lower risk. As θ increases however, there is a general convergence between the methods and a decrease in return and an increase in volatility.

Figure 46 below, indicates an inverse relationship between the level of volatility and HHI. Lower levels of volatility are observed for portfolios with higher levels of concentration. As HHI and thus concentration decrease and portfolios approach the equally weighted portfolio, volatility levels increase.

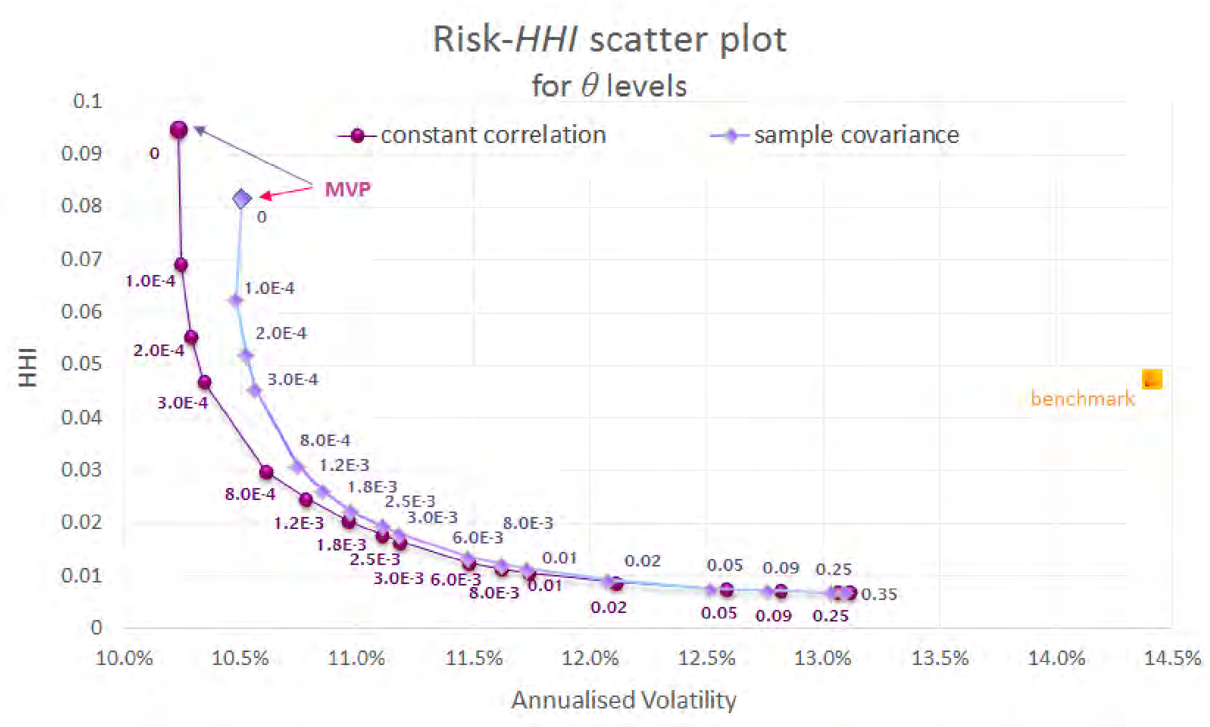

Figure 46 Risk-HHI plot for various θ levels

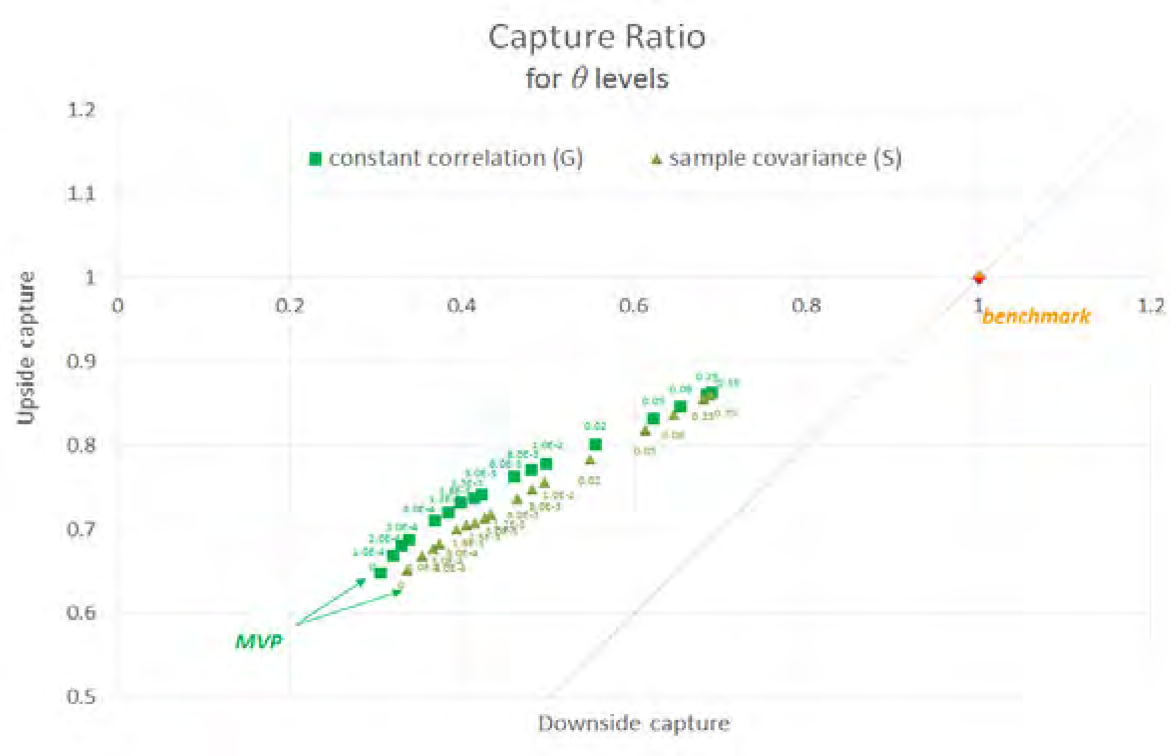

Figure 47 Capture Ratio - θ Levels

The capture ratio in Figure 47 indicates that when the benchmark is up, the θ portfolios generated under both covariance structures only participate partially during upturns, capturing between 64% (MVG_G) and 86% ($\theta = 0.35$). The downside capture ratio indicates that the tilt lowers risk, since the portfolios only capture between 30% and 70% in market drawdowns. Shrinkage is beneficial across all θ levels. Refer to all capture statistics in Appendix A, Table 14 and Table 15.

Turnover levels observed are lower in tilted portfolios - refer to Appendix B, Table 19. Shrinkage plays a more significant role in lowering resultant turnover in comparison to the standard sample covariance.

The relative maximum drawdown in Figure 48 below which measures the maximum drawdown differential between portfolios formulated using $\boldsymbol{G}$ and portfolios using $\boldsymbol{S}$, indicate that employing shrinkage produces portfolios with lower drawdown levels across the board. The MVP_G produced the smallest drawdown, this value increased as θ increased beyond 0 however. Refer to Appendix C, Table 23 and Table 24.

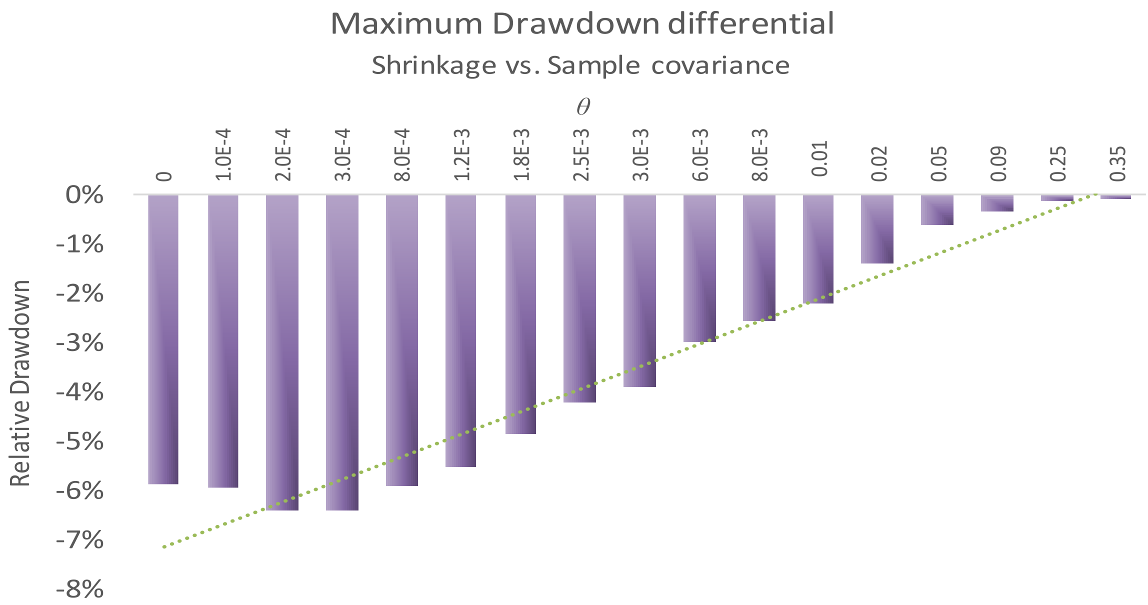

Figure 48 Relative drawdowns between the $\boldsymbol{G}$ and $\boldsymbol{S}$ for various θ levels

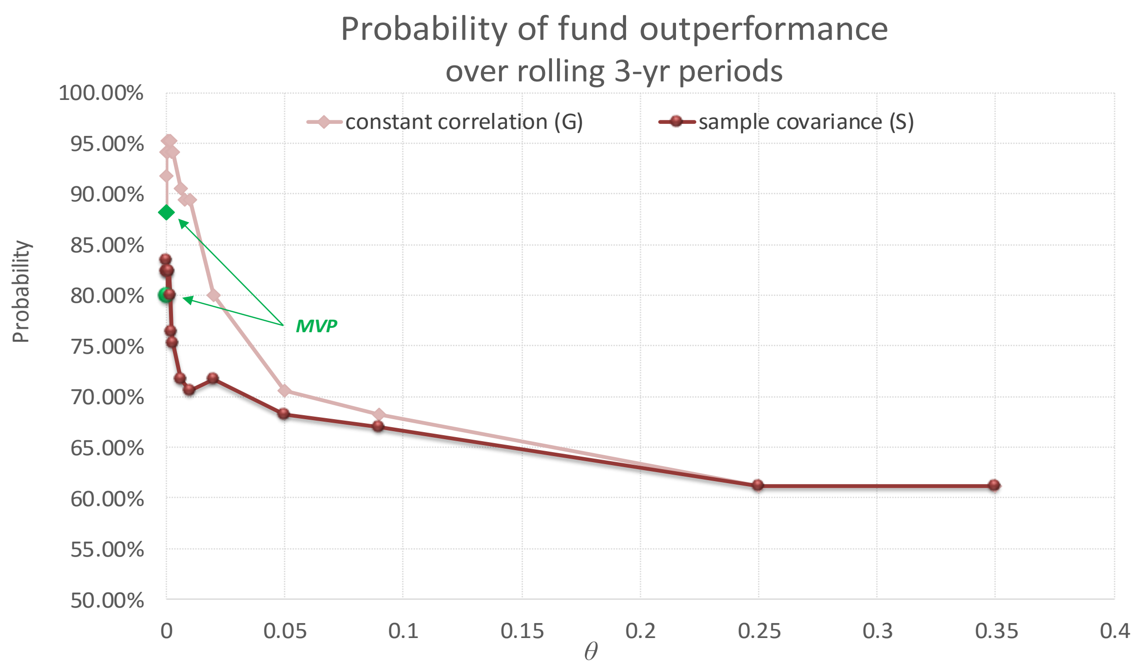

Figure 49 Probability of outperformance for θ various levels

The benefit of the tilt is clear from Figure 49 above, in that the portfolios have a high probability of outperforming the benchmark over three-year rolling periods. As θ increased, this percentage dropped significantly and approached 60%. Employing shrinkage is beneficial across all levels of θ.

5.2.3 Estimation of η (Active Distance Tilt)

Similar techniques employed in the section 5.2.1 above were used to determine the most optimal η ranges.

5.2.3.1 Broad η Range

The composition of the portfolios give an indication of the concentration or spread of holdings and are shown in Figure 50 below.

The MVP (η = 0) is once again the most concentrated portfolio. By adding even a small *tilt to the benchmark* (η > 0), the concentration is reduced by a significant margin. The most notable outcome is that at low levels of η, the largest 40 holdings account for a smaller proportion of the portfolio. As η increases, this number also increases, but slightly. This indicates that increasing the tilt to the benchmark may in fact decrease the diversity of the portfolio. Furthermore, it indicates that there may be benefit in narrowing the range over which to search for optimal levels of η, if portfolio diversification is one of the key goals.

