Robust Bayesian Portfolio Optimisation: Higher Moments and the Distorting Effects of Constraints

Dissertation submitted in fulfillment of the requirements for the degree of Master of Business Science

Byron Wilson
WLSBYR001

Supervisor: Adjunct Professor Gareth Witten

Faculty of Commerce
September 2012

The financial assistance of the National Research Foundation (NRF) towards this research is hereby acknowledged. Opinions expressed and conclusions arrived at, are those of the author and are not necessarily to be attributed to the NRF.

Plagiarism Declaration

1. I know the meaning of plagiarism and declare that all of the work in the document, save for that which is properly acknowledged, is my own
2. I have used the APA referencing guide for citation and referencing. Each contribution to, and quotation in this dissertation from the work(s) of other people has been contributed, and has been cited and referenced.
3. This dissertation is my own work.
4. I have not allowed, and will not allow, anyone to copy my work.

Signature: Signed by candidate

Date: 14/9/2012

Acknowledgements

I would like to express my appreciation to my family and friends who helped and inspired me during the course of writing my Masters dissertation. In particular, I would like to thank my supervisor, Dr Gareth Witten, for his guidance and support during the course of my studies as well as for his assistance in making my attendance at a conference in the USA possible - it was an invaluable experience.

In addition, I would like to thank the NRF, UCT Postgraduate Funding Office and the UCT Statistics department for their financial support during the course of my Masters dissertation.

Abstract

The aim of this thesis is to introduce the Bayesian approach to asset allocation. In particular, the Black-Litterman model is introduced as a powerful Bayesian asset allocation model that enables the incorporation of human decision making (in the form of views) within a portfolio optimisation framework. Several recommendations and adjustments are made within the Black-Litterman framework in order to improve its practical applicability. In particular, a major shortcoming of the Black-Litterman model is the normality of returns assumption. Robust estimates of higher moments and co-moments (co-skewness and co-kurtosis) are introduced and implemented within the Black-Litterman framework, thereby enabling the investor to express preferences for skewness and kurtosis as well as avoiding the pitfall of large negative returns that typically occur with a greater frequency than what is suggested by the normal distribution.

In addition, a suite of diagnostic tools aimed at analysing the individual contributions of the expressed views as well as constraints is developed. In particular, the diagnostic tools enable the investor to analyse the active weight and tracking error contributions of each view to the portfolio, therefore providing a transparent portfolio optimisation methodology whereby each particular driver of the asset allocations can be identified. More specifically, a novel approach is followed whereby the diagnostic tools are used to examine the severe distorting effects of the imposed constraints. Disturbing results are obtained whereby the imposed constraints effectively "drown out" the views expressed by the investor. For the particular example considered, the constraints account for over a third of the portfolio allocations and effectively change 32 out of the 40 expressed views. In order to mitigate the ill-effects of the imposed constraints, the long-only constraint is marginally relaxed. It was determined that a 123/23 portfolio resulted in a significant improvement in the expression of the investor views as well as a dramatic increase portfolio utility was observed.

In summary, incorporating higher order moments and applying the suite of diagnostic tools to the Black-Litterman framework, a transparent portfolio optimisation methodology that effectively utilises the human decision making and investor preferences is obtained.

Table of Contents

List of Tables and Figures

Chapter 1: Introduction

A significant majority of active portfolio managers use a fundamental analysis investment strategy in order to select assets and weight them in a portfolio. Performance forecasts on respective assets are generally qualitative and are usually based on fundamental indicators (Herold, 2003). In constructing portfolios, current techniques are largely heuristic based whereby quantitative optimisation models are used as a guideline in order to provide useful information on the possible location of the optimal portfolio for the investor (Schneeweis *et al.,* 2010). Quantitative asset allocation models are therefore used as tools whereby the recommended portfolio allocations are not necessarily rigorously followed and are largely subjected to the post-hoc adjustments of the asset manager. In particular, the asset manager typically makes various adjustments to the asset allocations in light of the solutions obtained from the quantitative model (Herold, 2003). These adjustments are largely based on the intuition and skill of the asset manager in following and understanding market drivers, which can be described as more of an art than a science. It is therefore evident that a gap exists in the process of asset allocation whereby the use of "scientific" approaches (based on various quantitative models) and the skilful "artful" adjustments (based on intuition of the asset manager) are often separated.

In addition to the divide that occurs between the art and science of asset allocation, the traditional quantitative asset allocation models lack "robustness" in the sense that extreme and unintuitive asset allocations are recommended as a result of small estimation errors inherent in the underlying data (Michaud, 1989). Estimation error results in parameter uncertainty and is typically the basis for many asset allocation models failing in practise (Avramov & Zhou, 2010). It is for this particular reason that many asset managers use quantitative asset allocation models as a guideline only and rely on post-hoc adjustments based on intuition (Schneeweis *et al.,* 2010).

In order to overcome the problem of parameter uncertainty as well as to facilitate the inclusion of subjective human decision making into the asset allocation process, the Bayesian approach to asset allocation has proven to provide a theoretically sound framework that overcomes many of these problems (Harvey *et al.,* 2008). It is for these particular reasons that the Bayesian approach to asset allocation has found a recent considerable increase in popularity amongst researchers and practitioners alike (Rachev *et al.,* 2008).

In this thesis, the Black-Litterman model is introduced as a powerful Bayesian asset allocation model that facilitates the combination of unique investor views and market equilibrium returns in a manner that results in asset allocations that intuitively reflect both sources of information. In particular, a notable strength of the Black-Litterman model is that correlations between assets are considered when investor views are expressed on specific assets in a portfolio, therefore considering relationships between the assets that would otherwise have been ignored. In addition, due to the Bayesian underpinning of the Black-

Litterman model, parameter estimates are considered as random variables (as opposed to fixed estimates), therefore incorporating the effects of parameter uncertainty directly into the optimisation process (Avramov & Zhou, 2010). Despite the notable strengths of the Black-Litterman model, it suffers from the underlying assumption that asset returns are normally distributed. However, in practise many asset returns do not follow a normal distribution and therefore, investors may exhibit preferences for higher-moments in addition to the mean and variance of a particular asset (Harvey & Siddique, 2000). In order to extend the Black-Litterman framework to include higher moments, robust estimates for higher order co-moments are utilised in order to achieve reverse-optimised higher moment implied equilibrium returns, which are then used as a neutral starting point for which the investor views can be applied. In this way, the extended higher moment Black-Litterman model satisfies investors who express risk aversions towards negative skewness and lower kurtosis inherent in asset returns.

In addition, there is currently very little literature available which analyses the impacts of the expressed views and constraints on the optimal asset allocations recommended by the Black-Litterman model. A suite of diagnostic tools are therefore developed in order to analyse the respective effects of the views and constraints on the optimal portfolio. More specifically, these developed tools analyse the percentage contribution of the views to the portfolio tracking error, active return and active weights. Fabozzi *et al.,* (2007) state that a rigorous risk management procedure is essential for any portfolio construction activity. Efficient risk management ensures that investors are aware of the possible risk consequences of their respective views. Therefore, the aim of the diagnostic tools is to aid the investor in avoiding excessive and unintentional risk taking as well as uncontrollable tracking error as a result of the expressed views and constraints. In this way, the investor is able to determine whether the impact of an expressed view is consistent with the intended risk and return objectives of the portfolio.

In addition to analysing how each individual view impacts the optimal portfolio, the impacts of constraints on the optimal portfolio allocations are analysed. More specifically, the distorting effects of constraints on the expressed views are exposed. A disturbing finding is observed whereby the views are distorted to the extent that they are completely drowned out by the constraints. The practical implications of such a finding is that a significant amount of market research and financial resources goes into generating these views and therefore, the investor would have been better off not expressing any views at all. In order to reduce the distorting effects of constraints, a marginal relaxation of the long-only constraint is examined.

Given the above extensions to the Black-Litterman model, investors are able to more effectively apply the Bayesian approach in practise. In this way, investors are able to bridge the gap between the art and science of asset allocation. In addition, investors are able to utilise an efficient risk management analysis (via risk diagnostic tools), therefore allowing for portfolios that are transparent in the sense that the risk and return effects of each expressed view can be monitored and calibrated accordingly. Portfolios obtained in this manner, intuitively represent the views of the investor, are robust and are well-balanced in terms of their respective risk and return profiles.

1.1 Objectives

The objectives of this thesis are to develop a robust portfolio optimisation methodology that appropriately takes into account the respective risk and return objectives of the investor. In addition, a suite of diagnostic tools will be developed in order for the investor to observe the direct impacts that specific views and constraints have on the optimal portfolio. In particular, the objectives of the thesis can be summarised as follows:

- Introduce the Bayesian approach to asset allocation and describe why a robust portfolio optimisation methodology is essential.
- Introduce higher moment utility functions and apply them to the Black-Litterman Model
- Develop a suite of portfolio diagnostic tools that identify the respective impacts that views and constraints have on the final optimal portfolio. In particular, the diagnostic tools will be used to illustrate the adverse effects of constraints on the optimal portfolio.
- Analyse how marginally relaxing the long-only constraint influences the optimal portfolio.

1.2 Thesis Outline

Chapter 2 introduces the notion that effective asset management is essentially an art and science combined, whereby investors cannot solely rely on optimisation models to produce *truly* optimal portfolios. In particular, a naive approach to asset allocation is described as one in which the investor relies solely on the scientific approach without the practical consideration of human decision making and individual preferences. In light of the aforementioned approach to asset allocation, the Bayesian approach to asset allocation is described as the key to achieving a balance between the art and science. The Black-Litterman model is then introduced as a powerful Bayesian asset allocation model that allows investors to combine their unique views on the market with reverse optimised implied equilibrium returns in order to result in a composite portfolio that intuitively reflects the various sources of information.