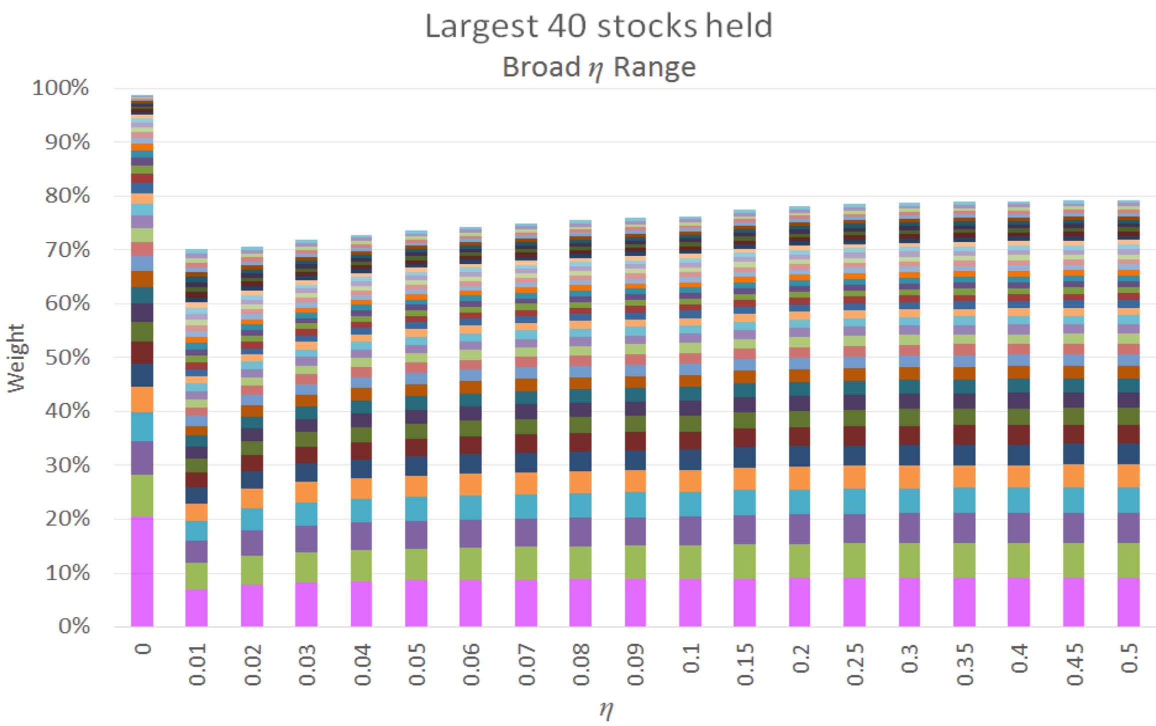

Figure 50 Weight distribution of largest 40 holdings for $\eta \geq 0$

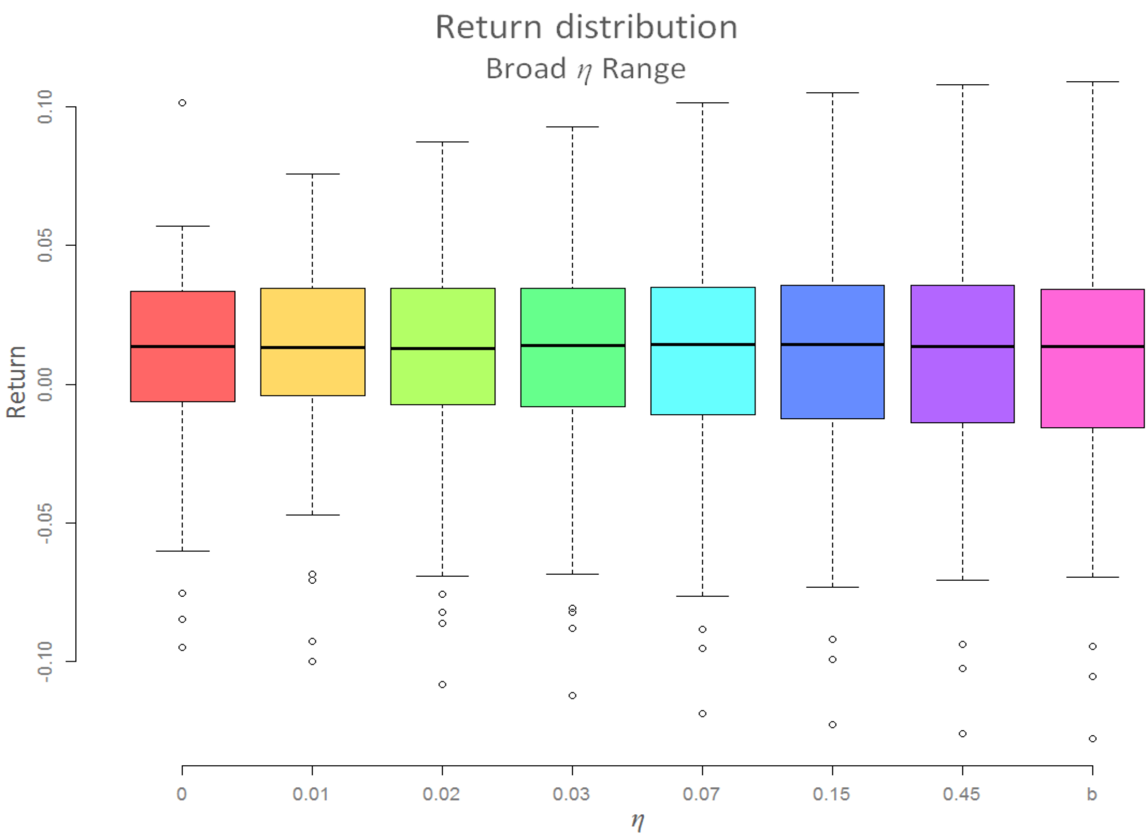

Figure 51 Box-plot of returns across a subset of η levels

An interesting aspect highlighted in the return distribution across portfolios (Figure 51 above), represented by the subset in Table 11 below, indicated that although the tilted portfolios have outperformed both the benchmark and MVP (η = 0) portfolios over the period - as η increased so does the dispersion of returns. At higher levels, the η portfolios suffered bigger drawdowns than the MVP. The range widened and there were far more distinct outliers as η increased.

Table 11 Broad η levels

η	0.01	0.02	0.03	0.07	0.15	0.45

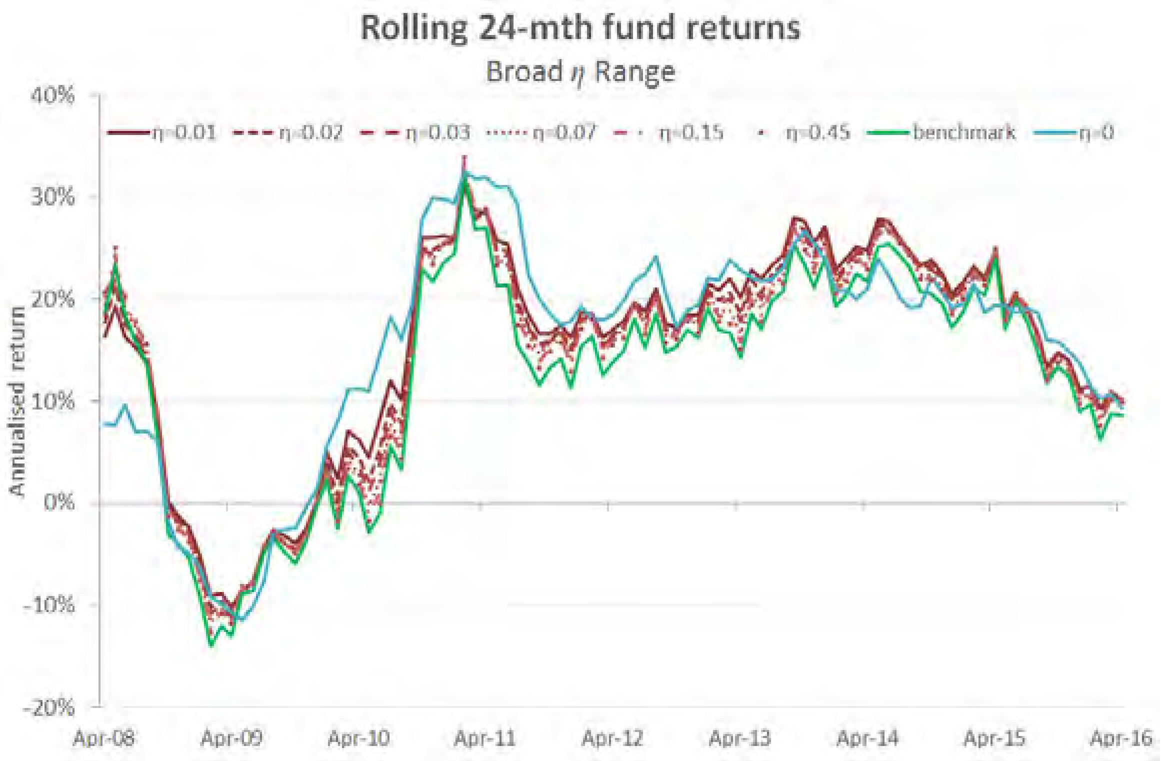

Figure 52 Rolling 24-month returns – Broad η range

The drawdown is evident in Figure 52 as well. The returns for the most part are bound between the benchmark and the MVP. Over the period June 2013 to April 2015, the benefit of holding any portfolio *tilted to the benchmark* was evident. All portfolios outperform the MVP and benchmark. Although the MVP is most distinctively different from the other portfolios, the return profiles are positively correlated.

Figure 53 Rolling 24-month volatility – Broad η range

The second point examined was risk. This is shown in Figure 53, above. There is a more significant difference in the annualised risk over the period in comparison to the returns. The η portfolios are far riskier than the MVP and approach the benchmark portfolio at higher levels of η.

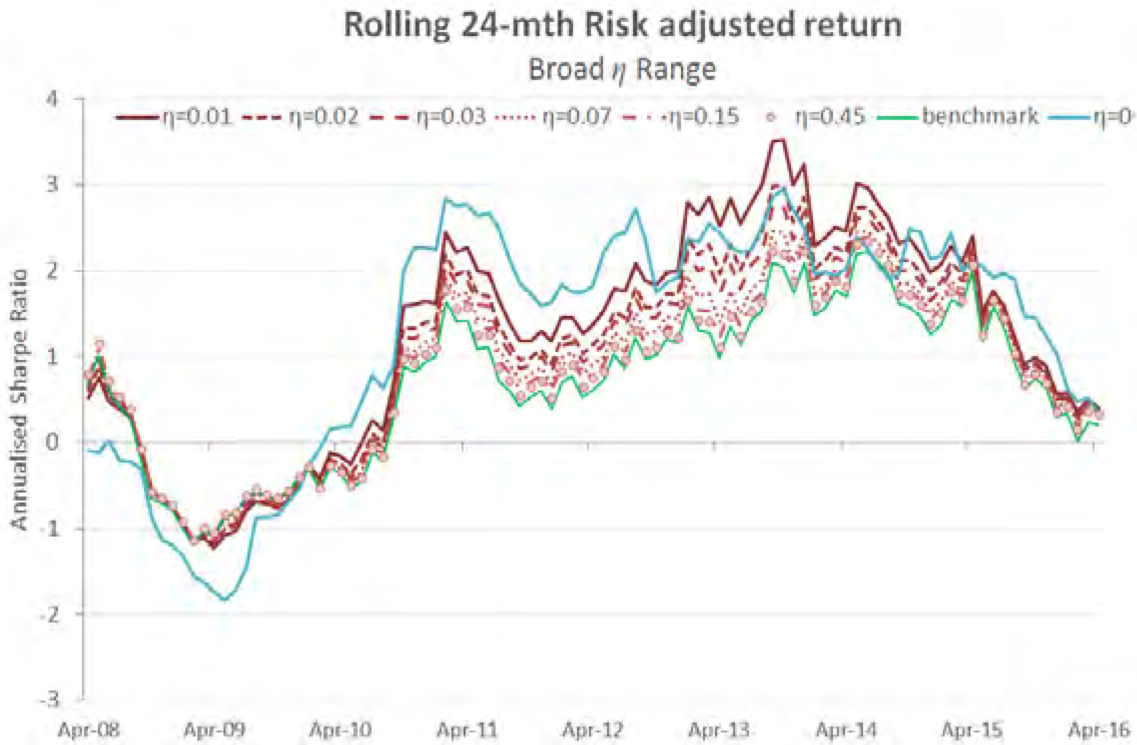

Figure 54 Rolling 24-month Sharpe Ratio – Broad η range

The impact of increased risk is highlighted in Figure 54, above. The 24-month rolling risk-adjusted returns measured for η portfolios is lower on average than that of the MVP. The effect of risk does impact and change the ranking of portfolios in terms of absolute performance.

5.2.3.2 Narrow η Range

Since the lower η levels performed better in most regards, the next step involved narrowing the range to determine the most optimal levels of η in the finer range.

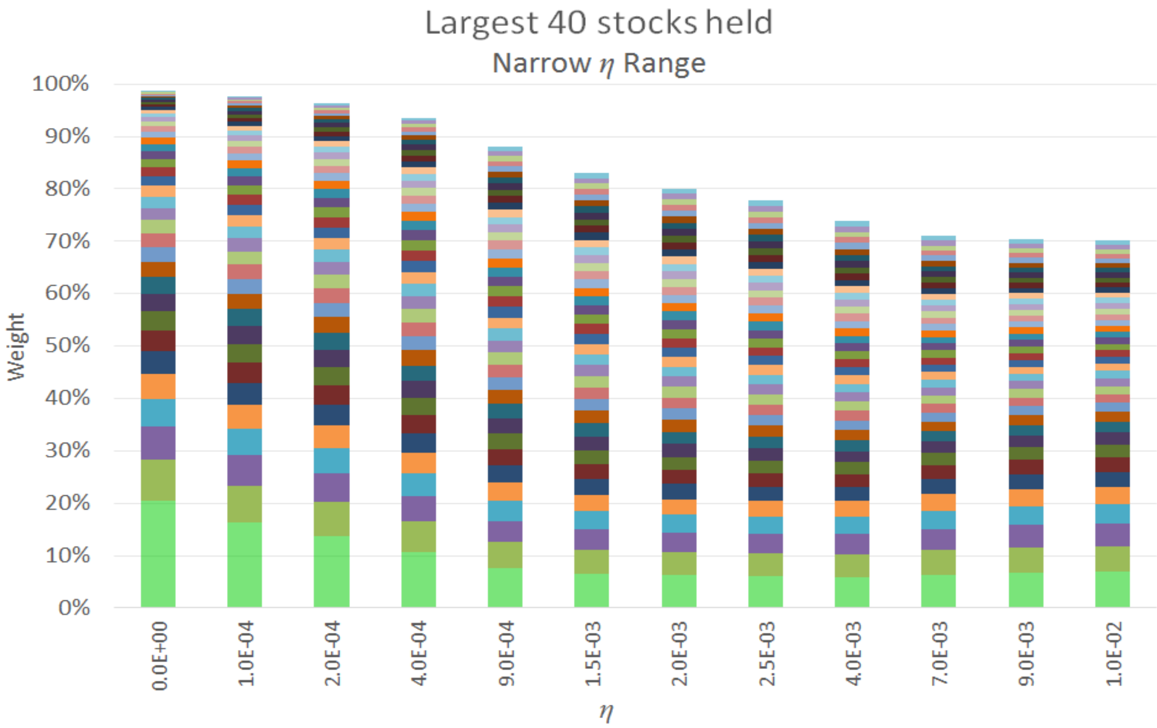

Figure 55 Largest 40 holdings for the Narrow η range

There is an increase in diversification as η increases beyond 0. As depicted in Figure 55 above, the largest 40 holdings of the portfolios formulated at lower levels of η, constituted close to 100% of the portfolio. Although this is the case, the introduction of the tilt increases diversity in the holdings. For η > 2.0E-3, exposure to the biggest stocks dropped to less than 80%, and continued to decline as η approached 0.01.

The cumulative return in Figure 56 below indicates that there is benefit in increasing η to levels beyond 3.0E-3. These portfolios start outperforming and return the highest overall returns.

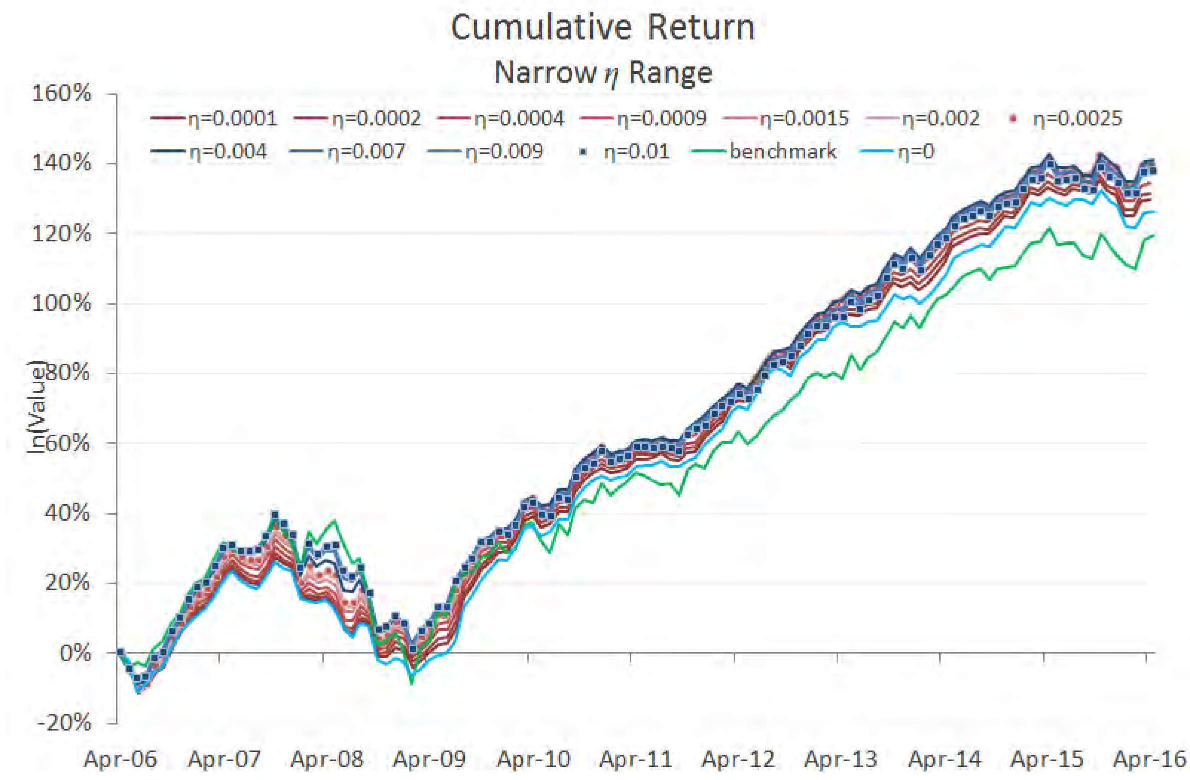

Figure 56 Cumulative Return – Narrow η Range

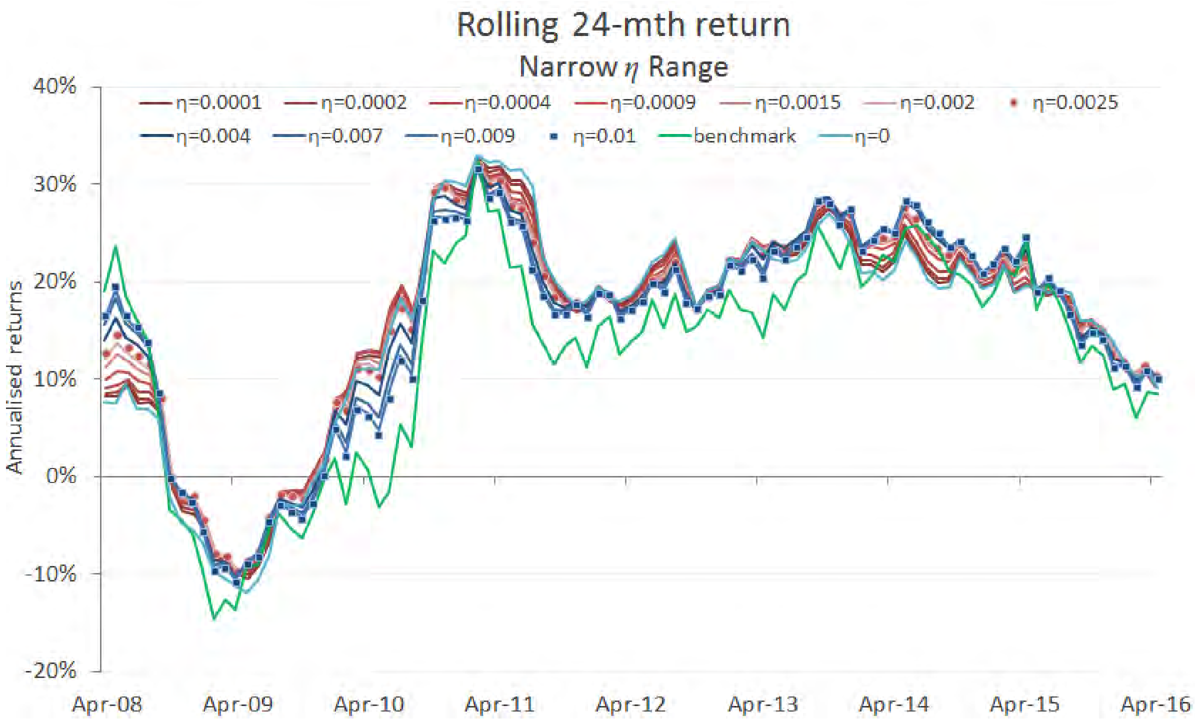

Figure 57 Rolling 24-month returns – Narrow η range

The spread of 24-month rolling returns across all portfolios was very similar (Figure 57).

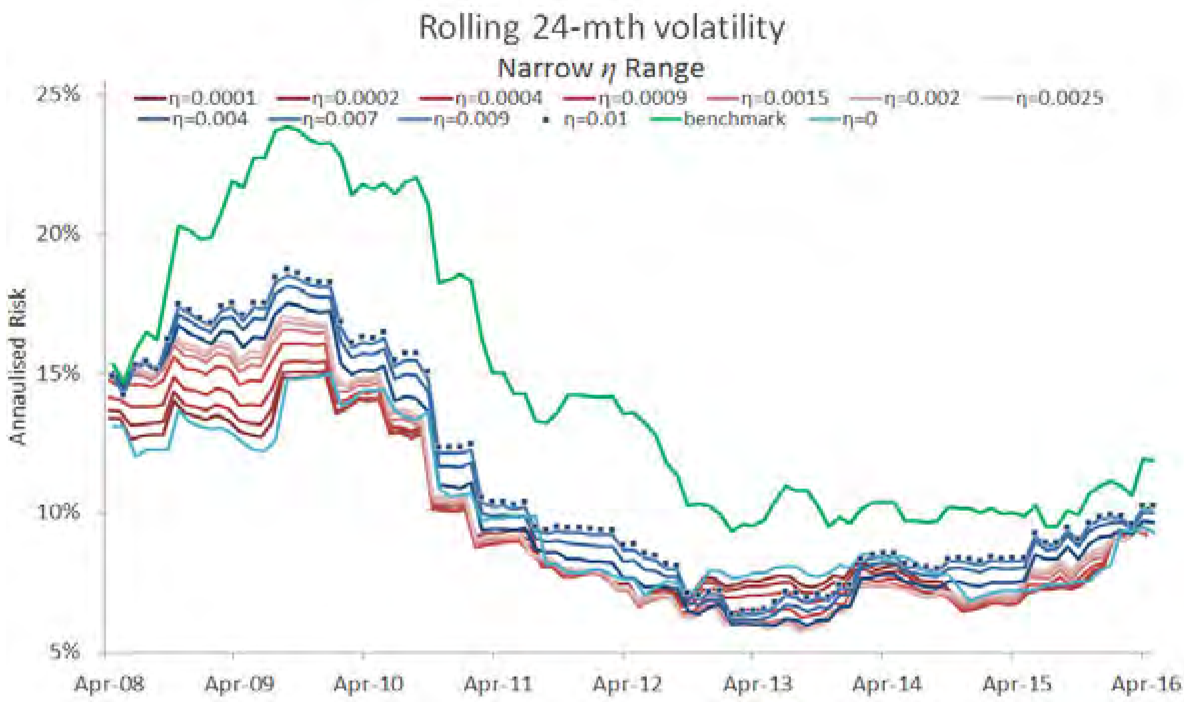

Figure 58 Rolling 24-month risk – Narrow η range

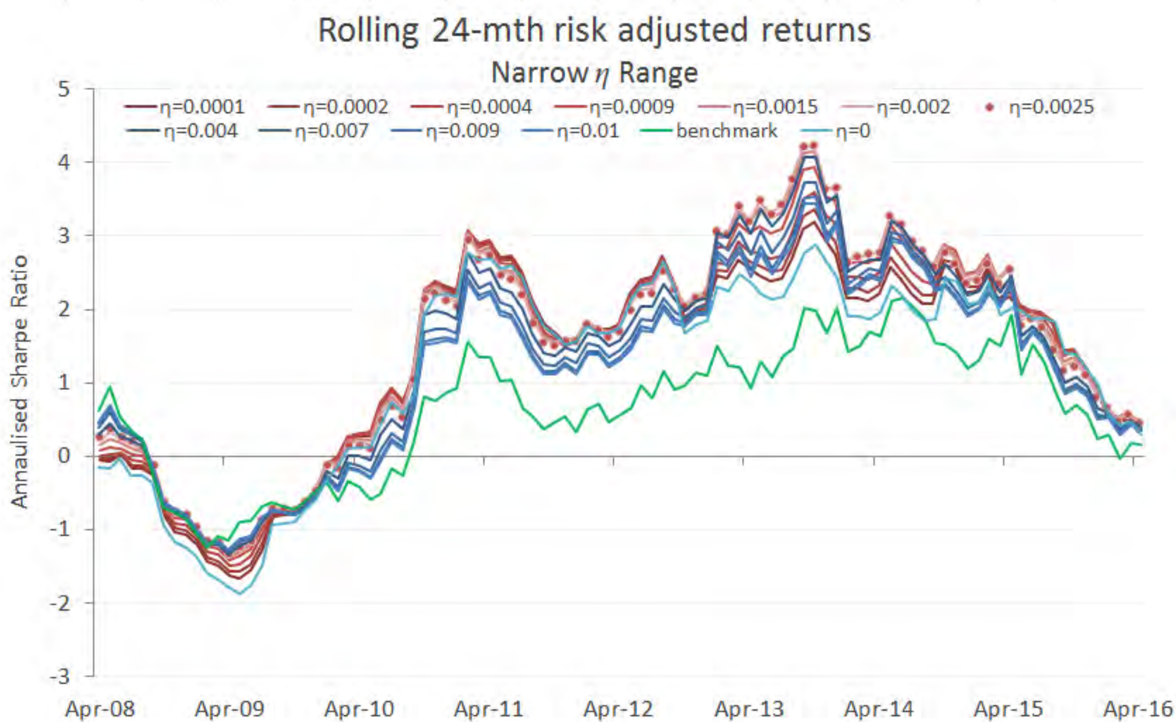

Figure 59 Rolling 24-month Sharpe Ratios– Narrow η range

Whilst the rolling risk measures are also alike, over certain periods they become more differentiated - for example, April 2008 to April 2010 (Figure 58 above). Over this period, the benchmark risk level was fairly high and the MVP was fairly low -

this would explain the dispersion in the series. There was a reduction in the risk-spread between the benchmark and MVP, as risk levels decreased over the latter period. An interesting point to note is that over June 2012 to June 2013, the η portfolios measured risk was lower than both the benchmark and MVP.

The risk reduction had most benefit on the Sharpe ratios, as seen in Figure 59, above. The risk-adjusted returns were on average higher than the MVP and benchmark. All portfolios tended to have a positive correlation with each other. When there was a directional change in the risk or return series, all portfolios tended to follow the same trajectory.

5.2.3.3 Outline of Results - η Level

The primary outcomes are outlined below and focus on the key aspects below (ranked in order of importance):

1. Highest return
2. Highest risk-adjusted return
3. Lowest risk

The returns were measured over the total period and are highlighted in Figure 60, below.