In light of the discussion of traditional asset allocation and its problems, **chapter 3** introduces the concept of Robust Statistics and explains the importance of robustness in asset allocation models. In addition, the Bayesian approach is recognised and described as a robust approach to developing portfolios that are intuitive and stable with regards to estimation errors.

Consistent with the subject of constructing robust portfolios that intuitively capture the respective preferences of investors, **chapter 4** introduces the various approaches to modelling investor utility. In particular, a Bayesian shrinkage approach to estimating higher-moment utility functions is introduced. In addition, the estimation of the implied equilibrium returns, used as a starting point for the Black-Litterman model, is extended to include the robust higher-moment estimates present in the underlying data.

In light of the subject of developing portfolios that intuitively represent investor views and preferences, **chapter 5** introduces various portfolio and constraint diagnostic tools that are developed with the aim of identifying and analysing how the views and constraints influence the optimal portfolio. In particular, strong attention is drawn to the distorting effects of constraints as well as how a mild relaxation of the long-only constraint results in a significant improvement in the efficiency of the portfolio.

Chapter 6 applies the concepts described in the previous chapters. In particular, two main examples are covered. The first example deals with analysing how the introduction of robust higher moments influences a Domestic Balanced Fund portfolio consisting of six asset classes. The second example applies the Black-Litterman model to a FTSE/JSE ALSI Top 40 asset portfolio and utilises the view and constraint diagnostic tools to analyse how each view influences the optimal portfolio.

Chapter 7 presents the final conclusion to the thesis. In addition, Appendix A includes an additional and more in-depth discussion of the framework of the Black-Litterman model. Appendix B discusses two popular competing robust optimisation methodologies. The attached CD includes Appendix C, the MATLAB programming applicable to the results obtained in this dissertation.

Chapter 2: The Art and Science of Asset Management

Asset allocation can be defined as a process of dividing funds among different investments, with the aim of achieving an optimal trade-off between risk and return as well as satisfying the personal objectives of the investor (Welch, 2008). In essence, asset management is based on the common notion that it is better to diversify ones funds amongst a variety of assets than to invest solely in one particular asset.

The science of asset allocation involves the concepts discussed in modern portfolio theory and various other portfolio optimisation techniques[1]. Despite the notable popularity of certain portfolio optimisation techniques, a common weakness occurs when human involvement is required, with many portfolio optimisers recommending "optimal" portfolios that are neither intuitive nor suitable for investment. In contrast to the science of asset allocation, the art of asset allocation lies in the asset manager's ability to impose forward-looking views on the portfolio as well as appropriately modifying portfolios to suit the personal requirements of the investor. In practice, asset managers typically use the traditional mean-variance methodology as a starting point for portfolio construction and then, via a post-hoc procedure, subjectively modify the asset allocations in order to suitably "tailor-make" the portfolio to suit the individual needs of the investor. The process of adjusting the portfolios from their mean-variance "optimal" positions to form a practically investable portfolio is largely a "coarse" process whereby nuisance weights are eliminated and other adjustments are made. The effects of such adjustments typically go un-monitored and a significant deviation from optimality may be observed (Fabozzi *et al.,* 2007).

The real challenge of asset allocation lies in the art of asset allocation, in which asset managers must determine which assets to invest in as well as formulate accurate risk and return expectations. Making informed judgements about future performance requires vast amounts of data analysis and market observation. Asset managers are required to make sense of all sources of information (market outlook as well as fundamental and economic factors) in order to form summarised judgements on the possible future of the market. As described by the Efficient Market Hypothesis, asset managers are compelled to make judgements about the future performance as investors cannot earn significant returns based on past performance alone. The efficient market hypothesis states that asset prices include all available information and that future asset returns are, by definition, unpredictable. Therefore, by applying an optimisation algorithm exclusively on past data will only ensure optimality for the past periods of observation and will not necessarily be optimal for future periods. Matching investors with suitable portfolios as well as incorporating appropriate trading strategies based on human expectations, risk

---

[1] Various portfolio optimisation techniques exist in literature. Examples of these include the traditional Markowitz mean-variance model, multi-moment and robust portfolio optimisation methodologies.

[2] Jagannathan and Ma (2003) discuss that imposing the long-only constraint reduces estimation error. The

tolerance and return objectives in a manner that effectively utilises the numerous scientific optimisation techniques, *consistently*, is truly an art form that is rarely mastered by asset managers.

It therefore follows that the scientific approach to asset allocation alone cannot result in *truly* optimal portfolios. Successful asset managers therefore possess the ability of effectively utilising and merging both the art and science approaches to asset allocation. This involves utilising various sources of information, in the form of qualitative and quantitative judgements in order to form views on the market that are accurate and consistent with the objectives of the investor and then making best use of the various optimisation techniques available. In essence, effective asset allocation lies in the asset manager's ability to understand financial market drivers as well as effectively utilise the (scientific) optimisation techniques at their disposal.

As described above, a disciplined asset allocation process involves a series of sequential steps aimed at satisfying the personal needs of the investor needs in a manner that is mathematically optimal. Effective asset allocation decisions are created on the foundation of a sound quantitative and qualitative decision framework. Effective asset allocation therefore proceeds beyond the mathematical optimisation of returns, standard deviations and correlations and aims to incorporate human judgement and interaction (Avramov & Zhou, 2010).

The following section outlines traditional asset allocation and its problems. The Bayesian approach to asset allocation is then introduced and it is then shown how the Bayesian framework is used to rectify some of the problems observed in traditional asset allocation as well as how the Bayesian methodology fits within the field of Decision Modelling and Behavioural Finance.

2.1 Traditional Asset Allocation and its Problems

The original mean-variance optimisation model formulated by Harry Markowitz in 1951 (Markowitz, 1951) quantified the trade-off between the risk and return of a security for investors. The original version of the mean-variance framework assumed that there were no transaction costs or taxes and investors had identical information about assets (and thus shared the same estimates for the expected return, standard deviations, and correlation across assets). Despite these assumptions, the mean-variance optimisation framework has been used extensively in portfolio management. Various authors (Michaud, 1989), Jobson and Korkie (1981), Britten-Jones (1999) have criticised the implementation of the mean-variance framework to real-world portfolios (securities) because it results in counter-intuitive asset allocations of extreme security weightings due to the sensitivity of the inputs to the model (expected returns, volatilities and correlation). Best and Grauer (1991) showed that the portfolio weights are particularly sensitive to the expected return inputs, exposing an investor to unintended risks. Furthermore, Chopra and Ziemba (1993) examined the impact of errors in the parameter inputs on the optimal portfolio weights and demonstrated that errors in the expected returns are a maximum of eleven times more important than errors in the variances between security returns.

Various methodologies and extensions have been proposed to overcome the problems of the mean-variance framework. One attempt to resolve the occurrence of extreme solutions is to impose constraints on the asset allocations, however, imposing constraints forces the solution to a pre-specified allocation, nullifying the impact of the optimisation procedure. In an attempt to reduce the estimation error present in the estimation of the variance-covariance matrix, factor models that capture various risk sources have been used (Fabozzi, 1998). Jagannathan and Ma (2003) found that, under certain conditions[2], imposing the long-only constraint, the sample covariance matrix performs as well as factor model and shrinkage estimates. An additional approach used to mitigate the estimation error of input parameters is the use of robust optimisation techniques. Robust optimisation attempts to minimise the worst case return for a given confidence interval subject to the usual constraints (no short selling and the portfolio weights should sum to one). However, Scherer (2002) shows that robust optimisation is computationally intensive and that is equivalent to the shrinkage estimation methods proposed by Jorion (1986), which are significantly easier to implement in practice[3]. A third approach, as proposed by Michaud (1998), was the use of a portfolio resampling approach in order to mitigate some of the estimation error involved in computing the portfolio parameters. Michaud (1998) illustrated that portfolios on the resampled frontier are more diversified than the portfolios on the traditional mean-variance frontier. Scherer (2002) illustrated that resampled long-only portfolios may possibly (although it is very unclear) lead to more diversified out-of-sample portfolios that outperform out-of-sample Markowitz portfolios for the long-only case.

2.2 The Bayesian Approach to Asset Allocation

In contrast to the traditional frequentist approach, Bayesian statistics assumes that unknown parameters are considered as random variables that follow a particular distribution. In the context of portfolio selection, Bayesian statistics allows us to impose a prior view, and then, upon the arrival of new data, to alter our view in light of newly acquired information. In other words, the Bayesian framework allows investors to introduce non-sample information in the form of a prior distribution which can be optimally combined with sample information in order to result in improved estimates.

[2] Jagannathan and Ma (2003) discuss that imposing the long-only constraint reduces estimation error. The corresponding trade-off is an increase in specification error. Imposing constraints is therefore likely to be helpful if there is significant estimation error present in the sample covariance matrix. They further show that imposing portfolio weight constraints produces performance results that are equivalent to the results obtained from using alternative methods aimed at reducing estimation error, such as factor models and shrinkage estimators.

[3] Scherer (2002) illustrates mathematically that robust optimisation is equivalent to a shrinkage estimator. He further discusses the fact that robust portfolio construction is no different form Jorion's (1986) Bayes-Stein shrinkage estimator, which interpolates between the minimum variance portfolio and the maximum Sharpe ratio portfolio. Shrinkage estimators are significantly easier to implement in practice.

Prior information can be interpreted as the odds an investor would be willing to accept if forced to bet on the true parameters before investigating the data (Avramov & Zhou, 2010). The prior distribution can represent information regarding particular events, macroeconomic news or any other additional information having an influence on the dynamics of the asset returns. The prior distribution is then combined with the sample information in order to form a predictive distribution of asset returns. The predictive distribution, incorporating the prior distribution, integrates out the parameter space and characterises the uncertainty about future asset returns. The optimal Bayesian portfolio is then obtained by maximising the expected utility with respect to the predictive distribution (Avramov & Zhou, 2010).

Supporters of the traditional frequentist approach will often criticise the Bayesian approach for lacking objectivity and as well as introducing bias into the estimates as a result combining a subjectively obtained prior distribution with an objectively obtained sampling distribution. This is due to the fact that the frequentist approach follows a strictly rigid and purely objective approach whereby a maximum likelihood approach is used to produce point estimates[4]. It must be noted that this approach is not appropriate for asset allocation as the point estimates are treated with 100% certainty, which is practically infeasible within the asset allocation context. In addition, parameter estimates obtained via the frequentist approach are known to exhibit over sensitivity, therefore sacrificing specification accuracy in order to avoid creating a bias in the estimates (Avramov & Zhou, 2010).

In addition, the frequentist approach restricts investors solely to the information within the sample and are unable to utilise any non-sample information; such as information based on professional experience or insights. In the context of asset allocation, Nobel laureate, Harry Markowitz (1959), stated that "The rational investor is Bayesian", thereby inferring that it would be irrational for investors to ignore additional information not available within the sample. Scherer (2007) states that investors will not be able to deal with the effects of parameter uncertainty if they are confined solely to the information available within the sample.

Michaud (2008) states that the use of historically estimated asset returns is generally unreliable and that in practise, asset managers often replace historically estimated average returns with exogenous estimates based on economic and fundamental views of markets and asset returns. In this way, asset managers act as "Bayesian agents" by improving the forecast value of historically estimated returns (Michaud & Michaud, 2008).

The following sections further describe the Bayesian approach to asset allocation. The aim of the following sections is to illustrate how the Bayesian approach bridges the gap between the art and science of asset allocation. The Bayesian approach allows investors to combine their unique views on the characteristics of asset returns in an intuitive manner, which effectively reflects the investor's confidence in the various sources of information via the Bayesian improved estimates.

---

[4] Maximising the likelihood function maximises the probability that the data have been generated from a distribution with the given estimated parameters.

2.2.1 Prior Probabilities in the Asset Allocation Context

In many Bayesian models, the prior distribution forms the avenue whereby the asset manager can express subjectively (or objectively) obtained views[5] and alter the historically estimated returns in light of the views or alternative sources of information.

A notable difficulty of the Bayesian framework is the particular choice of a prior distribution. Finding the correct method for constructing an objective prior has been the pursuit of many statistical theorists, such as Laplace, John Maynard Keynes and Harold Jeffreys, who have suggested various complex methods[6] for constructing objective priors (Lee, 2007). It must be noted that each of the methods proposed have been useful in the use of the Bayesian methodology, however, there is still no universal methodology for constructing prior distributions (Lee, 2007).

Asset managers therefore need to express either informative priors, thereby using relevant expertise, or choose diffuse priors, which do not express any information on the parameters but instead impose a level of uncertainty in the estimation. The following sections describe the range of prior distributions available to users of the Bayesian framework:

Informative Priors

Prior beliefs can be described as informative if they significantly alter the information contained in the data (Fabozzi *et al.,* 2007). The most common approach to constructing an informative prior is to specify a particular distribution for the unknown parameters as well as specify hyper parameters which appropriately reflect the beliefs of the asset manager. An example of an informative prior distribution would be the prior view that asset returns are expected to follow the generalised extreme value distribution[7]. In this particular case, the hyper parameters for the location, scale and shape parameters are used to describe the investor's view and confidence regarding the performance of the respective assets before the data has been observed.

Non-informative

In many cases, the situation arises whereby the investor is unsure of the distribution of asset returns or does not possess a particular view. Non-informative priors are designed to reflect uncertainty about the model parameters without significantly altering the posterior parameter distribution. It therefore follows that non-informative priors do not add information to the data and are also known as vague or diffuse

[5] One exception is the Black-Litterman model, which uses market equilibrium implied views as a prior distribution and investor views as a frequency distribution.

[6] The particular methods include: maximum entropy, transformation group analysis and reference analysis techniques.

[7] The generalised extreme value distribution is a family of continuous probability distributions, comprised of the Gumbel, Fréchet and Weibull distributions. Extreme value theory is directly suited for modelling the tail-end of return distributions – which is most appropriate for modelling the return distributions of hedge funds, which are known to exhibit fat left tails.

priors. Non-informative priors typically follow flat, uniformly distributions. The Jeffery's' prior is additional example of a popular non-informative prior and has the key feature that it is objectively obtained and is invariant under re-parameterisation of the parameter vector; which makes it especially appealing for its use with scale parameters (Fabozzi *et al.,* 2007).

Michaud (2008) describes a potential pitfall of the Bayesian approach in that a perverse (incorrect) prior distribution imposed on the parameter values may result in poor estimates of the asset parameters. In other words, if the investor specifies incorrect views on the parameter space, then the resulting posterior inferences will be significantly inferior compared to the estimating the parameters from the data alone. Michaud (2008) further states that the existence of a perverse Bayesian prior is a significant concern in statistical estimation theory. In this regard, Markowitz and Usmen (2003) use a non-informative prior when estimating Bayesian parameter values in the absence of investment information.

Conjugate Priors

A prior distribution is known as a conjugate prior distribution if it is of the same distribution as the posterior distribution. Therefore, a conjugate distribution does not alter the posterior distribution (Fabozzi *et al.,* 2007). For example, the normal distribution is said to be conjugate to itself as the posterior distribution will also be normally distributed. Conjugate priors are convenient as they allow for closed form analytical posterior expressions to be obtained. In the absence of a conjugate prior, a complex and unidentifiable posterior distribution may be obtained, for which numerical methods will be required in order to evaluate the distribution.

Given that the choice of a prior distribution significantly influences the posterior distribution, an informative prior should only be used when the investor exhibits useful and reliable information regarding the distribution of future asset returns. In contrast, when no reliable information is present, an uninformative prior should be used in order to express uncertainty regarding the asset returns.

2.2.2 Bayesian Model Averaging

Bayesian modelling techniques have long been used in the study of traffic congestion in the transport sector, heart attacks in medicine, as well as economic growth in macroeconomic literature. As a result of its strong rooting in decision theory, the Bayesian approach to asset allocation has shown significant growth in popularity in recent years. Avramov and Zhou (2010) state that there are three appealing characteristics of Bayesian techniques that have lead to its increasing popularity in the field of asset management, namely:

- The flexibility of the Bayesian approach to handle intricate and realistic financial models.

- Robust estimations[8] that explicitly incorporate estimation error.
- A theoretically sound framework that enables the optimal combination of various sources of information.

In the asset management context, Bayesian model averaging provides for a flexible framework for modelling investor uncertainty with regards to potentially relevant predictive factors in forecasting models. In particular, Bayesian model averaging assigns posterior probabilities to a wide variety of competing forecasting models and uses the respective probabilities in order to achieve a composite weighted model. The optimally weighted Bayesian averaged model is then used for asset allocation decisions. This contrasts significantly with the traditional frequentist approach for asset allocation model selection. The frequentist approach uses a particular model assessment criterion in order to select a single model, which is deemed to be the best model and is assumed to be correct. All other competing models are assumed to be incorrect and are discarded. In this way, model uncertainty is not accounted for (Avramov & Zhou, 2010). Bayesian Model Averaging therefore seeks to obtain an optimal combination of all competing sources of information. The following section introduces and describes the Black-Litterman Model, which is a popular Bayesian optimisation model.

2.2.3 The Black-Litterman Model

The Black-Litterman model, created by Fischer Black and Robert Litterman (1992), is a portfolio optimisation model that provides a Bayesian approach that enables investors to combine their unique views on specific assets with market equilibrium returns in a manner that results in intuitive and diversified portfolios. In the context of portfolio selection, Bayesian statistics allows us to impose a prior view, and then, upon the arrival of new data, to alter our view in light of the newly acquired information.

The Black-Litterman model can be further described as Bayesian benchmark-based optimisation model that effectively results in a portfolio that is a weighted average of the equilibrium expected return and the investor's views on the expected return. Depending on the confidence of the individual views of the investor, the Black-Litterman model tilts the optimal portfolio away from the market portfolio (benchmark weights) in the direction of the expressed views. The assumption is that the market portfolio weight vector is a sensible starting point which involves no extreme positions.

The Black-Litterman methodology can be summarised as follows:

[8] It must be noted that in the context of this thesis, the term "robust estimates" and "robust optimisation" refers to achievement of parameter estimates that remain stable over time and are not easily influenced by outliers. This is to avoid confusion with another optimisation procedure (also named Robust Optimisation) which specifically refers to maximising the worst case scenario.

- An equilibrium point for optimisation is introduced. This equilibrium point is the market portfolio of implied equilibrium excess returns[9]. The implied equilibrium returns forms the prior distribution and is a neutral starting point for the Black-Litterman model.
- The investor then forms views about the asset returns and assigns a confidence level to each of them. These views are unique to the investor and contrast the equilibrium views of the market. In the Bayesian context, the views of the investor form the frequency distribution of security returns and are used to update the prior implied equilibrium returns[10].
- The Black-Litterman methodology then combines the two sources of information, using a Bayesian methodology to form a posterior distribution of expected excess asset returns in light of the investor views. The posterior Black-Litterman expected returns can be interpreted as a weighted average of the implied market equilibrium and the investor views. The degree of the tilt between the implied equilibrium returns and the investor views depends on the volatility and correlations of each asset as well as the specified confidence level in the respective investor views. The greater the confidence in the views, the greater the extent to which the posterior portfolio tilts towards the investor views
- The posterior Black-Litterman expected returns, which combine the two sources of information, are then used to solve for the optimal portfolio weights using a standard optimiser.

The Black-Litterman model highlights the fact that because assets in a portfolio are correlated, adjustments to the expected excess returns involved in an active investment view should also result in adjustments to the expected excess returns of assets that are not explicitly involved in the active investment views (Herold, 2003). As a result, it has been shown that the Black-Litterman methodology leads to more stable, tractable and intuitive portfolios (Da Silva *et al.,* 2009).