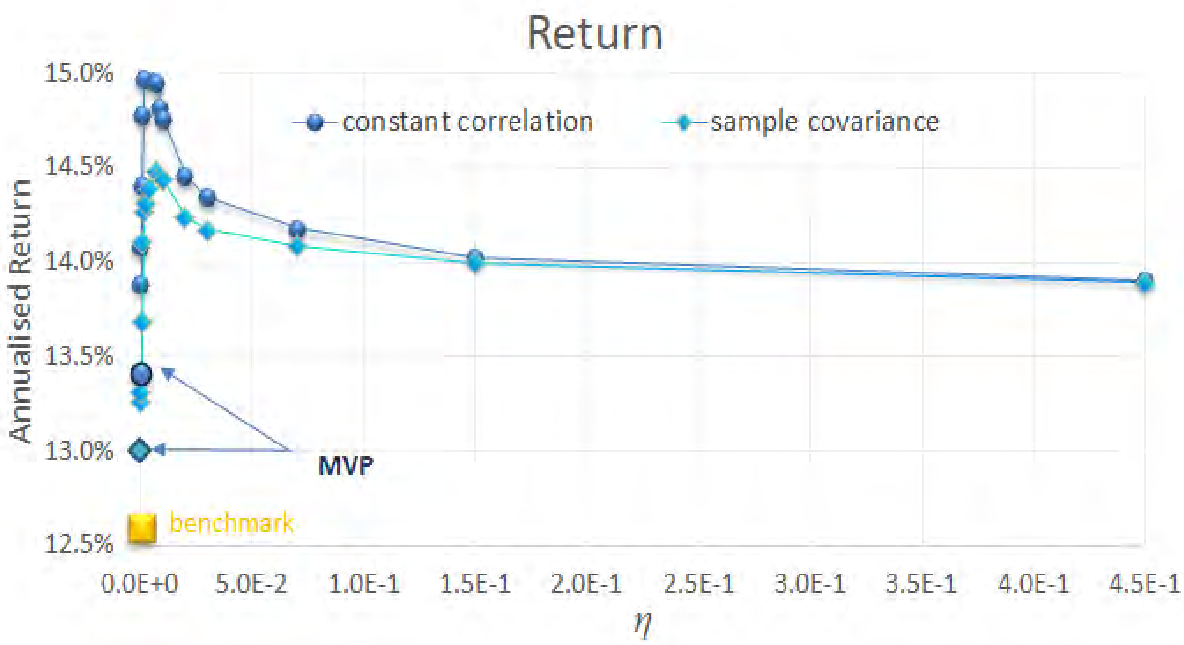

Figure 60 Annualised return measured over combined η range

Introducing the active distance tilt has the advantage of increasing the returns across all levels of η. This has resulted in the outperformance of both the MVP and benchmark over the total measurement period. The returns produced by $portfolio(\eta, \boldsymbol{G})$ were superior to returns produced by $portfolio(\eta, \boldsymbol{S})$. This is the general observation for η levels below 0.15. As the risk parameter's influence decreases with increasing levels of η, the returns tend to converge. The returns are non-monotonic, since for $portfolio(\eta, \boldsymbol{G})$, with η levels below 0.004, the returns increased and approached a maximum at η levels around 0.0018, before they declined sharply as η increased beyond 0.01. A similar trend is observed for $portfolio(\eta, \boldsymbol{S})$. As η increases and portfolios approach the SWIX 150 weightings, the returns decline but still outperform the MVP and benchmark.

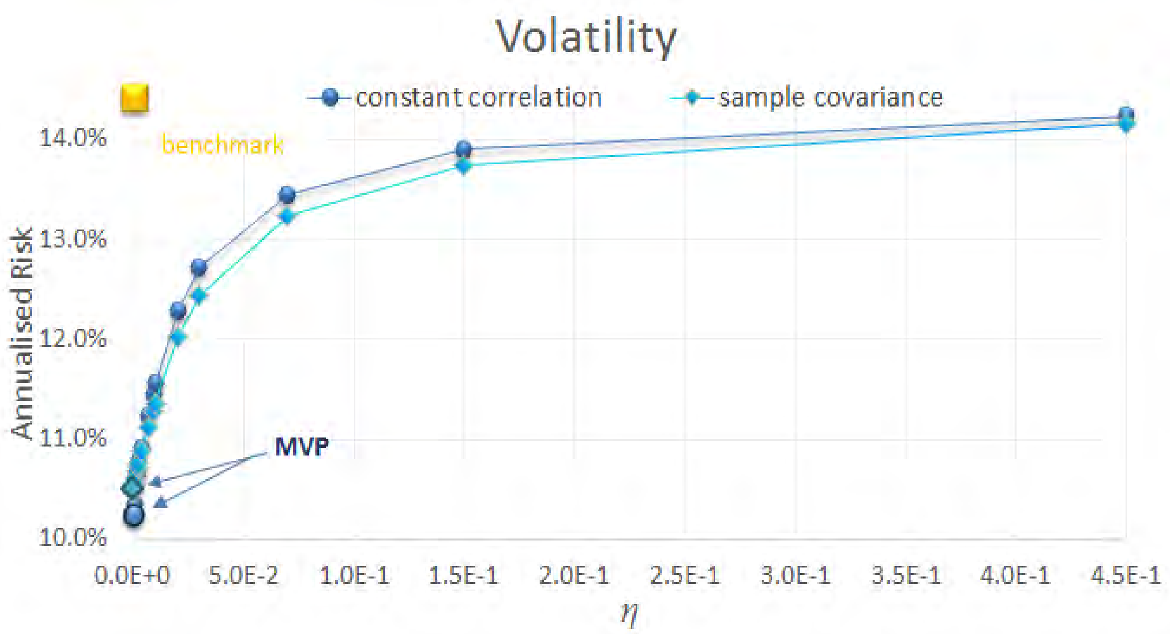

Figure 61 Annualised volatility measured over combined η range

In Figure 61, the volatility increases monotonically but at higher levels of η, portfolio weightings quickly approached their SWIX 150 weightings.

The lowest risk is observed for the $MVP(\theta, \boldsymbol{G})$ (the MVP formulated under $\boldsymbol{G}$), but as η increases, there is a reversal in the rank of risk measured for $portfolio(\eta, \boldsymbol{G})$, which returns higher levels of risk relative to the risk returned from $portfolio(\eta, \boldsymbol{S})$. At high levels of η, the risk measured under both constructs is relatively similar.

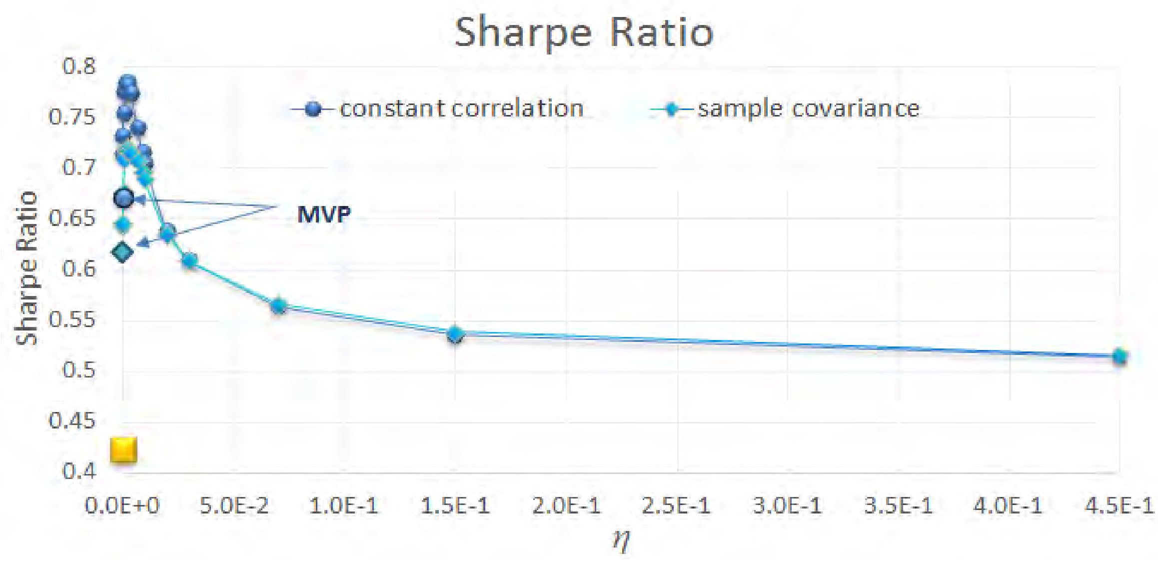

Figure 62 Annualised risk-adjusted returns measured over combined η range

Figure 62 above and Figure 63 (below) highlights the above relationship. Since realised risk monotonically increases and rises at a faster rate than return as η increases, the maximum Sharpe ratio is attained at a lower η threshold. The MVP benefitted from lower realised risk and was better on a risk-adjusted basis than the portfolios generated at η levels greater than 0.015 for $portfolio(\eta, \boldsymbol{G})$ and 0.002 for $portfolio(\eta, \boldsymbol{S})$.

Finally, the relationship between $portfolio(\eta, \boldsymbol{G})'s$ risk-adjusted return (measured by the Sharpe ratio) and η can be represented as follows -

$$Sharpe_i = \begin{cases} f(\eta_{i+1}) > f(\eta_i) & 0 < \eta_i < 0.002 \\ f(\eta_{i+1}) < f(\eta_i) & \eta_i > 0.002 \end{cases} \qquad 5.4$$

Where $Sharpe_i$ represents the Sharpe ratio for η_i where $i = 1, 2, \dots n$.

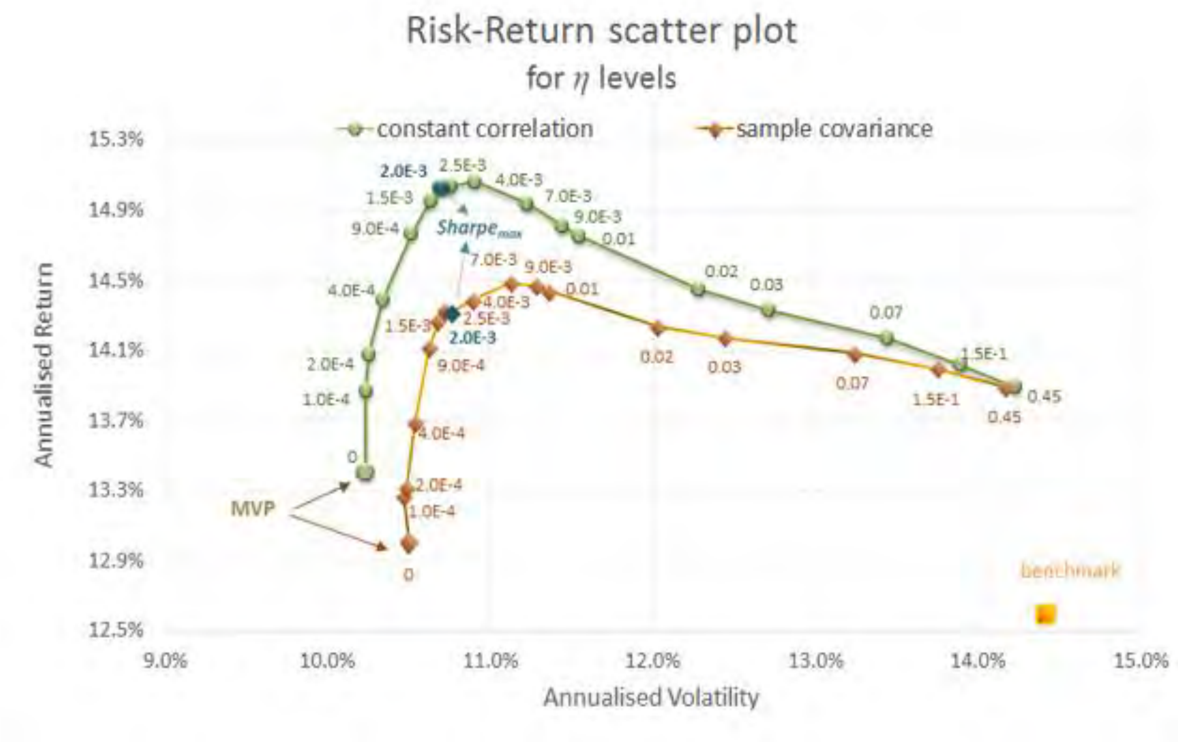

Figure 63 Risk-return scatter plot for various η levels

For more detail around other relevant statistics refer to 8.1-8.5 below.

This section provided an in depth analysis of the results obtained from the shrinkage parameter estimation and preference parameter estimation processes respectively. The results generally support the use of factor tilts to enhance the returns of the MVP, in addition employing covariance shrinkage further enhances portfolio statistics across all portfolios, in contrast to the standard covariance matrix. These results are summarised in section 6 below.

6. Conclusion

A plethora of research is available on the benefits of minimum variance investing in both the international and South African environment. However, some concerns around this, which include highly concentrated portfolios, illiquidity and estimation error maximisation, still exist. This study investigated theoretically sound yet pragmatic methods that could be employed by investment practitioners, which would enhance the returns of the MVP and address some of these key concerns. This was carried out through enhancements to the objective function formulation - which included tilts to factors that have historically exhibited attributes that have enhanced portfolio performance, as well as through improved parameter input estimation by means of covariance shrinkage. The MVP tilt enhancement employed is academically sound, in line with capital market theory and improves on more ad hoc approaches which impose pre-existing views on the overall process. These views will inadvertently introduce hindsight or anchoring bias. Some of these judgemental approaches include blending the MVP with certain factor portfolios as well as the introduction of additional constraints into the framework which would mitigate the effect of extremities on the optimal portfolio.

The key findings are summarised below:

6.1 A Note on Absolute and Risk-Adjusted Performance

6.1.1 Enhancement through Average Volume Indicator

Adding the enhancement through the average volume indicator, results in an outperformance of the MVP (for portfolios formulated at lower levels of the tilt), on both an absolute and risk-adjusted basis. There is a decay in the return, as portfolios approach the volume-weighted portfolio, however all portfolios outperform the benchmark regardless of the level of tilt employed.

Over rolling two-year periods, there was an improvement in the Sharpe Ratio for the majority of the period.

6.1.2 Enhancement through diversification (Herfindahl–Hirshman Index)

Introducing the enhancement through diversification, results in improved absolute performance across all enhanced portfolios - surpassing both the MVP and benchmark over the entire measurement period. The relationship between the level of tilt and return is non-monotonic, at lower levels - returns increase until a *sweet-spot* is attained. Beyond this level however, returns start to decrease as portfolios approach the equally weighted portfolio.