Due to the fact that the Black-Litterman model uses a Bayesian approach, it can be noted that any estimation errors present in the data are spread out over all the assets in the portfolio. It therefore follows that the Black-Litterman expected return vector exhibits a lower sensitivity to estimation errors in the data. This is a very useful property as error maximisation in the final optimisation process is significantly reduced (Fabozzi *et al.,* 2007).

Many studies, inspired by the Bayesian framework, have been conducted in order to advance our understanding and implementation of the Black-Litterman framework. Lee (2000) and Satchell and

[9] Assuming that the market is in a state of equilibrium, the implied equilibrium returns are the returns that are assumed to clear the market, where demand equals supply. The implied equilibrium returns are assumed to be the returns that are optimal, given the observed market capitalisation weights.

[10] It can be noted that the Black-Litterman approach differs slightly from traditional Bayesian statistics. In traditional Bayesian statistics, the observed data forms the prior distribution and the subjective information forms the frequency distribution. The Black-Litterman model assumes that in the absence of investor information, the investor holds the market portfolio (as a prior portfolio) and then, as new information arrives, the portfolio allocations are updated accordingly.

Scowcroft (2000) extended the theoretical framework, while others (Bevan and Winkelmann (1998), He and Litterman (1999), Herold (2003), Idzorek (2005)) focused on the implementation of the Black-Litterman framework. The Bayesian approach of the Black-Litterman framework proves to be a key factor in obtaining portfolios that are more robust and are less sensitive to estimation errors present in the expected excess return inputs (Da Silva *et al.,* 2009).

In applying the Black-Litterman model to active management, the investor needs to exercise caution. Da Silva *et al.,* (2009) warns that changing the posterior weight optimiser to a mean-tracking error optimisation (typical of the active management objective), the investor could be exposed to the phenomenon of unintentional trades and increased risk. This is a direct result of the mismatch between objective functions of the "Sharpe ratio optimised" implied equilibrium returns and the "information ratio optimised" portfolio weights (Da Silva *et al.,* 2009)[11]. The following two sections illustrate the original Black-Litterman formulation as well as the new *active management* Black-Litterman model formulation, as proposed by Da Silva *et al.,* (2009).

The table below provides a description of the notations of respective components that will be used in expressing the Black-Litterman model in its original formulation as well as its active management formulation:

$\mathbf{E[R]}$	Is the new (posterior) vector of Black-Litterman returns (Nx1 Column Vector)
$\boldsymbol{\tau}$	Is the covariance shrinkage factor (a scalar)
$\boldsymbol{\Sigma}$	Is the covariance matrix of excess returns (NxN matrix)
$\mathbf{P}$	Is the view participation matrix that identifies the assets involved in the views (KxN matrix) for K views on N assets.
$\boldsymbol{\Omega}$	Is the diagonal covariance matrix of error terms from the expressed views. It represents the uncertainty in each view (KxK matrix)
$\boldsymbol{\Pi}$	The implied equilibrium return vector (N x 1 column vector)
$\mathbf{Q}$	Is the View Vector (K x 1 column Vector)
$\boldsymbol{\lambda}$	Is the Risk Aversion Coefficient
$\mathbf{w_{mkt}}$	The Market Capitalization Weight vector[12] (Nx1 Column Vector)
$\boldsymbol{w_{BL}}$	Final Black-Litterman recommended posterior weight vector
$\boldsymbol{w_{\alpha}}$	The active weights in excess of the benchmark weights

[11] The same holds true for other utility functions. In order to avoid information-less trades, the objective functions for the (prior) implied equilibrium returns and the posterior weight objective function are required to be consistent.
[12] An alternative to market capitalisation weights include presumed efficient benchmark weights.

The Original Black-Litterman Model

Using the notation above, the Black-Litterman posterior expected return vector can be expressed as follows:

$$E[R] = [(\tau\Sigma)^{-1} + P^T\Omega^{-1}P]^{-1}[(\tau\Sigma)^{-1}\Pi + P^T\Omega^{-1}Q] \quad (2.2.1)$$

The prior distribution of implied equilibrium excess returns, representing the implied views of the market, is obtained via the following reverse optimised formula:

$$\Pi = \lambda\Sigma w_{mkt} \quad (2.2.2)$$

The vector of implied equilibrium returns represents the expected returns that are optimal, given the available benchmark weights, risk aversion and covariance matrix.

Once the posterior expected return vector, given the market and investor views has been obtained, it is now possible to obtain the optimal Black-Litterman posterior weight vector using the following unconstrained optimisation formula[13]:

$$w_{BL} = (\lambda\Sigma)^{-1}E[R] \quad (2.2.3)$$

The above vector of posterior Black-Litterman portfolio weights are the optimal portfolio weights that consider the information supplied by the investor as well as information present in the market views of the implied equilibrium returns. As previously stated, the above recommended portfolio weights are not subject to any investment constraints, such as the long-only and tracking error constraints. The above posterior Black-Litterman portfolio weights satisfy the maximum Sharpe ratio objective of an unconstrained portfolio.

[13] The formula below is equivalent to the solution to the following unconstrained objective function:
$max(E[R] - \lambda w\Sigma \mathrm{w})$

Figure 1 below, adapted from Padberg (2007), graphically illustrates the Black-Litterman approach, whereby the implied equilibrium returns (Π) are combined with the investor views (Q) in order to form the Black-Litterman expected return vector ($E[R]$).

Figure 1: Components of the Black-Litterman Framework

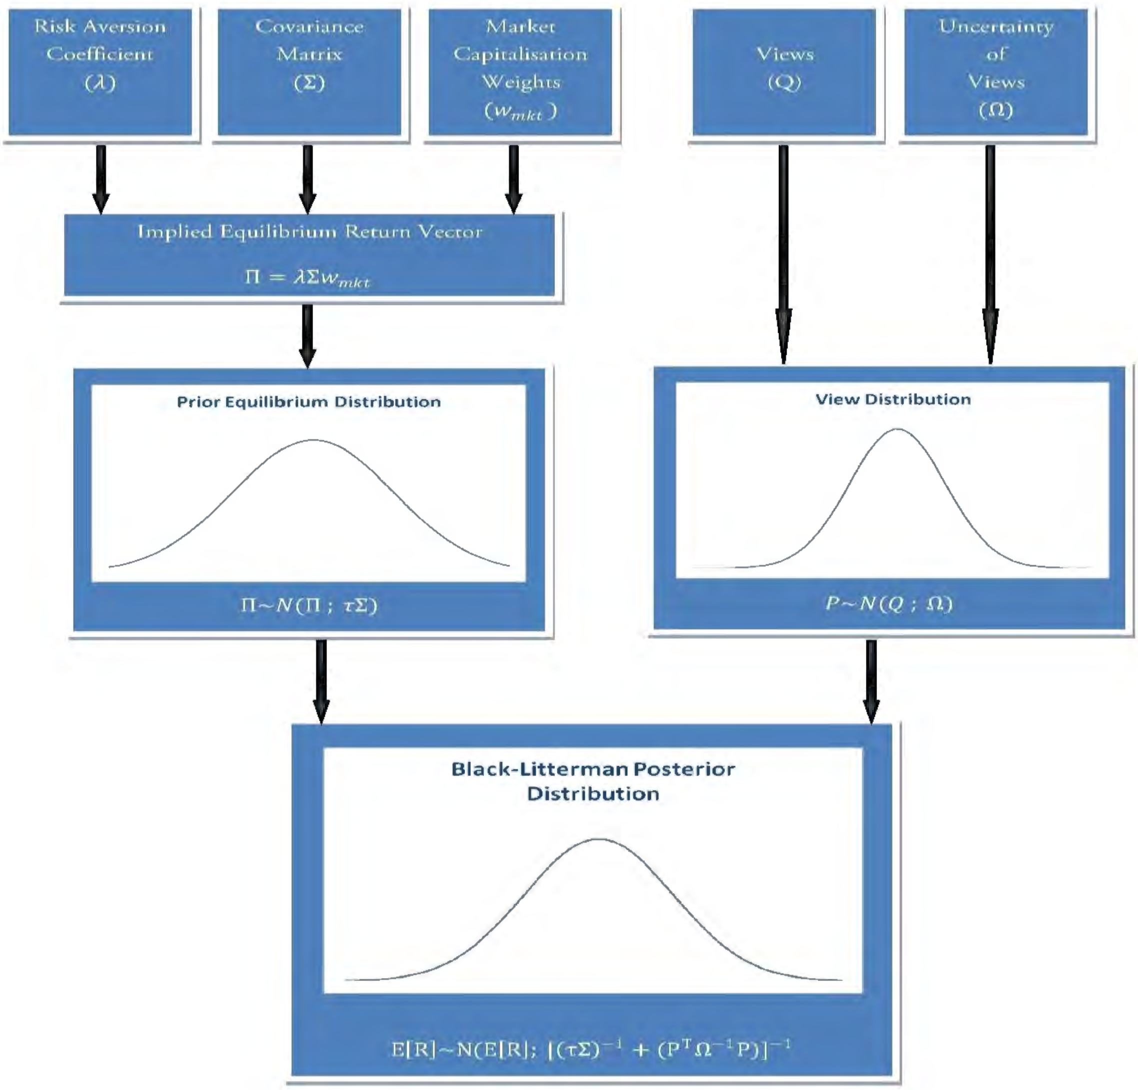

The Black-Litterman Model Applied to Active Management

It must be noted that the original Black-Litterman framework was developed under the maximum Sharpe Ratio objective, which follows a traditional mean-variance framework (Da Silva *et al.,* 2009). More specifically, the implied equilibrium returns are derived using a mean-variance objective function and the final posterior Black-Litterman weight vector was then obtained via the same mean-variance objective function. However, in the active management setting, investors are typically interested in maximising the portfolio alpha for a given level of tracking error, thereby satisfying the maximum *information ratio* criteria (Witten & Wilson, 2011). Da Silva *et al.,* (2009) illustrate the adverse effects that occur when a maximum information ratio objective function is naively applied to the posterior expected returns in order to obtain the final vector of poster Black-Litterman weights.