The risk-adjusted returns follow a similar pattern described in the paragraph above. The important differentiator however is that at higher levels of the tilt, the Sharpe Ratio degenerates rapidly and drops below that of the MVP.

Over rolling two-year periods, there was an improvement in the Sharpe Ratio for the majority of the period.

6.1.3 Enhancement through active distance

Similar results to those described in 6.1.2 above are obtained across all portfolios formulated using sequential η levels.

All portfolios including the MVP outperform the benchmark.

6.2 A Note on Volatility, Concentration, Turnover and Drawdowns

The MVP produced the lowest realised volatility. This was expected since the optimiser prioritises risk minimisation and negates all other impacts.

All portfolios, regardless of the enhancement employed, were less risky than the benchmark over the total measurement period. The same outcome was observed over rolling two-year periods; where the realised risk was consistently lower than the benchmark.

The active risk measured over rolling two-year periods was substantially lower than the MVP.
The concentration levels of portfolios employing the diversification tilt were lower in comparison to the MVP, regardless of the covariance structure used. This was evident from the distribution of portfolio weights, which was far less concentrated and the lower HHI. Supplementary to this however, as the HHI and thus concentration decreased and portfolios approached the equally weighted portfolio, volatility levels escalated. Other enhancements produced similar results.
Turnover directly impacts costs. Thus in general, strategies that require lower levels of turnover tend to be more desirable. All enhancements introduced lowered turnover levels relative to the MVP. As the tilt level increased, a decrease in turnover was observed, across all portfolios, in comparison to the MVP.
The MVP recorded the smallest drawdown in comparison to other portfolios. As the level of tilt employed increased, the drawdowns worsened monotonically. The benchmark returned the worst drawdown over the measurement period.

Whilst it may be argued that a constrained minimum-variance approach would be a more suitable benchmark, and that all the tilts end up achieving a decrease in concentration in the portfolio, and that the decrease in concentration rather than the tilt is responsible for the improved metrics across all the tilts investigated, it is worth noting that constraining the portfolio would introduce a bias to the optimisation framework. There have been studies that introduce constraints to the problem, to deal with the concentration issue. But this will impose a pre-existing view on the overall process which will inadvertently introduce hindsight or anchoring bias. (Refer to King, 2007)
The HHI tilt, by its very formulation (definition) prioritises an equal weighting which represents the most diversified mix of assets. This will directly impact the final portfolio weightings and tilt it toward a more diversified asset mix, since you are incorporating an element (equally weighted factor value) around which the final portfolio weights are determined.
Over certain rolling periods the tilted portfolios underperformed the MVP. The reduction in concentration, does seem to benefit the portfolios, but the degree

of the outperformance or underperformance over the various periods are variable ,this indicates that there are other factors at play which are affecting overall fund performance between the different tilts.

6.3 General Observations on Covariance Shrinkage

The use of covariance shrinkage is beneficial across all the enhancements introduced to the MV optimisation framework. Both absolute and risk-adjusted returns are superior in comparison to the returns obtained when the standard sample covariance matrix is employed.

The MVP formulated through shrinkage resulted in the lowest realised volatility across all portfolios generated using both covariance structures.

An interesting observation however was the relationship between the level of the respective tilt and volatility. At lower levels of the tilt, the risk of portfolios formulated using shrinkage were lower than for those formulated under the standard covariance matrix. As this level increased however, this relationship reversed and as the tilts were prioritised over risk minimisation, portfolios generated under shrinkage became marginally riskier. Turnover levels reduced when shrinkage was employed.

Drawdowns were less severe when shrinkage was used - the benefit was far greater at lower levels of the tilt, where risk played a greater role in the optimisation process.

Portfolio participation in different market phases was improved when shrinkage was employed. The downside capture ratio was less extreme in market drawdowns and portfolios had a greater participation in market recoveries in comparison to portfolios formulated under the sample covariance matrix.

6.4 Limitations and directions for further research

The outcome of the analysis indicated that the technique has greater feasibility and merit for specific levels of the tilt only and a general decay in performance is noted as exposure to the factor is increased.

The study also covers a limited period and is only tested on the FTSE-JSE SWIX Index constituents, which dates back to January 2002. In order to comprehensively assess the validity of the hypotheses introduced, one could explore other individual or combinations of asset classes or complete a similar study on the FTSE-JSE ICB[4] Industry or Sectors, and devolve the period back to June 1995, when the ALSI was established. Additionally, other jurisdictions could be examined. The Yanushevsky & Yanushevsky (2015) study was limited to the ten largest holdings of the US Vanguard Windsor II Fund and King (2007) focused specifically on regional developed equity markets.

The methodology is innovative in that it offers a platform to introduce multiple factor tilts to the MV optimisation framework which could be incorporated on a stand-alone basis or in combination with other factors.
However this research doesn't rule out the possibility that the outperformance could to some extent be caused by:

- The outperformance could be due to sampling effects
- The outperformance could be explained by not controlling for systematic risk.

With the rise of *smart beta* products, this enhanced framework could be used to improve the product offering of suppliers who base products on MV optimisers. Further research could include an evaluation of the enhanced framework examined in this study in comparison to more passive indexing alternatives.

Furthermore, the reduction in turnover in comparison to the MVP was generally observed across all enhancements. The impact of this on a net-fee basis was not examined since different tax regimes are generally applied to institutional investors like pension funds in contrast to individual investors. This could be further investigated however, since lower *portfolio churn* [5] directly lowers trade costs which will enhance net-of-fee returns more extensively.

The outcomes of this study are generally in line with and support both existing international and domestic research by Yanushevsky & Yanushevsky (2015), Ledoit and Wolf (2004), Munro and Bradfield (2016) and King (2007), to mention a few. Firstly, the introduction of the $J\ Index$ creates an innovative platform through which

[4] Industry Classification Benchmark
[5] Rate of turnover of portfolio

multiple factor tilts can be introduced and tested in a multicriteria MVP optimisation framework. Secondly, it was found that the methodology can be further augmented by employing a shrinkage technique like the constant correlation model in estimating the covariance structure. And finally, there is benefit in including factor tilts like - volume, concentration and active distance - since improved performance was observed over specific degrees of the factor tilt, on both an absolute and risk-adjusted basis, over both the MVP and benchmark.

7. References

Anagnostopoulos, K. & Mamanis, G. 2010. A portfolio optimization model with three objectives and discrete variables. *Computers & Operations Research.* 37(7), 1285-1297.

Bacon, C. 2004. *Practical Portfolio Performance Measurement and Attribution.* 2nd edition. London, Wiley.

Baker, N.L. & Haugen, R.A. 2012. Low risk stocks outperform within all observable markets of the world. Available at: http://ssrn.com/abstract=2055431.

Baks, K., Busse, J.A. & Green, T.C. 2007. Fund Managers Who Take Big Bets: Skilled or Overconfident. Available at: http://ssrn.com/abstract=891727.

Belsley, D., Kuh, E., and Welsch, R. 1980. Regression diagnostics: Identifying influential data and sources of collinearity. New York, John Wiley.

Best, M.J. & Grauer, R. 1991. On the sensitivity of mean variance efficient portfolios to changes in asset means: Some analytical and computational results. *Review of Financial Studies.* 4(2), 315-342.

Black, F. & Litterman,R. 1990. Asset Allocation: Combining Investors Views with Market Equilibrium. Fixed Income Research. Goldman, Sachs & Company.

Black, F. & Litterman, R. 1992. Global Portfolio Optimization. *Financial Analysts Journal.* 48(5), 28-43.

Bonini, S., Zanetti, L., Bianchini, R. & Salvi, A. 2010. Target Price Accuracy in Equity Research. *Journal of Business Finance & Accounting.* 37(9-10), 1177-1217.

Bouchaud, J.P., Potters, M. & Aguilar, J.P. 1997. Missing Information and Asset Allocation. Available at: http://xxx.lanl.gov/cond-mat/9707042.

Bradfield, D., Maritz, J. & Swartz, J. 2005. Managing Tracking Error - from Mandate to Active Bet. Research Report, Cape Town, Cadiz Financial Strategists Research.

Bradfield D. & Swartz, J. 2002. Demonstration of the Black-Litterman Portfolio Design Procedure in the South African Setting. Research Report, Cape Town, Cadiz Financial Strategists Research

Bradshaw, M.T. 2004. How Do Analysts Use Their Earnings Forecasts in Generating Stock Recommendations? *The Accounting Review.* 79(1), 25-50.

Britten-Jones, M. 1999. The sampling error in estimates of mean-variance efficient portfolio weights. *The Journal of Finance.* 54(2), 655-671.

Carhart, M.M. 1997. On Persistence in Mutual Fund Performance. *The Journal of Finance.* 52(1), 57-82.

Chen, X. & Lai, Y. 2015. On the concentration of mutual fund portfolio holdings: Evidence from Taiwan. *Research in International Business and Finance.* 33, 268-286

Chopra, V.K. & Ziemba, W.T. 1993. The Effect of Errors in Means, Variances, and Covariances on Optimal Portfolio Choice. *Journal of Portfolio Management.* 19(2), 6-11.

Cremers, M.K.J. & Petajisto, A. 2009. How Active Is Your Fund Manager? A New Measure That Predicts Performance. *Review of Financial Studies.* 22(9), 3329-3365.

Crouch, R.L. 1970. The Volume of Transactions and Price Changes on the New York Stock Exchange. *Financial Analysts Journal.* 26(4), 104-109.

Ehrgott, M., Klamroth, K. & Schwehm, C. 2004. An MCDM approach to portfolio optimization. *European Journal of Operational Research.* 155(3), 752-770.

Elton, E.J. & Gruber, M.J. 1997. Modern portfolio theory, 1950 to date. *Journal of Banking & Finance.* 21(11-12), 1743-1759.

Elton, E.J. & Gruber, M.J. 1973. Estimating the Dependence Structure of Share Prices-Implications for Portfolio Selection. *Journal of Finance.* 28(5), 1203-1232.

Fabozzi, F.J., Kolm, P.N., Pachamanova, D.A. & Focardi, S.M. 2007. Robust Portfolio Optimisation. *Journal of Portfolio Management.* 33(3), 40-48.

Fama, E. & French, K. 1993. Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics.* 33(1), 3-56.

Fletcher, J. 2009. Risk Reduction and Mean-Variance Analysis: An Empirical Investigation. *Journal of Business Finance & Accounting.* 36(7-8), 951-971.

Frazzini, A., Friedman, J. & Pomorski, L. 2016. Deactivating Active Share. *Financial Analysts Journal.* 72(2), 14-21.

Godfrey, M.D., Granger, C.W.J. & Morgenstern, O. 1964. The Random-walk hypothesis of stock market behavior. *Kyklos.* 17(1), 1-30.

Grinold, R.C. & Kahn, R.N. 1999. *Active Portfolio Management.* 2nd edition. New York , McGraw-Hill.

Haugen, R. & Baker, N. 1991. The efficient market inefficiency of capitalization-weighted stock portfolios. *The Journal of Portfolio Management.* 17(3), 35-40.

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

Jobson, J.D. & Korkie, B. 1981. Putting Markowitz Theory to Work. *Journal of Portfolio Management.* 7(4), 70-74.

Jorion, P. 1985. International Portfolio Diversification with Estimation Risk. *Journal of Business.* 58(3), 259-278.

Kacperczyk, M., Sialm, C. & Zheng, L. 2005. On the Industry Concentration of Actively Managed Equity Mutual Funds. *The Journal of Finance.* 60(4), 1983-2011.

Karpoff, J.M. 1987. The Relation between Price Changes and Trading Volume: A Survey. *Journal of Financial and Quantitative Analysis.* 22(01), 109-126.

Khuzwayo, B. 2011. Diversification in Portfolio Construction: Can it lead to outperformance? Research Report. Cape Town, Cadiz Quantitative Research.

King, D. 2007. Portfolio optimisation and diversification. *Journal of Asset Management.* 8(5), 296-307.

Kraft, A., Leone, A., and Wasley, C. 2006. An analysis of the theories and explanations offered for the mispricing of accruals and accrual components. Journal of Accounting Research 44 (2), 297–339

Kopman, L. & Liu, S. 2009. Maximizing the Sharpe Ratio and Information Ratio in the Barra Optimizer. Research Report. New York, MSCI Barra Research.

Kruger, R, Strugnell, D and Gilbert, E., 2011. Beta, size and value effects on the JSE, 1994-2007. *Investment Analysts Journal.* 74, 1-17.

Kruger, R. & Van Rensburg, P. 2008. Evaluating and constructing equity benchmarks in the South African portfolio management context. *Investment Analysts Journal.* 67, 5-17.

Ledoit, O. & Wolf, M. 2003. Improved estimation of the covariance matrix of stock returns with an application to portfolio selection. *Journal of Empirical Finance.* 10(5), 603-621.

Ledoit, O. & Wolf, M. 2004. Honey, I Shrunk the Sample Covariance Matrix. *The Journal of Portfolio Management.* 30(4), 110-119.

Leone, A., J., Minutti-Meza, M., and Wasley, C. 2015. Influential Observations and Inference in Accounting Research. Simon Business School Working Paper No. FR 14-06. Available at SSRN: https://ssrn.com/abstract=2407967

Markowitz, H. 1952. Portfolio Selection. *The Journal of Finance.* 7(1), 77-91.

Massa, M. & Simonov, A. 2006. Hedging, Familiarity and Portfolio Choice. *Review of Financial Studies.* 19(2), 633-685.

Matoti, L. 2009. Building a statistical linear factor model and a global minimum variance portfolio using estimated covariance matrices. Unpublished thesis. University of Cape Town.

Michaud, R. 1989. The Markowitz Optimisation Enigma: Is 'Optimized' Optimal? *Financial Analysts Journal.* 45(1), 31-42.

Michaud, R.O. 2002. An Introduction to Resampled Efficiency. Available at: http://ssrn.com/abstract=2405555

Michaud, R.O. & Ma, T. 2001. Efficient Asset Management: A Practical Guide to Stock Portfolio Optimization and Asset Allocation. *Review of Financial Studies.* 14(3), 901-904.

Munro, B. & Bradfield, D. 2016. Putting the squeeze on the sample covariance matrix for portfolio construction. *Investment Analysts Journal.* 45(1), 47-62.

Oladele, O. & Bradfield, D. 2016. Low volatility sector-based portfolios: a South African case. *Orion.* 32(1), 55-78.

Panulo, B., 2014. Risk Parity and other Risk Based Portfolio Allocation Approaches in South African Equity Markets. Master's Thesis. University of Cape Town.

Peterson, B.G. & Carl, P. 2014. PerformanceAnalytics: Econometric tools for performance and risk analysis.

Sarkar, D. 2008. Lattice: Multivariate Data Visualization with R. New York, Springer.

Scherer, B. 2007. Can robust portfolio optimisation help to build better portfolios? *Journal of Asset Management.* 7(6), 374-387.

Scherer, B. 2011. A note on the returns from minimum variance investing. *Journal of Empirical Finance.* 18(4), 652-660.

Schlanger, T., Philips, C.B. & LaBarge, K.P. 2012. The search for outperformance: Evaluating 'Active Share'. Research Report. Valley Forge, Vanguard Research.

Sharpe, W.F. 1963. A Simplified Model for Portfolio Analysis. *Management Science.* 9(2), 277-293.

Skylogiannis, V. & Xu, J. 2009. Mean Variance Optimization and Beyond. Master's Thesis. Stanford University.

Stein, C. 1955. Inadmissibility of the usual estimator for the mean of a multivariate normal distribution. Proceedings of the 3d Berkeley Symposium on Probability and Statistics, Berkeley, University of California Press.

Walter, C. 2016. The financial Logos: The framing of financial decision-making by mathematical modelling. *Research in International Business and Finance.* 37(C), 597-604.

Turlach B, A., & Weingessel, A. (2013). Quadprog: Functions to solve quadratic programming problems. S original by Berwin, A. Turlach, R port by A. Weingessel.

Westerfield, R. 1977. The Distribution of Common Stock Price Changes: An Application of Transactions Time and Subordinated Stochastic Models. *Journal of Financial and Quantitative Analysis.* 12(05), 743-765.

Wilcox, D. & Gebbie, T. 2007. An analysis of cross-correlations in an emerging market. Physica A. 375 (2), 584-598.

Yanushevsky, R. & Yanushevsky's, D. 2015. An approach to improve mean-variance portfolio optimization model. *Journal of Asset Management.* 16(3), 209-219.

Ying, C.C. 1966. Stock Market Prices and Volumes of Sales. *Econometrica.* 34(3), 676-685.

8. Appendix

8.1 Appendix A

Table 12 Capture Ratio Statistics[6] for portfolio formulated under covariance structure, G for various λ levels

portfolio(λ,G)	Up Capture	Down Capture	Up Number	Down Number	Up Percent	Down Percent
0	64.6%	30.7%	83.5%	65.9%	32.9%	90.2%
1.0E-4	66.6%	32.4%	84.8%	63.4%	34.2%	90.2%
2.0E-4	67.7%	33.4%	84.8%	63.4%	34.2%	90.2%
3.0E-4	68.4%	34.4%	84.8%	63.4%	34.2%	90.2%
6.0E-4	70.2%	36.8%	86.1%	61.0%	34.2%	87.8%
8.0E-4	71.0%	38.0%	86.1%	61.0%	36.7%	87.8%
9.0E-4	71.4%	38.5%	86.1%	61.0%	38.0%	87.8%
1.0E-3	71.6%	39.0%	86.1%	61.0%	38.0%	85.4%
2.0E-3	73.2%	42.9%	87.3%	61.0%	43.0%	85.4%
3.0E-3	74.3%	45.6%	88.6%	61.0%	43.0%	82.9%
4.0E-3	75.3%	48.1%	87.3%	61.0%	43.0%	82.9%
8.0E-3	78.3%	54.9%	89.9%	65.9%	41.8%	80.5%
0.01	79.4%	57.5%	89.9%	68.3%	40.5%	80.5%
0.02	83.2%	65.9%	89.9%	80.5%	44.3%	73.2%
0.03	85.4%	70.5%	89.9%	82.9%	46.8%	70.7%
0.04	86.7%	73.5%	88.6%	82.9%	48.1%	73.2%
0.15	90.7%	81.7%	88.6%	85.4%	43.0%	68.3%

Table 13 Capture Ratio Statistics for portfolio formulated under covariance structure, $\boldsymbol{S}$ for various λ levels

portfolio(λ,S)	Up Capture	Down Capture	Up Number	Down Number	Up Percent	Down Percent
0	64.4%	33.9%	84.8%	65.9%	32.9%	90.2%
1.0E-4	66.0%	35.7%	87.3%	68.3%	32.9%	90.2%
2.0E-4	66.9%	37.1%	87.3%	70.7%	32.9%	90.2%
3.0E-4	67.7%	38.0%	87.3%	63.4%	32.9%	87.8%
6.0E-4	69.3%	39.6%	88.6%	63.4%	35.4%	87.8%
8.0E-4	69.9%	40.6%	88.6%	65.9%	34.2%	87.8%
9.0E-4	70.1%	41.0%	88.6%	65.9%	35.4%	87.8%
1.0E-3	70.2%	41.4%	88.6%	65.9%	36.7%	87.8%
2.0E-3	71.0%	44.4%	88.6%	68.3%	38.0%	85.4%
3.0E-3	71.9%	46.6%	88.6%	68.3%	39.2%	82.9%
4.0E-3	72.9%	48.5%	86.1%	68.3%	36.7%	82.9%
8.0E-3	76.0%	54.2%	88.6%	75.6%	39.2%	80.5%
0.01	77.1%	56.5%	88.6%	78.1%	39.2%	80.5%
0.02	81.0%	64.2%	88.6%	80.5%	43.0%	73.2%
0.03	83.3%	68.6%	88.6%	82.9%	43.0%	73.2%
0.04	84.9%	71.6%	88.6%	85.4%	44.3%	73.2%
0.15	89.9%	80.8%	88.6%	85.4%	44.3%	68.3%

[6] Refer to Bacon (2004)

*Table 14 Capture Ratio Statistics for portfolio formulated under covariance structure, **G** for various θ levels*

portfolio(θ,G)	Up Capture	Down Capture	Up Number	Down Number	Up Percent	Down Percent
0	64.7%	30.6%	83.5%	65.9%	32.9%	90.2%
1.0E-4	66.9%	32.1%	84.8%	63.4%	31.7%	90.2%
2.0E-4	67.9%	33.1%	86.1%	63.4%	34.2%	90.2%
3.0E-4	68.7%	33.9%	86.1%	63.4%	34.2%	90.2%
8.0E-4	70.9%	37.0%	87.3%	61.0%	36.7%	87.8%
1.2E-3	72.0%	38.5%	88.6%	61.0%	39.2%	85.4%
1.8E-3	73.0%	40.0%	88.6%	63.4%	40.5%	85.4%
2.5E-3	73.6%	41.6%	89.9%	61.0%	43.0%	82.9%
3.0E-3	74.1%	42.4%	89.9%	61.0%	41.8%	82.9%
6.0E-3	76.1%	46.1%	89.9%	61.0%	40.5%	85.4%
8.0E-3	77.1%	48.1%	89.9%	61.0%	41.8%	85.4%
0.01	77.8%	49.9%	89.9%	61.0%	40.5%	85.4%
0.02	80.1%	55.6%	88.6%	65.9%	39.2%	85.4%
0.05	83.1%	62.3%	87.3%	73.2%	40.5%	78.1%
0.09	84.5%	65.5%	87.3%	75.6%	41.8%	78.1%
0.25	85.9%	68.5%	88.6%	82.9%	40.5%	75.6%
0.35	86.2%	69.1%	88.6%	82.9%	40.5%	75.6%

*Table 15 Capture Ratio Statistics for portfolio formulated under covariance structure, **S** for various θ levels*

portfolio(θ,S)	Up Capture	Down Capture	Up Number	Down Number	Up Percent	Down Percent
0	65.1%	33.7%	84.8%	65.9%	32.9%	90.2%
1.0E-4	66.8%	35.3%	87.3%	68.3%	32.9%	90.2%
2.0E-4	67.7%	36.6%	87.3%	70.7%	32.9%	90.2%
3.0E-4	68.3%	37.3%	87.3%	65.9%	34.2%	87.8%
8.0E-4	70.0%	39.4%	87.3%	65.9%	34.2%	87.8%
1.2E-3	70.5%	40.5%	87.3%	68.3%	35.4%	87.8%
1.8E-3	70.8%	41.6%	88.6%	68.3%	38.0%	85.4%
2.5E-3	71.4%	42.7%	88.6%	70.7%	40.5%	85.4%
3.0E-3	71.8%	43.4%	88.6%	68.3%	39.2%	85.4%
6.0E-3	73.7%	46.4%	89.9%	65.9%	38.0%	85.4%
8.0E-3	74.7%	48.2%	89.9%	61.0%	38.0%	85.4%
0.01	75.6%	49.6%	88.6%	63.4%	38.0%	85.4%
0.02	78.3%	54.9%	87.3%	68.3%	38.0%	85.4%
0.05	81.8%	61.4%	87.3%	75.6%	39.2%	80.5%
0.09	83.7%	64.7%	87.3%	75.6%	40.5%	75.6%
0.25	85.6%	68.1%	88.6%	82.9%	39.2%	75.6%
0.35	85.9%	68.8%	88.6%	82.9%	40.5%	75.6%

Table 16 Capture Ratio Statistics for portfolio formulated under covariance structure, $\boldsymbol{G}$ for various η levels

portfolio (η,G)	Up Capture	Down Capture	Up Number	Down Number	Up Percent	Down Percent
0	64.7%	30.6%	83.5%	65.9%	32.9%	90.2%
1.0E-4	67.0%	32.1%	84.8%	63.4%	32.9%	90.2%
2.0E-4	68.3%	33.2%	86.1%	63.4%	34.2%	92.7%
4.0E-4	70.4%	35.2%	86.1%	63.4%	34.2%	90.2%
9.0E-4	73.7%	39.1%	88.6%	61.0%	38.0%	90.2%
1.5E-3	76.3%	42.7%	91.1%	61.0%	38.0%	90.2%
2.0E-3	77.8%	45.2%	93.7%	63.4%	38.0%	92.7%
2.5E-3	79.0%	47.4%	93.7%	68.3%	39.2%	90.2%
4.0E-3	81.7%	52.3%	94.9%	68.3%	40.5%	90.2%
7.0E-3	85.0%	59.1%	94.9%	73.2%	39.2%	90.2%
9.0E-3	86.5%	62.7%	96.2%	78.1%	39.2%	87.8%
1.0E-2	87.2%	64.2%	97.5%	80.5%	40.5%	87.8%
0.02	91.6%	74.2%	97.5%	87.8%	44.3%	87.8%
0.03	94.0%	79.1%	98.7%	92.7%	45.6%	87.8%
0.07	97.7%	86.6%	100.0%	97.6%	51.9%	85.4%
0.15	99.7%	90.9%	100.0%	100.0%	54.4%	85.4%
0.45	101.1%	94.0%	100.0%	100.0%	58.2%	85.4%

Table 17 Capture Ratio Statistics for portfolio formulated under covariance structure, $\boldsymbol{S}$ for various η levels

portfolio (η,S)	Up Capture	Down Capture	Up Number	Down Number	Up Percent	Down Percent
0	65.1%	33.7%	84.8%	65.9%	32.9%	90.2%
1.0E-4	66.9%	35.5%	87.3%	68.3%	31.7%	90.2%
2.0E-4	67.8%	37.0%	87.3%	70.7%	31.7%	90.2%
4.0E-4	69.7%	38.2%	88.6%	63.4%	32.9%	87.8%
9.0E-4	72.4%	40.7%	88.6%	70.7%	35.4%	87.8%
1.5E-3	74.1%	42.9%	89.9%	70.7%	36.7%	87.8%
2.0E-3	75.2%	44.6%	89.9%	70.7%	38.0%	87.8%
2.5E-3	76.2%	46.5%	92.4%	70.7%	39.2%	87.8%
4.0E-3	78.7%	50.8%	92.4%	73.2%	38.0%	87.8%
7.0E-3	82.1%	56.5%	93.7%	80.5%	39.2%	90.2%
9.0E-3	83.6%	59.4%	93.7%	82.9%	38.0%	87.8%
1.0E-2	84.3%	60.8%	94.9%	82.9%	38.0%	87.8%
0.02	88.9%	70.5%	97.5%	90.2%	40.5%	92.7%
0.03	91.7%	75.8%	97.5%	92.7%	44.3%	92.7%
0.07	96.2%	84.4%	100.0%	92.7%	48.1%	90.2%
0.15	98.8%	89.5%	100.0%	97.6%	51.9%	87.8%
0.45	100.8%	93.5%	100.0%	100.0%	53.2%	85.4%

8.2 Appendix B

Table 18 Average Annual Turnover measured over the period May 2006 - April 2016 for various λ levels

λ	portfolio (λ,G)	portfolio (λ,S)
0	153.5%	190.8%
1.0E-4	146.6%	182.0%
2.0E-4	141.4%	175.4%
3.0E-4	137.2%	170.2%
6.0E-4	127.3%	157.1%
8.0E-4	122.8%	150.1%
9.0E-4	121.0%	147.2%
1.0E-3	119.5%	144.9%
2.0E-3	111.4%	130.9%
3.0E-3	109.0%	126.0%
4.0E-3	108.7%	123.6%
8.0E-3	107.9%	118.7%
0.01	108.0%	117.4%
0.02	108.4%	114.3%
0.03	108.4%	112.7%
0.04	108.3%	111.9%
0.15	108.5%	109.3%
benchmark	8.73%	8.73%

Table 19 Average Annual Turnover measured over the period May 2006 - April 2016 for various θ levels