Da Silva *et al.,* (2009) point out that a mismatch between the implied equilibrium returns objective function and the poster weight vector objective function occurs when the two have separate objectives, namely, the maximum Sharpe ratio objective and the maximum information ratio objective. Returns that are optimal under the maximum Sharpe ratio objective are incorrectly perceived to be sub-optimal as per the maximum information ratio objective. Da Silva *et al.,* (2009) mathematically demonstrate that the straight forward implementation of the maximum alpha-tracking error objective function into the Black-Litterman methodology results in information-less trades and unintentional risk taking as a result of the mismatch in objective functions. In other words, the posterior portfolio optimiser will initiate active trades in order to "correct" the apparent sub-optimality, even in the absence of investor views.

For the purpose of illustration and simplicity, assume that no investor views are expressed, such that $E[R] = \Pi$. Therefore, if the "*Sharpe ratio optimised*" implied equilibrium returns are substituted into a maximum information ratio objective posterior weight optimiser, a mismatch between the objective functions occurs. The "*Sharpe ratio optimised*" implied equilibrium returns are by definition sub-optimal according to the maximum information ratio objective of the constrained posterior weight vector optimiser. The posterior optimiser then perceives the mismatch as an alpha opportunity and initiates active trades in order to optimise the implied equilibrium returns according to the information ratio objective. These active trades are unintentional and information-less since no investor views were expressed (Witten & Wilson, 2011). The same phenomenon of information-less active trades occurs for when investor views are expressed and is most prevalent for highly volatile securities in the portfolio (Da Silva *et al.,* 2009). As illustrated by Da Silva *et al.,* (2009), further inconsistencies that arise are increased portfolio volatility as well as additional (unintentional) exposure to the benchmark.

In order to rectify the problem of unintentional trades, it is necessary to ensure that both optimisation problems are consistent. This therefore requires the implied equilibrium returns to be backed out according to the same maximum information ratio objective as the optimiser used to construct the active portfolio. However, as illustrated by Da Silva *et al.,* (2009), in the presence of constraints, the optimisation process for Π seems to become very complicated. They further illustrated that if Π is set to any vector with its elements all the same, the optimisation problem simplifies and essentially becomes a tracking error minimization problem, regardless of the constraints imposed. They further illustrate that in the absence of investor views, the prior belief for the implied equilibrium excess returns should be an uninformative one whereby all securities in the portfolio are expected to have the same excess returns. The most intuitive uninformative prior would be to assume that all securities earn zero excess returns in the absence of investor views, therefore setting Π equal to the zero vector.

The posterior Black-Litterman *active* return vector, assuming an uninformative prior, can therefore be expressed as:

$$E[R] = [\tau\Sigma^{-1} + P'\Omega^{-1}P]^{-1}[P'\Omega^{-1}P] \tag{2.2.4}$$

As proposed by Herold (2003), the final optimal active allocations (w_α) can be extracted via the following (maximum information ratio) tracking error/alpha objective function:

$$\max_{w_\alpha}\left(w_\alpha' E[R] - \frac{\lambda}{2} w_\alpha' \Sigma w_\alpha\right) \tag{2.2.5}$$

Subject to the following constraints:

$w_\alpha' \Sigma w_\alpha \leq a^2$ (Tracking Error (a) Constraint)
$\sum_{i=1}^{s} w_{\alpha i} = 0$ (Active weights sum to zero)
$w_{mkt} + w_\alpha = w_{BL} \geq 0$ (Long-only Constraint)

The final Black-Litterman portfolio weights are given as the summation of the market capitalisation weights and the resulting optimal active allocations $(w_{BL} = w_{mkt} + w_\alpha)$, where w_α represent the active allocations.

The above discussion highlights the fact that implementing the Black-Litterman model in practice is not straightforward. In practice, many portfolios are subject to various investment constraints and direct application of the Black-Litterman model can cause unintentional trades and risk taking. In contrast to the original unconstrained Black-Litterman model, the constrained active management model assumes an uninformative prior for Π, which solves the mismatch between objective functions.

2.2.4 Bayesian Decision Science and Behavioural Finance

One of the basic mechanisms of learning is collecting, analysing and utilising new information arriving from the external sources; this is the basic underpinning of the Bayesian decision-making framework. A Bayesian decision maker learns by adjusting past estimates in light of new data that becomes available. Bayes' theorem therefore provides a means of implementing the learning mechanism by combining knowledge about the distribution of asset returns with the information regarding the parameters enclosed in the data (Avramov & Zhou, 2010).

As previously discussed, the Black-Litterman model is a Bayesian portfolio optimisation model that requires investment input in the form of views (estimates and judgements) on the returns of the assets in the portfolio. Mankert (2006) notes that there is a scarcity of literature available which studies the behaviour of the users of the model as well as the context in which the Black-Litterman model is applied.

Traditional finance theory (Markowitz model) assumes that investors exhibit quadratic utility. More specifically, utility is defined in absolute terms whereby investors evaluate investment portfolios in terms of absolute wealth and therefore have no reference point for which portfolio returns are evaluated. This particular utility function differs both in shape and domain in which Tversky and Kahneman (1984) define utility. Tversky and Kahneman (1984) state that investor utility is not defined in absolute terms but rather as gains and losses relative to a point of reference. The relative loss aversion that the investor experiences, as implied by the relative utility function, is defined as the aversion to achieving a return lower than the benchmark portfolio. The magnitude of the relative loss aversion is greater than the positive psychological effects of corresponding gains. This implies that investors would frequently choose to keep the benchmark allocations out of fear of obtaining inferior relative performance (Montier, 2002).

In the context of the Black-Litterman model, the market (or benchmark) portfolio acts as a point of reference for which investors portfolio deviate according to the specified views and associated view confidence levels. This property of the Black-Litterman model is very appealing due to the fact that many financial managers' portfolio performances are evaluated relative to a reference point, such as a benchmark portfolio. Similarly, portfolios recommended by the Black-Litterman model are known to be more realistic[14] and intuitive due to the model's consistency with the relative utility function of that proposed by Tversky and Kahneman (1984).

Howson and Urbach (1989) show that the Bayesian approach improves upon traditional statistical approaches in aiding the decision-making process. While the Bayesian approach often subjectively determines prior probabilities, Howson and Urbach (1989) argue that the Bayesian approach evaluates evidence in such a manner that users with opposing views will most likely approach a common view as the evidence is accumulated and incorporated into the Bayesian framework. The result of such a procedure is the reduction of the affects of initial biases and heuristics inherent in human decision making.

As a result of requiring decision makers to utilise base probabilities, the Bayesian methodology requires a reasoning process that causes decision makers to thoroughly evaluate all sources of information. This therefore provides for a sound reasoning process which yields results that are easier to understand and interpret. Zlotnick (1972) summarises three benefits of using the Bayesian approach as a decision support system:

- Decision makers are forced to quantify judgements for which they would not usually assign numerical values.
- The decision maker is able to evaluate evidence against a particular hypothesis as opposed to treating evidence as given and therefore making impulsive conclusions as a result.

[14] This is as opposed to the portfolios recommended by the traditional Markowitz model, which is known to generate portfolios exhibiting extreme positions that are highly concentrated in a few assets.

- The most appealing feature of the Bayesian approach is that it considers the information segment by segment. This therefore avoids the scenario whereby the decision maker is forced to consider huge amounts of information at once and thereby basing decisions on information that is yet to be fully explored. Instead, the mathematics involved in the Bayesian methodology does the summing up and tells the analyst "If these are your judgements (or readings) on the individual items of information, then the conclusion is as follows...".

From the above discussion, it is therefore evident that the Bayesian framework allows asset managers to observe how each individual analyst evaluated and weighted the importance of their data as well as the manner in which they arrived at their respective conclusions. The Bayesian methodology therefore provides a transparent system of accounting for analytical judgements that can be achieved. In summary, the Bayesian methodology helps ensure that information, both supporting and contradictory are not overlooked.

Chapter 3: Robust Statistics

3.1 Overview of Robust Statistics

The performance and value of an optimised portfolio largely depends on the efficient estimation and optimisation of the respective input parameters (Michaud, 2008). Asset managers usually base the portfolio risk and return estimates on sample means, standard deviations, and correlations, computed from historic data and then make post-hoc adjustments for current information (Michaud, 2008). This multivariate estimation of the sample parameters is therefore very likely to result in significant estimation errors whereby the sample estimates deviate significantly from the true parameter estimates. Fabozzi *et al*., (2007), describe the "garbage-in-garbage-out" principle whereby an optimised portfolio based on poor parameter estimates is likely to result in a portfolio that is sub-optimal.

A notable problem with classical estimators, such as the sample mean and covariance, is that they are very sensitive to outliers in the data and are known to change significantly for small changes in the underlying data (Fabozzi *et al.,* 2007). The objective of robust estimation is to achieve estimates that are relatively *insensitive* to small changes in the underlying data or the particular sample choice. A *robust* parameter is one which is stable and insensitive to small changes in the data[15]. Robust statistics is a technique used to find models that are robust in the sense that the same results will be achieved even if the samples change or if the respective underlying assumptions regarding the shape of the distribution are not correct (Fabozzi *et al.,* 2007). Robust statistics can be viewed as a method used to separate the contribution of the majority of data from that of the extremes. This is particularly relevant in finance, as many asset return distributions are known to exhibit fat left tails (Hwang & Satchell, 1999). Robust statistics therefore allows the portfolio manager to discriminate between usual and extreme events present in the data. As described by Fabozzi *et al.,* (2007), an investor seeking a robust portfolio is one who would prefer a portfolio that performs well under a number of different scenarios and is well-protected from estimation and model risk. A robust portfolio is one that is assured to perform well even if the optimisation model is severely misspecified.