θ	portfolio (θ,G)	portfolio (θ,S)
0	154.2%	191.6%
1.0E-4	147.1%	182.2%
2.0E-4	141.2%	174.9%
3.0E-4	136.2%	168.9%
8.0E-4	118.0%	145.8%
1.2E-3	109.3%	133.5%
1.8E-3	99.8%	121.3%
2.5E-3	92.5%	112.3%
3.0E-3	88.6%	107.4%
6.0E-3	74.9%	89.5%
8.0E-3	70.1%	82.5%
1.0E-2	66.8%	77.5%
0.02	58.2%	65.2%
0.05	52.2%	55.2%
0.09	50.4%	52.0%
0.25	49.7%	50.1%
0.35	49.7%	50.0%
benchmark	8.73%	8.73%

Table 20 Average Annual Turnover measured over the period May 2006 - April 2016 for various η levels

η	portfolio (η,G)	portfolio (η,S)
0	154.2%	191.6%
1.0E-4	147.1%	182.6%
2.0E-4	141.4%	176.4%
4.0E-4	131.9%	164.3%
9.0E-4	114.9%	141.4%
1.5E-3	102.6%	125.7%
2.0E-3	95.6%	117.0%
2.5E-3	90.0%	110.3%
4.0E-3	77.5%	95.0%
7.0E-3	63.7%	78.6%
9.0E-3	58.0%	71.9%
1.0E-2	55.6%	69.2%
0.02	41.8%	52.1%
0.03	34.7%	43.4%
0.07	23.6%	28.8%
0.15	17.9%	20.7%
0.45	14.3%	15.2%
benchmark	8.73%	8.73%

8.3 Appendix C

Table 21 Monthly Drawdown statistics for λ portfolios, formulated under covariance structure, $\boldsymbol{G}$

portfolio(λ,G)	Semi Deviation	Gain Deviation	Loss Deviation	Downside Deviation (0%)	Maximum Drawdown
0	2.27%	1.70%	2.28%	1.78%	28.00%
1.0E-4	2.30%	1.69%	2.32%	1.80%	27.97%
2.0E-4	2.32%	1.67%	2.35%	1.81%	28.05%
3.0E-4	2.34%	1.66%	2.38%	1.83%	28.18%
6.0E-4	2.39%	1.69%	2.45%	1.89%	28.61%
8.0E-4	2.42%	1.70%	2.49%	1.92%	28.76%
9.0E-4	2.44%	1.70%	2.52%	1.93%	28.89%
1.0E-3	2.45%	1.70%	2.54%	1.94%	29.03%
2.0E-3	2.54%	1.74%	2.65%	2.04%	30.55%
3.0E-3	2.59%	1.78%	2.70%	2.09%	31.40%
4.0E-3	2.62%	1.80%	2.75%	2.13%	31.68%
8.0E-3	2.70%	1.90%	2.79%	2.20%	32.28%
0.01	2.73%	1.94%	2.79%	2.23%	32.81%
0.02	2.84%	2.02%	2.83%	2.34%	34.54%
0.03	2.91%	2.10%	2.83%	2.40%	35.31%
0.04	2.95%	2.14%	2.84%	2.44%	35.88%
0.15	3.09%	2.35%	2.90%	2.57%	37.53%
benchmark	3.06%	2.56%	2.67%	2.53%	38.20%

*Table 22 Monthly Drawdown statistics for λ portfolios, formulated under covariance structure, **S***

portfolio(λ,S)	**Semi Deviation**	**Gain Deviation**	**Loss Deviation**	**Downside Deviation (0%)**	**Maximum Drawdown**
0	2.36%	1.72%	2.45%	1.89%	34.95%
1.0E-4	2.37%	1.69%	2.44%	1.90%	34.77%
2.0E-4	2.39%	1.66%	2.44%	1.91%	34.98%
3.0E-4	2.40%	1.70%	2.44%	1.92%	35.00%
6.0E-4	2.44%	1.72%	2.49%	1.95%	34.98%
8.0E-4	2.46%	1.71%	2.53%	1.97%	35.02%
9.0E-4	2.46%	1.71%	2.54%	1.98%	35.00%
1.0E-3	2.47%	1.72%	2.56%	1.99%	35.01%
2.0E-3	2.55%	1.73%	2.69%	2.08%	35.37%
3.0E-3	2.60%	1.76%	2.76%	2.13%	35.29%
4.0E-3	2.62%	1.75%	2.79%	2.16%	35.27%
8.0E-3	2.70%	1.83%	2.85%	2.23%	35.09%
0.01	2.73%	1.85%	2.85%	2.26%	35.22%
0.02	2.83%	1.98%	2.86%	2.34%	35.87%
0.03	2.89%	2.05%	2.87%	2.40%	36.16%
0.04	2.93%	2.09%	2.87%	2.43%	36.43%
0.15	3.08%	2.32%	2.90%	2.56%	37.65%
benchmark	3.06%	2.56%	2.67%	2.53%	38.20%

*Table 23 Monthly Drawdown statistics for θ portfolios, formulated under covariance structure, **G***

portfolio(θ,G)	**Semi Deviation**	**Gain Deviation**	**Loss Deviation**	**Downside Deviation (0%)**	**Maximum Drawdown**
0	2.26%	1.70%	2.29%	1.77%	27.75%
1.0E-4	2.29%	1.69%	2.35%	1.79%	27.07%
2.0E-4	2.31%	1.67%	2.38%	1.81%	26.85%
3.0E-4	2.33%	1.66%	2.41%	1.83%	26.84%
8.0E-4	2.41%	1.68%	2.53%	1.91%	27.56%
1.2E-3	2.46%	1.70%	2.61%	1.96%	28.15%
1.8E-3	2.51%	1.70%	2.69%	2.01%	29.02%
2.5E-3	2.54%	1.74%	2.74%	2.05%	29.85%
3.0E-3	2.56%	1.74%	2.76%	2.07%	30.30%
6.0E-3	2.63%	1.80%	2.84%	2.13%	31.45%
8.0E-3	2.66%	1.82%	2.87%	2.16%	31.86%
0.01	2.68%	1.84%	2.89%	2.19%	32.20%
0.02	2.75%	1.90%	2.92%	2.25%	33.38%
0.05	2.82%	1.97%	2.88%	2.31%	34.66%
0.09	2.86%	2.02%	2.86%	2.34%	35.29%
0.25	2.90%	2.04%	2.85%	2.38%	35.89%
0.35	2.91%	2.05%	2.84%	2.38%	36.00%
benchmark	3.06%	2.56%	2.67%	2.53%	38.20%

Table 24 Monthly Drawdown statistics for θ portfolios, formulated under covariance structure, $\boldsymbol{S}$

portfolio(* θ *,S)	**Semi Deviation**	**Gain Deviation**	**Loss Deviation**	**Downside Deviation (0%)**	**Maximum Drawdown**
0	2.34%	1.72%	2.44%	1.87%	33.65%
1.0E-4	2.35%	1.69%	2.45%	1.87%	33.02%
2.0E-4	2.37%	1.66%	2.45%	1.89%	33.25%
3.0E-4	2.38%	1.69%	2.47%	1.90%	33.25%
8.0E-4	2.43%	1.69%	2.55%	1.95%	33.49%
1.2E-3	2.46%	1.68%	2.61%	1.98%	33.69%
1.8E-3	2.50%	1.70%	2.68%	2.02%	33.87%
2.5E-3	2.53%	1.70%	2.73%	2.06%	34.07%
3.0E-3	2.55%	1.73%	2.77%	2.08%	34.19%
6.0E-3	2.62%	1.80%	2.88%	2.15%	34.43%
8.0E-3	2.66%	1.84%	2.92%	2.18%	34.43%
0.01	2.68%	1.83%	2.94%	2.20%	34.42%
0.02	2.75%	1.86%	2.96%	2.26%	34.76%
0.05	2.82%	1.93%	2.91%	2.32%	35.26%
0.09	2.85%	2.00%	2.88%	2.35%	35.62%
0.25	2.90%	2.03%	2.86%	2.38%	36.01%
0.35	2.90%	2.05%	2.85%	2.38%	36.08%
benchmark	3.06%	2.56%	2.67%	2.53%	38.20%

Table 25 Monthly Drawdown statistics for η portfolios, formulated under covariance structure, $\boldsymbol{G}$

portfolio(* η *,G)	**Semi Deviation**	**Gain Deviation**	**Loss Deviation**	**Downside Deviation (0%)**	**Maximum Drawdown**
0	2.26%	1.70%	2.29%	1.77%	27.75%
1.0E-4	2.28%	1.69%	2.34%	1.78%	26.98%
2.0E-4	2.30%	1.67%	2.37%	1.80%	26.72%
4.0E-4	2.33%	1.65%	2.43%	1.82%	26.83%
9.0E-4	2.38%	1.70%	2.53%	1.87%	27.70%
1.5E-3	2.41%	1.75%	2.58%	1.89%	28.36%
2.0E-3	2.43%	1.77%	2.58%	1.90%	28.70%
2.5E-3	2.44%	1.75%	2.58%	1.91%	29.01%
4.0E-3	2.46%	1.80%	2.54%	1.93%	29.73%
7.0E-3	2.51%	1.88%	2.54%	1.97%	30.83%
9.0E-3	2.54%	1.93%	2.56%	2.00%	31.46%
0.01	2.56%	1.96%	2.56%	2.02%	31.76%
0.02	2.68%	2.11%	2.58%	2.13%	33.43%
0.03	2.76%	2.20%	2.61%	2.20%	34.32%
0.07	2.88%	2.37%	2.65%	2.32%	35.70%
0.15	2.96%	2.47%	2.68%	2.39%	36.50%
0.45	3.02%	2.56%	2.71%	2.45%	37.09%
benchmark	3.06%	2.56%	2.67%	2.53%	38.20%

Table 26 Monthly Drawdown statistics for η portfolios, formulated under covariance structure, $\boldsymbol{S}$

portfolio(η,S)	**Semi Deviation**	**Gain Deviation**	**Loss Deviation**	**Downside Deviation (0%)**	**Maximum Drawdown**
0	2.34%	1.72%	2.44%	1.87%	33.65%
1.0E-4	2.35%	1.69%	2.45%	1.87%	32.98%
2.0E-4	2.36%	1.65%	2.44%	1.88%	33.15%
4.0E-4	2.38%	1.70%	2.48%	1.89%	33.06%
9.0E-4	2.40%	1.66%	2.52%	1.90%	32.84%
1.5E-3	2.42%	1.69%	2.58%	1.92%	32.66%
2.0E-3	2.43%	1.70%	2.60%	1.93%	32.62%
2.5E-3	2.44%	1.73%	2.62%	1.94%	32.65%
4.0E-3	2.46%	1.73%	2.61%	1.96%	32.68%
7.0E-3	2.50%	1.77%	2.59%	1.98%	32.71%
9.0E-3	2.52%	1.81%	2.59%	2.00%	32.89%
0.01	2.54%	1.85%	2.59%	2.01%	33.00%
0.02	2.65%	2.01%	2.62%	2.11%	33.96%
0.03	2.72%	2.11%	2.63%	2.18%	34.58%
0.07	2.85%	2.34%	2.66%	2.30%	35.76%
0.15	2.94%	2.45%	2.69%	2.37%	36.50%
0.45	3.01%	2.54%	2.71%	2.44%	37.08%
benchmark	3.06%	2.56%	2.67%	2.53%	38.20%

8.4 Appendix D

Table 27 Portfolio Statistics for various λ portfolios, formulated under covariance structure, $\boldsymbol{G}$

portfolio(λ,G)	***Quartile 1***	***Median***	***Quartile 3***	***SE Mean***	***Skewness***	***Kurtosis***
0	-0.65%	1.37%	3.31%	0.27%	-0.75	1.84
1.0E-4	-0.58%	1.37%	3.42%	0.27%	-0.85	1.52
2.0E-4	-0.49%	1.32%	3.44%	0.27%	-0.90	1.44
3.0E-4	-0.48%	1.32%	3.52%	0.27%	-0.93	1.41
6.0E-4	-0.55%	1.33%	3.49%	0.28%	-0.98	1.47
8.0E-4	-0.55%	1.30%	3.53%	0.28%	-1.00	1.55
9.0E-4	-0.48%	1.30%	3.51%	0.28%	-1.01	1.60
1.0E-3	-0.48%	1.30%	3.53%	0.29%	-1.02	1.64
2.0E-3	-0.45%	1.42%	3.39%	0.29%	-1.05	1.82
3.0E-3	-0.44%	1.55%	3.31%	0.30%	-1.05	1.89
4.0E-3	-0.52%	1.62%	3.28%	0.30%	-1.05	1.91
8.0E-3	-0.66%	1.55%	3.26%	0.31%	-0.97	1.69
0.01	-0.71%	1.53%	3.26%	0.32%	-0.93	1.57
0.02	-1.02%	1.51%	3.57%	0.34%	-0.82	1.24
0.03	-1.02%	1.47%	3.95%	0.35%	-0.76	1.06
0.04	-1.13%	1.54%	4.20%	0.35%	-0.72	0.97
0.15	-1.48%	1.46%	4.39%	0.38%	-0.61	0.73
benchmark	-1.54%	1.37%	3.40%	0.38%	-0.35	0.71