Estimation error has always been a significant problem in portfolio construction (Scherer, 2007). As previously discussed, a problem with the Markowitz mean-variance optimisation model is that it is known to be overly sensitive to the respective inputs. The mean-variance model is known to be unstable in the sense that small changes in the inputs can result in significant changes to the optimal allocations of the portfolio, often resulting in extreme allocations and poor out-of-sample performance. As a result of the adverse effects of estimation errors on the optimal portfolio, portfolio optimisation using the classical

[15] To better understand the concept of robustness, consider the sample mean and median. The sample mean is not a robust statistic as it is very sensitive to outliers. A change in a single observation can have an uncontrolled effect on the mean while the median is insensitive to changes of up to half of the sample.

mean-variance framework has been termed "error maximisation" (Michaud, 1989). Michaud (1989) argues that the classical mean-variance optimisation tends to overweight assets with with a large ratio of estimated return to estimated variance and that these assets are the most likely to have large estimation errors. Due to the significant problem of estimation error and parameter uncertainty in traditional portfolio optimisation, robust statistics plays an important role in portfolio optimisation. Parameter uncertainty can be regarded as an additional form of risk to the investor, which is ignored by classical portfolio optimisation models as the parameters are typically treated with 100% certainty.

In addition to Bayesian estimators described in previous sections, portfolio resampling is another popular robust optimisation methodology that aims to produce an improved optimiser which accounts for parameter uncertainty. As described by Scherer (2007) and Michaud (2008), various methodologies have been devised in an attempt aimed at reducing the effects of estimation error on the optimal portfolio allocations, with the bulk of research directed at improving the respective inputs before inserted into an optimiser. Michaud (2008) suggests that an improved optimiser is an additional method for improving portfolio performance in the presence of estimation error. There are various optimisation models available that explicitly incorporate estimation error directly into the optimisation process. Michaud (2008) further states that it is possible to combine both methodologies (refined inputs and an improved optimiser) to result in a more efficient and superior portfolio. A description of portfolio resampling and its comparison to Bayesian methods is included in Appendix B. The following section briefly describes the Bayesian approach to obtaining robust parameter estimates.

3.2 The Bayesian Approach to Parameter Uncertainty

As discussed by Scherer (2007), the effect of parameter uncertainty and estimation error on the optimal portfolio choice cannot be appropriately resolved by focusing solely on the available sample information. He further discusses the fact that additional information, in the form of investor experience, is required to overcome the problem and that it would be irrational to ignore this additional information. The optimal combination of sample and non-sample information can be achieved using Bayesian statistics.

Bayesian statistics treats parameter values as random variables that follow a particular distribution. In this way, Bayesian statistics is used to combine sample information with non-sample information. The non-sample information is known as prior information and can be interpreted as the investor's probability estimates of the parameters before observing the data (Avramov & Zhou, 2010). This particular methodology proves to be extremely useful in portfolio choice when the exact value of the inputs is uncertain. In contrast, the frequentist approach often uses maximum likelihood methods to create point estimates of the parameters, where maximising the likelihood function maximises the probability that the data has been drawn from a particular distribution with the estimated parameters. Estimates of the parameter values are treated as 100% certain as per the frequentist approach, which is a poor assumption as in practise, parameter values are typically very uncertain.

Prior knowledge that represents the investors' unique insights of the expected return, especially when future returns cannot be elicited from sample information, becomes extremely useful in investment management. The Bayesian framework therefore provides portfolio managers with the ability of utilising external sources of information. Therefore, Bayesian methods in portfolio management is a manner in which portfolio managers can better control the quantitative framework of the optimisation process via subjective involvement (Fabozzi *et al.,* 2007).

As described above, Bayesian methods can be used to achieve superior risk and return estimates that incorporate parameter uncertainty as well as incorporate the investor's unique insights on the data. Bayesian methods are typically used to provide superior risk and return estimates that perform better in practise as opposed to raw sample information that is known to contain a high degree of estimation error.

3.2.1 Shrinkage Methods

Shrinkage estimation is a subset of the Bayesian approach to obtaining robust parameter estimates (Avramov & Zhou, 2010). Shrinkage estimation aims to improve raw estimates by combining sample based estimates with additional information in order to achieve final estimates that exhibit decreased estimation error (Fabozzi et al, 2007). More specifically, shrinkage methods aim to "shrink" sample estimates towards some specified target estimator obtained from additional information not present in the sample. A shrinkage estimator is a weighted average of a sample estimator and another, more structured estimator[16]. The resultant shrinkage estimator ($\hat{\mu}_{shrink}$) is dependent on the sample estimator ($\hat{\mu}$), the shrinkage target (μ_0) and the level of shrinkage intensity (ϕ) (Fabozzi et al, 2007). Shrinkage estimators, although not unbiased, possess more desirable properties than the sample estimator. The general form of a shrinkage estimator can be written as:

$$\hat{\mu}_{shrink} = (1 - \phi)\hat{\mu}_{sample} + \phi\mu_{target}. \qquad (3.2.1)$$

The shrinkage target should have some basic properties in common with the quantity being estimated and should typically be robust with a lot of structure (Fabozzi et al., 2007). The shrinkage intensity is the relative weighting that is placed on the shrinkage target and is dependent on the degree of accuracy involved in the sample estimates. The shrinkage problem is therefore to identify the optimal trade-off between the bias (of the target estimator) and the estimation error (of the sample estimator). In particular, the shrinkage intensity therefore defines the optimal trade-off between the bias and estimation error present in the

[16] In particular, shrinkage methods aim to combine an unbiased (but very variable) sample covariance estimate towards a biased (but less variable) estimate in a optimal manner in order to obtain a more efficient estimator.

target and sample estimates. The greater the shrinkage intensity, the greater the structure imposed in the shrinkage estimator.

In addition to shrinkage estimators for the sample mean, various shrinkage estimators involving factor models have been proposed for improving the efficiency of the sample covariance matrix (Wolf, 2001) as well as for higher co-moments (Martellini & Ziemann, 2010). In particular, Ledoit and Wolf (2001) suggest that the sample covariance matrix alone should not be used for the purpose of portfolio optimisation due to the extent of the estimation error present in the data. In place of the sample covariance matrix, they suggest that a shrinkage-enhanced covariance matrix estimator should be used, thus imposing structure on the estimator[17].

There are many different types of shrinkage estimators proposed in literature. For example, the Bayes-Stein shrinkage estimator by Jorion (1986) is one of the most well-known shrinkage estimators and is used to estimate expected returns (Fabozzi et al., 2007). The Bayes-Stein estimator shrinks the sample expected returns towards the minimum variance portfolio and has been shown to result in decreased portfolio weight variability as well as superior out-of-sample risk adjusted performance, as compared to the sample mean (Jorion, 1986). An additional shrinkage estimator (for expected returns), as described in section 2.2.3, is the Black-Litterman model where the market implied equilibrium returns form the prior distribution. The Black-Litterman model uses the market implied equilibrium returns as a neutral starting point for the portfolio optimisation. Using the Bayesian methodology, the investor views are shrunk towards the market implied returns. This has shown to result in intuitive and diversified portfolios (Idzorek, 2005).

3.2.2 Pitfalls of the Bayesian Approach

Despite the notable success of the various shrinkage methods as well as the overall Bayesian approach, there are some major practical difficulties involved in specifying a particular shrinkage target (Wolf, 2001) or Bayesian prior. Scherer (2007) points out that in order to result in improved estimates, an appropriate prior/shrinkage target needs to be defined; as an incorrect or perverse prior may result in worse estimates than only using the information available in the data (Michaud & Michaud, 2004). For a shrinkage target to be effective, the assets in a portfolio are required to exhibit similar characteristics in order to be combined with a particular prior distribution (Scherer, 2007). This becomes particularly difficult when a multi-asset portfolio is involved. In the context of the Black-Litterman model, an appropriate prior distribution may be difficult to specify when no benchmark portfolio exists.

Appendix B further describes additional robust portfolio optimisation methodologies available and compares them to the Bayesian approach by discussing the respective strengths and weaknesses of each approach.

[17] The target estimator is typically estimated via a low-dimensional factor model with uncorrelated residuals. The lower the number of factors used to estimate the target covariance matrix, the stronger the structure imposed.

Chapter 4: Modelling Investor Utility

In the presence of non-normally distributed asset returns and non-quadratic investor utility preferences, an optimal portfolio selection technique requires not only the first two moments (mean and variance), but also estimates of higher order moments and co-moments of the return distribution (Martellini & Ziemann, 2010). Therefore, it can be noted that the mean-variance utility function of the Markowitz model is only appropriate when asset returns follow a normal distribution and when investors are risk neutral to skewed returns and abnormal kurtosis. However, Martellini and Ziemann (2010) discuss the fact that research has shown that typical investors exhibit non-trivial utility preferences, whereby many investors would accept a lower mean return for positive skewness and a lower kurtosis.

Many portfolio selection models are based on the assumption of normally distributed returns as it significantly simplifies the optimisation procedure by allowing the return distribution to be fully described by the first two moments, namely the mean and variance. However, in practise return distributions in many financial markets do not follow a normal distribution, but instead are known to be negatively skewed and exhibit fat left tails (Harvey & Siddique, 2000). For example, hedge funds typically exhibit return distributions that are non-normal and possess a significant positive or negative skewness with a high kurtosis (Popova *et al.,* 2007). A portfolio optimisation model based on the standard mean-variance trade-off will therefore be insufficient in accurately determining the optimal utility of the investor as well as capturing the unique skewness and kurtosis present in the return distribution. As a consequence of the mean-variance framework ignoring the presence of higher-order moments, investors who are concerned about skewness and kurtosis would be extremely sceptical with regards to the unrealistically high allocation towards hedge funds that standard mean-variance optimisation models typically suggest (Popova *et al.,* 2007). It is therefore easy to see that the standard mean-variance framework will be insufficient in quantifying the risks of assets in a portfolio that exhibit non-normal return distributions (Fabozzi *et al.,* 2007). In other words, mean-variance models penalise portfolios that possess occasional positive return surprises while significantly underestimating the asymmetric downside risk of the portfolio.

In order to more appropriately capture the information inherent in the data and to better model the investor utility, Fabozzi *et al.,* (2007) suggest that higher moments of the return distribution be incorporated into the utility function. Higher moments, such as the kurtosis and skewness coefficients will better capture the risk aversion characteristics of the investor. These higher moments will, for example, capture the investor's aversion to extreme losses as well as penalise assets in the portfolio that exhibit severe negative skewness, thus exhibiting an increased probability of a negative return.

Various utility functions aimed at better modelling the investor's utility have been proposed. Popova *et al.,* (2007) suggest including a variety of benchmark-based objectives dependent on the respective scenario and tastes of the investor. In particular, an investor's utility function could be specified as maximising the probability of obtaining returns above a particular benchmark return while minimising some pre-specified risk measure.

Examples of alternative risk measures that have been proven to be particularly useful in modelling an investor's aversion to risk are the portfolio Value-at-Risk (VaR) and Conditional Value-at-Risk (CVaR)[18]. These risk measures are typically aimed at minimising the portfolio extreme downside risk and extensively used by investment banks (Harvey & Siddique, 2000). The following sections deal with various utility functions that are appropriate for a number of investment scenarios.

4.1 Benchmark Relative Utility Functions

The traditional use of utility functions involves maximising return while minimising some measure of portfolio risk. While this particular objective is appropriate for many situations, investors often exhibit other preferences for portfolio characteristics. For example, many investors are interested in maximising the probability of beating a particular benchmark by a certain percentage. In practise, particular investment funds are created with the sole objective of outperforming the CPI index, thereby achieving a positive real *return* on their investment. In addition to optimising returns relative to a benchmark, it is possible to optimise a portfolio relative to a number of benchmark portfolios as a result of combining multiple utility functions. Using Monte-Carlo simulation, Popova *et al.,* (2007) propose the following dual benchmark optimisation:

$$\max_{w} U = P(w'E[R] \geq b_1) - \lambda E[b_2 - wE[R]]^+ \qquad (4.1.1)$$

Where $P(w'E[R] \geq b_1)$ represents the probability of outperforming the benchmark return b_1 and $E[b_2 - wE[R]]^+$ represents the minimum shortfall of the portfolio beyond another benchmark b_2 where in this case, λ represents the priority. The value of λ can be varied in order to obtain an efficient frontier, which illustrates the optimal trade-off for one benchmark optimisation over another. The values for b_1 and b_2 can be set so as to achieve any particular objective. For example, b_1 can be set to outperform the FTSE/JSE All Share Index, while b_2 can be set to zero in order to strive to obtain a positive return.

The inclusion of alternative utility functions such as the benchmark based utility function above is that investors are able to express a broader and more concise picture of their risk and return objectives whenever a particular benchmark is involved. In addition to the above utility function, risk aversion to higher moments of the return distribution may also be included (Popova *et al.,* 2007).

[18] It can be noted that CVaR (as opposed to VaR) is an attractive risk measure as it is a coherent risk measure that satisfies properties of monotonicity, sub-additivity, homogeneity, and invariance.

4.2 Portfolio Selection with Higher Moments

The mean-variance optimisation framework proposed by the Markowitz model assumed that investors preferred higher expected returns for the lowest level of portfolio risk. The Markowitz model assumes that returns are normally distributed and that investors are not concerned about the skewness and kurtosis of the returns (Markowitz, 1959). However, in practise, asset returns are known to exhibit significant negative skewness as well as the existence of fat left tails- indicating large negative losses occurring at the tail-end of the distribution (Martellini & Ziemann, 2010). Intuitively, investors are risk averse towards negative skewness and would prefer positively skewed asset returns. In the study of Mutual funds, Levy and Sarnat (1984) find a strong preference for positively skewed returns. Harvey and Siddique (2000) introduce conditional skewness into an asset pricing model and show that investors may be willing to expect a negative return in order to achieve positive return skewness. This is due to the fact that investors are willing to trade a portion of their expected return in order to decrease the probability of experiencing a large reduction value of their portfolio.

In theory, higher-order moment utility functions should always lead to improved investor utility, as their investor preferences would be more accurately estimated. However, in practise, estimation of the additional moments results in a severe dimensionality problem (Martellini & Ziemann, 2010). For example, optimising a portfolio of twenty assets would require the estimation of 210 variance-covariance parameters, 1540 skewness-coskewness[19] parameters and 8855 kurtosis-cokurtosis parameters; which makes one question whether implementing higher order moments into a utility function would be feasible given the degree of estimation error involved. In order to decrease estimation error present in the moment estimators, Martellini and Ziemann (2010) provide shrinkage methods for improved estimation of the higher moments and demonstrate how the implementation of higher order moments dominates mean-variance optimisation from an out-of-sample perspective.

Following the methodology of Martellini and Ziemann (2010), the following sections outline how robust estimates of the higher co-moments can be computed using a Bayesian shrinkage approach. In addition, the following sections illustrate how the robust higher co-moment estimates can be included in the portfolio utility function.

[19] Co-skewness can be interpreted as correlated extreme returns. An aversion to negative co-skewness implies that investors are willing to reduce their average wealth level in order to decrease the probability of a large negative loss occurring. Co-Skewness describes the “skewness” relationship between two assets. Similarly, co-kurtosis describes the co-tendency of the assets to be correlated at the extremes as well as to be centred about the mean.

4.2.1 Four Moment Utility Function Decomposition

As briefly described above, a non-normal return distribution can typically be better explained via the inclusion of higher-order moments in the utility function. The first two moments (mean and variance) describe the location and spread of the asset returns. The third moment, portfolio skewness, describes the degree of asymmetry that occurs in the probability distribution, thereby better describing the lengths of the tails of the distribution. The fourth moment, kurtosis, describes the level of "peakedness" of the probability distribution. Similar to the concept of covariance between assets in a portfolio, the co-skewness and co-kurtosis parameters describe how the assets are correlated at the extremes as well as the co-tendency of assets to be centred about the mean.

The portfolio kurtosis can additionally be considered as a measure of the shape of the distribution and, more specifically, a measure of the relative weight of the tails of the distribution. For example, a distribution with a high kurtosis will have a high peak with the majority of the data situated close to the mean and a variance that is typically composed of more values that are extreme. A significantly large kurtosis therefore implies that the distribution contains significant fat tails (Fabozzi *et al.*, 2007). A rational investor therefore prefers positive skewness and lower kurtosis.

In order to evaluate the impact of higher order moments on portfolio selection techniques, a standard expected utility framework is considered. Assuming an infinitely differentiable utility function U, the investor's terminal wealth (W) can be approximated using the following Taylor series approximation:

$$U(W) = \sum_{k=0}^{\infty}\left[\frac{U^{(k)}\big(E(W)\big)}{k!}(W - E(W))^k\right]$$

It therefore follows that expected utility for the first four moments can be approximated as[20]:

$$E[U(W)] \approx U\big(E(W)\big) + \frac{U^{(2)}\big(E(W)\big)}{2}\mu^{(2)} + \frac{U^{(3)}\big(E(W)\big)}{6}\mu^{(3)} + \frac{U^{(4)}\big(E(W)\big)}{24}\mu^{(4)}$$

Hwang and Satchell (1999) illustrate that investors typically exhibit positive preferences for positive skewness and lower kurtosis for asset returns[21]. The portfolio choice can now be regarded as tangency points in a four-dimensional space and is no longer a trade-off between expected return and standard deviation of the asset returns.

---

[20] With $U^{(k)}$being the k^{th} derivative of the utility function and $\mu^{(k)}$ being the k^{th} order centred moment, represented as $\mu^{(k)} = E((W - E(W))^k)$

[21] It can be noted that the signs of the respective Taylor series expansion illustrate the investor's preference for lower volatility, positive skewness and lower kurtosis.

$$\frac{U^{(2)}\big(E(W)\big)}{2}\mu^{(2)} < 0 \qquad \frac{U^{(3)}\big(E(W)\big)}{6}\mu^{(3)} > 0 \qquad \frac{U^{(4)}\big(E(W)\big)}{24}\mu^{(4)} < 0$$

The moments in the above expected utility expression are a function of the portfolio weights as well as the first, second, third and fourth moments of the individual asset return distributions. In order to ensure that the relationship is more explicit, Martellini and Ziemann (2010) introduce higher-order moment tensors for the assets under consideration. They define the higher order tensors[22] as:

$$M_2 = E\left[\left(R - E(R)\right)\left(R - E(R)\right)'\right] \tag{4.2.1}$$
$$M_3 = E\left[\left(R - E(R)\right)\left(R - E(R)\right)' \otimes \left(R - E(R)\right)'\right]$$
$$M_4 = E\left[\left(R - E(R)\right)\left(R - E(R)\right)' \otimes \left(R - E(R)\right)' \otimes \left(R - E(R)\right)'\right]$$

Where M_2 is the standard variance-covariance matrix and $\otimes$ is the Kronecker product. It therefore follows that the moments of the individual asset returns can be written as functions of the higher moment tensors and the vector of asset weights:

$$\mu^{(2)} = w'M_2w \tag{4.2.2}$$
$$\mu^{(3)} = w'M_3(w \otimes w)$$
$$\mu^{(4)} = w'M_4(w \otimes w \otimes w)$$

Using the above expressions for the moments, the utility function can be re-written as:

$$E[U(W)] \approx U\left(E(\mu'w)\right) + \frac{U^{(2)}\left(E(\mu'w)\right)}{2} w'M_2w + \frac{U^{(3)}\left(E(\mu'w)\right)}{6} w'M_3(w \otimes w) + \frac{U^{(4)}\left(E(\mu'w)\right)}{24} w'M_4(w \otimes w \otimes w)$$

Where μ represents the mean vector of asset returns.