Table 28 Portfolio Statistics for various λ portfolios, formulated under covariance structure, $\boldsymbol{S}$

portfolio(λ,S)	Quartile 1	Median	Quartile 3	SE Mean	Skewness	Kurtosis
0	-0.60%	1.33%	3.29%	0.28%	-0.87	1.97
1.0E-4	-0.67%	1.31%	3.42%	0.28%	-0.91	1.66
2.0E-4	-0.67%	1.39%	3.38%	0.28%	-0.93	1.53
3.0E-4	-0.66%	1.33%	3.38%	0.28%	-0.93	1.48
6.0E-4	-0.54%	1.35%	3.31%	0.28%	-0.96	1.53
8.0E-4	-0.50%	1.35%	3.23%	0.29%	-0.98	1.61
9.0E-4	-0.49%	1.35%	3.22%	0.29%	-0.99	1.66
1.0E-3	-0.48%	1.37%	3.23%	0.29%	-1.00	1.70
2.0E-3	-0.52%	1.41%	3.08%	0.30%	-1.07	2.02
3.0E-3	-0.48%	1.52%	3.07%	0.30%	-1.08	2.11
4.0E-3	-0.47%	1.56%	3.12%	0.30%	-1.08	2.11
8.0E-3	-0.75%	1.60%	3.30%	0.31%	-1.03	1.98
0.01	-0.82%	1.53%	3.28%	0.32%	-1.00	1.85
0.02	-0.93%	1.52%	3.51%	0.33%	-0.88	1.47
0.03	-1.05%	1.50%	3.63%	0.34%	-0.82	1.27
0.04	-1.09%	1.53%	3.92%	0.35%	-0.77	1.14
0.15	-1.44%	1.49%	4.39%	0.37%	-0.63	0.79
benchmark	-1.54%	1.37%	3.40%	0.38%	-0.35	0.71

Table 29 Portfolio Statistics for various θ portfolios, formulated under covariance structure, $\boldsymbol{G}$

portfolio(θ,G)	Quartile 1	Median	Quartile 3	SE Mean	Skewness	Kurtosis
0	-0.65%	1.37%	3.31%	0.27%	-0.75	1.88
1.0E-4	-0.54%	1.37%	3.39%	0.27%	-0.86	1.61
2.0E-4	-0.51%	1.32%	3.43%	0.27%	-0.92	1.53
3.0E-4	-0.50%	1.33%	3.43%	0.27%	-0.96	1.51
8.0E-4	-0.48%	1.43%	3.49%	0.28%	-1.04	1.67
1.2E-3	-0.42%	1.43%	3.51%	0.28%	-1.08	1.86
1.8E-3	-0.38%	1.48%	3.45%	0.29%	-1.12	2.07
2.5E-3	-0.34%	1.44%	3.47%	0.29%	-1.14	2.18
3.0E-3	-0.31%	1.49%	3.47%	0.29%	-1.14	2.21
6.0E-3	-0.23%	1.65%	3.46%	0.30%	-1.14	2.29
8.0E-3	-0.35%	1.64%	3.34%	0.31%	-1.13	2.29
0.01	-0.44%	1.62%	3.32%	0.31%	-1.12	2.28
0.02	-0.58%	1.67%	3.40%	0.32%	-1.05	2.07
0.05	-0.74%	1.60%	3.61%	0.33%	-0.93	1.72
0.09	-0.95%	1.57%	3.75%	0.34%	-0.87	1.55
0.25	-1.16%	1.48%	3.80%	0.34%	-0.81	1.40
0.35	-1.19%	1.48%	3.80%	0.35%	-0.80	1.37
benchmark	-1.54%	1.37%	3.40%	0.38%	-0.35	0.71

Table 30 Portfolio Statistics for various θ portfolios, formulated under covariance structure, $\boldsymbol{S}$

portfolio(θ ,S)	Quartile 1	Median	Quartile 3	SE Mean	Skewness	Kurtosis
0	-0.60%	1.34%	3.29%	0.28%	-0.88	2.07
1.0E-4	-0.66%	1.35%	3.42%	0.28%	-0.93	1.79
2.0E-4	-0.64%	1.44%	3.38%	0.28%	-0.95	1.67
3.0E-4	-0.60%	1.37%	3.37%	0.28%	-0.96	1.61
8.0E-4	-0.42%	1.44%	3.26%	0.28%	-1.00	1.70
1.2E-3	-0.32%	1.47%	3.22%	0.29%	-1.05	1.88
1.8E-3	-0.43%	1.46%	3.18%	0.29%	-1.09	2.12
2.5E-3	-0.36%	1.56%	3.24%	0.29%	-1.12	2.25
3.0E-3	-0.35%	1.57%	3.19%	0.29%	-1.13	2.29
6.0E-3	-0.35%	1.54%	3.25%	0.30%	-1.16	2.44
8.0E-3	-0.37%	1.54%	3.31%	0.31%	-1.16	2.46
0.01	-0.42%	1.53%	3.37%	0.31%	-1.15	2.44
0.02	-0.65%	1.66%	3.31%	0.32%	-1.10	2.28
0.05	-0.77%	1.58%	3.58%	0.33%	-0.97	1.87
0.09	-0.89%	1.58%	3.67%	0.34%	-0.90	1.66
0.25	-1.15%	1.49%	3.80%	0.34%	-0.83	1.44
0.35	-1.18%	1.47%	3.80%	0.34%	-0.82	1.40
benchmark	-1.54%	1.37%	3.40%	0.38%	-0.35	0.71

Table 31 Portfolio Statistics for various η portfolios, formulated under covariance structure, $\boldsymbol{G}$

portfolio(η ,G)	Quartile 1	Median	Quartile 3	SE Mean	Skewness	Kurtosis
0	-0.65%	1.37%	3.31%	0.27%	-0.75	1.88
1.0E-4	-0.51%	1.42%	3.38%	0.27%	-0.86	1.60
2.0E-4	-0.45%	1.41%	3.38%	0.27%	-0.92	1.54
4.0E-4	-0.39%	1.50%	3.44%	0.27%	-0.98	1.56
9.0E-4	-0.29%	1.43%	3.36%	0.28%	-1.03	1.68
1.5E-3	-0.12%	1.47%	3.35%	0.28%	-1.03	1.71
2.0E-3	-0.11%	1.53%	3.31%	0.28%	-1.01	1.67
2.5E-3	-0.16%	1.50%	3.35%	0.28%	-1.00	1.62
4.0E-3	-0.20%	1.55%	3.46%	0.29%	-0.94	1.41
7.0E-3	-0.27%	1.40%	3.41%	0.30%	-0.83	1.11
9.0E-3	-0.35%	1.35%	3.37%	0.30%	-0.78	1.02
1.0E-2	-0.37%	1.33%	3.46%	0.30%	-0.76	0.99
0.02	-0.67%	1.29%	3.44%	0.32%	-0.63	0.82
0.03	-0.78%	1.41%	3.43%	0.33%	-0.57	0.77
0.07	-1.07%	1.44%	3.49%	0.35%	-0.46	0.73
0.15	-1.24%	1.41%	3.53%	0.37%	-0.41	0.72
0.45	-1.37%	1.35%	3.54%	0.37%	-0.37	0.72
benchmark	-1.54%	1.37%	3.40%	0.38%	-0.35	0.71

Table 32 Portfolio Statistics for various η portfolios, formulated under covariance structure, $\boldsymbol{S}$

portfolio(η ,S)	Quartile 1	Median	Quartile 3	SE Mean	Skewness	Kurtosis
0	-0.60%	1.34%	3.29%	0.28%	-0.88	2.07
1.0E-4	-0.66%	1.38%	3.35%	0.28%	-0.93	1.79
2.0E-4	-0.62%	1.42%	3.28%	0.28%	-0.95	1.68
4.0E-4	-0.54%	1.35%	3.20%	0.28%	-0.97	1.66
9.0E-4	-0.47%	1.47%	3.16%	0.28%	-1.00	1.75
1.5E-3	-0.30%	1.49%	3.16%	0.28%	-1.03	1.86
2.0E-3	-0.25%	1.53%	3.23%	0.28%	-1.03	1.89
2.5E-3	-0.33%	1.58%	3.27%	0.28%	-1.03	1.89
4.0E-3	-0.30%	1.57%	3.26%	0.29%	-1.01	1.77
7.0E-3	-0.40%	1.51%	3.42%	0.29%	-0.93	1.48
9.0E-3	-0.37%	1.43%	3.34%	0.30%	-0.88	1.34
1.0E-2	-0.41%	1.40%	3.31%	0.30%	-0.86	1.30
0.02	-0.87%	1.30%	3.52%	0.32%	-0.72	1.07
0.03	-0.87%	1.41%	3.38%	0.33%	-0.65	0.96
0.07	-1.16%	1.39%	3.43%	0.35%	-0.52	0.84
0.15	-1.27%	1.41%	3.51%	0.36%	-0.44	0.78
0.45	-1.37%	1.35%	3.53%	0.37%	-0.38	0.74
benchmark	-1.54%	1.37%	3.40%	0.38%	-0.35	0.71

8.5 Appendix E

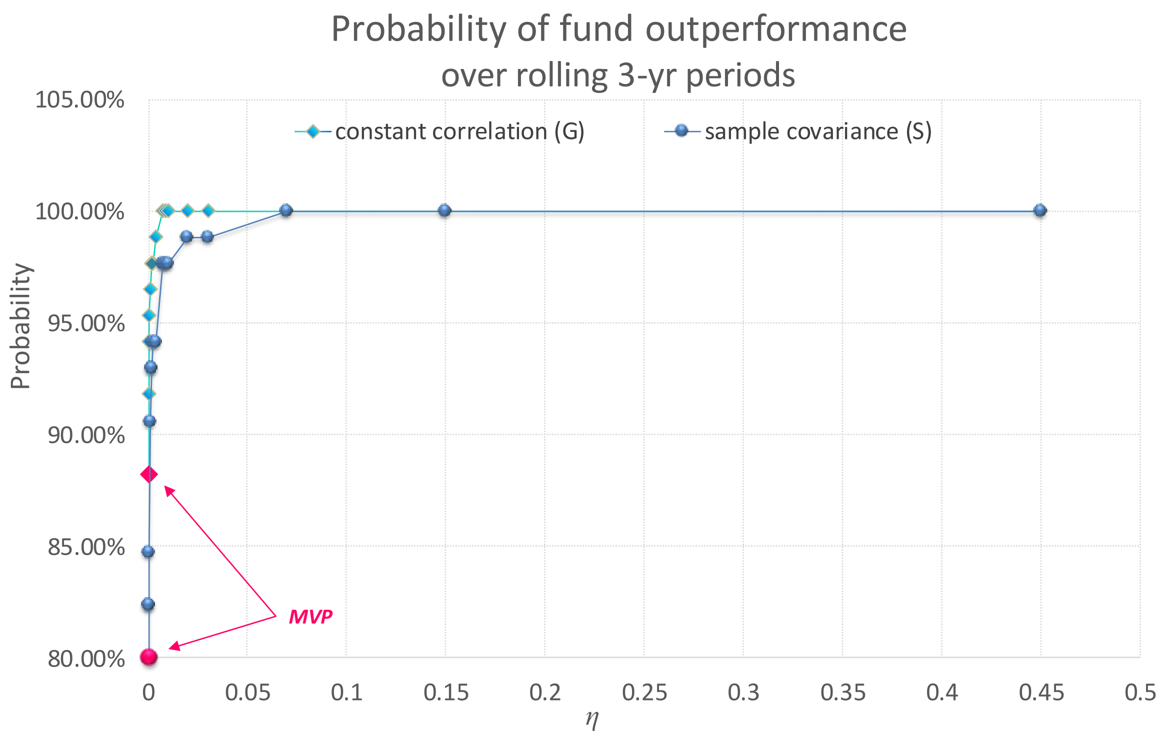

Figure 64 Probability of outperformance for η various levels

The benefit of the tilt is clear from Figure 64 above, in that the portfolios have a high probability of outperforming the benchmark over three-year rolling periods. As η increases this percentage increases significantly and approaches 100%. This could indicate that timing of rebalancing added more value than the actual tilt. Employing shrinkage is beneficial across all levels of η.

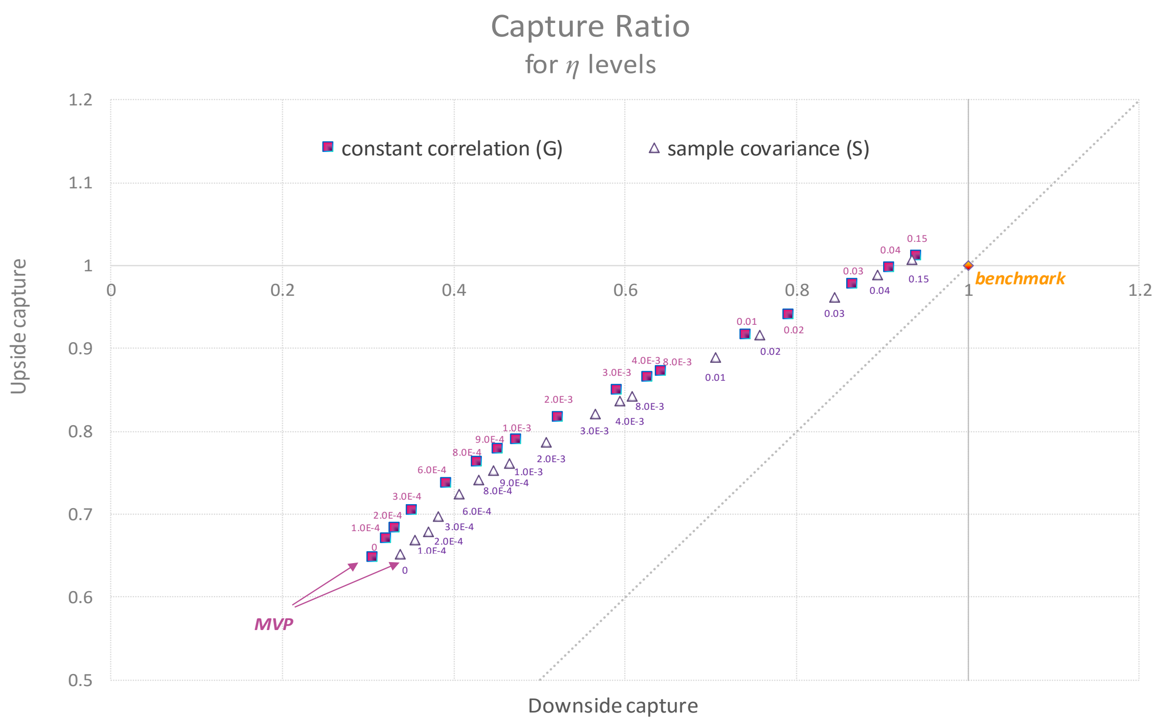

Figure 65 Capture Ratio - η Levels