4.2.2 Estimating the Four Moments

It must be noted that estimation of the above four moments becomes increasingly difficult when many assets are included in the portfolio. Martellini and Ziemann (2010) state that large amounts of data for each asset in the portfolio will be required in order to ensure the number of asset return observations exceeds the number of parameters to be estimated. For example, a portfolio of twenty assets would require 45 years of monthly data in order to ensure that the number of observations exceeds the number of parameters required to estimate the moments in the utility model.

In order to obtain improved higher order moment estimators, Martellini and Ziemann (2010) analyse and compare the performance results on portfolios for using various types of structural and shrinkage estimators for the higher order moments. Their findings suggest that a combination of structured and shrinkage estimators results in significantly improved moment estimators for portfolios that are robust

[22] The third moment tensor (Co-Skewness tensor) can be pictured as a symmetric three-dimensional cube with height, width, and depth equal to the number of assets in the portfolio. The fourth moment tensor (Co-kurtosis tensor) can similarly be pictured as a cube in the four dimensional space.

over time. In addition, the use of improved four-moment estimators has shown to always outperform the standard mean-variance optimisation when high risk aversion to skewness and kurtosis as well as non-normality is observed. However, it must be noted that their findings also show that the four moment optimisation significantly underperforms the standard mean variance optimisation *when sample estimates*[23] *are used*.

In order to reduce the dimensionality of the required estimation of the higher order moments, Martellini and Ziemann's (2010) findings suggest the shrinkage approach, proposed by Ledoit and Wolf (2003), with a single factor model as a shrinkage target[24], imposes the necessary structure on the sample estimates and results in estimates that are robust over time and result in superior outperformance[25].

Single Factor Model Shrinkage Target

The single factor model, introduced by Sharpe (1963), summarises asset returns by the following regression equation:

$$R_{it} = c + \beta_i F_t + \varepsilon_{it} \tag{4.2.3}$$

$$\varepsilon_{it} \sim N(0\,;\,\psi)$$

Where ε_{it} is the residual term and the factor F_t is assumed to be the benchmark index of asset returns. β_i is the vector of regression coefficients. The usual regression assumptions apply, whereby the ε_{it} are independently normally distributed with a mean of zero and a constant variance (ψ).

Following the methodology provided by Martellini and Ziemann (2010), the above regression expression can be substituted into equation set 4.2.1 in order to obtain the following single factor moment tensor expressions:

$$\begin{aligned} M_2^{SF} &= \beta\beta'\mu_0^{(2)} + \psi \\ M_3^{SF} &= E[(\beta\bar{F} + \varepsilon)(\beta\bar{F} + \varepsilon)'\otimes(\beta\bar{F} + \varepsilon)'] \\ &= \beta\beta'\otimes\beta'\mu_0^{(3)} + \Phi \\ M_4^{SF} &= E[(\beta\bar{F} + \varepsilon)(\beta\bar{F} + \varepsilon)'\otimes(\beta\bar{F} + \varepsilon)'\otimes(\beta\bar{F} + \varepsilon)'] \\ &= \beta\beta'\otimes\beta'\otimes\beta'\mu_0^{(4)} + \Upsilon \end{aligned} \tag{4.2.4}$$

Where $\bar{F}$ is the the vector of centred benchmark returns ($\bar{F} = F - \mu_0$) and $\mu_0^{(k)}$is the k^{th} centred moment of the benchmark returns.

[23] The four-moment optimisation underperforms the mean-variance optimisation when structured estimates are not used and instead, sample estimates are used.

[24] The single-factor approach significantly reduces the dimensionality of the four-moment estimation. For example, for a portfolio of 20 assets, a total of 83 parameters are required to be estimated – compared to 10605 parameters if the sample estimator is used.

[25] This particular approach was also found to outperform the constant correlation counterpart, which was shown to result in significant misspecification error.

The values of ψ, Φ and Υ are interpreted and expressed as follows:

- ψ is a NxN covariance matrix of residual returns and is given by the sample estimate:

$$\psi_{ii} = E\left[\varepsilon_i^2\right]$$
$$= \frac{1}{T}\sum_{t=1}^{T} \hat{\varepsilon}_{it}^2$$
$$\psi_{ij} = 0\ for\ all\ i \neq j$$

- Φ can be interpreted as an NxN^2 return co-skewness matrix, whereby the individual skewnesses are represented by the third-order sample moments on the super diagonal of the matrix. Due to the independence assumption of the regression residuals, all off-super diagonal elements have an expectation of zero. The co-skewness matrix can be written as:

$$\Phi_{iii} = E\left[\varepsilon_i^3\right]$$
$$= \frac{1}{T}\sum_{t=1}^{T} \hat{\varepsilon}_{it}^3$$
$$\Phi_{iij} = 0$$
$$\Phi_{ijk} = 0\ for\ all\ i \neq j \neq k$$

- Υ is represented by an NxN^3 matrix, whereby the individual asset return kurtosis' are represented on the super diagonal of the matrix. The entries for the Υ matrix are as follows:

$$\upsilon_{iiii} = E\left[\varepsilon_i^4\right]$$
$$= \frac{1}{T}\sum_{t=1}^{T} \hat{\varepsilon}_{it}^4$$
$$\upsilon_{iiij} = 3\beta_i\beta_j\mu_0^{(2)}\psi_{ii}$$
$$\upsilon_{iijj} = \beta_i^2\mu_0^{(2)}\psi_{jj} + \beta_j^2\mu_0^{(2)}\psi_{ii} + \psi_{ii}\psi_{jj}$$
$$\upsilon_{iijk} = \beta_j\beta_k\mu_0^{(2)}\psi_{ii}$$
$$\upsilon_{ijkl} = 0\ for\ all\ i \neq j \neq k \neq l$$

Substituting the above expressions for the moment tensors into equation set 4.2.2, the following expressions for the moments, using the single factor approach, can be expressed as:

$$\mu^{(2)} = w'(\beta\beta'\mu_0^{(2)} + \psi)w \qquad (4.2.5)$$
$$\mu^{(3)} = w'(\beta\beta'\otimes\beta'\mu_0^{(3)} + \Phi)w\otimes w$$
$$\mu^{(4)} = w'(\beta\beta'\otimes\beta'\otimes\beta'\mu_0^{(4)} + \Upsilon)w\otimes w\otimes w$$

The above methodology illustrates how a single factor model for the asset returns can be used to significantly reduce the dimensionality involved in estimating the moments of the asset return distributions. As previously discussed, the single factor approach to estimating the four moments results in estimates that are well structured and result in portfolios that exhibit significantly less estimation error and are robust over time. The following section illustrates how the single factor moment estimates can be used as the shrinkage targets for which shrinkage estimates can be obtained via shrinking the sample estimates towards the structured estimates.

Shrinkage Estimates

As a result of the imposed structure of the single-factor model, the structured estimates exhibit a lower level of estimation risk. However, due to the artificial structure imposed, there may be a significant amount of specification error involved in the estimates. In light of the above weakness of structured estimates, shrinkage estimates aim to achieve an optimal balance between sample risk and specification error.

The optimal shrinkage intensity determines the optimal weight distribution between a highly structured target estimator and the sample estimates. The shrinkage intensity therefore defines the optimal trade-off between the bias and estimation error present in the target and sample estimates. The greater the shrinkage intensity, the greater the structure imposed in the shrinkage estimator. The problem essentially becomes that of shrinking an unbiased (but very variable) sample moment estimate towards a biased (but less variable) structured moment estimator in an optimal manner, in order to obtain more efficient estimates.

Martellini and Ziemann (2010) provide a modified methodology of that proposed by Ledoit and Wolf (2001) for objectively estimating the optimal shrinkage intensity (ϕ_k) for the structured target estimator M_k^{SF} and the sample estimator M_k^{sample} for the k^{th} moment tensor. The optimal shrinkage intensity can be determined by minimising a loss function that is a quadratic measure of the distance between the true (φ_k) and estimated moments (M_k^{sample}) , based on the Frobenius norm. The loss function can be expressed as:

$$L(\phi_k) = \left\| \phi_k M_k^{SF} + (1 - \phi_k) M_k^{sample} - \varphi_k \right\| \tag{4.2.6}$$

Minimising the above loss function is achieved by obtaining an appropriate asymptotic estimator of the shrinkage intensity. Since the Frobenius norm is not restricted to square matrices, Martellini and Ziemann (2010) obtain the following expression for the optimal shrinkage intensity for the higher order moment tensors:

$$\hat{\phi}_k = \frac{1}{T} \frac{\hat{\pi}_k - \hat{\rho}_k}{\hat{\gamma}_k} \tag{4.2.7}$$

Where π is the asymptotic variance of the sample estimator, γ is the specification error of the structured estimator, ρ is the asymptotic covariance between the structured and sample estimators and T is the number of time period observations[26].

As previously discussed, the Martellini and Ziemann (2010) illustrate that the shrinkage estimators extended for the higher moment tensors, using the single factor model , results in improved estimates that strongly reduce estimation error as well as improve the robustness of the overall portfolio.

[26] The calculations π, γ and especially ρ are rather complex and will not be included in this thesis. Martellini and Ziemann (2010) provide a very extensive overview of the computations.

4.3 Black-Litterman Four Moment Implied Equilibrium Returns

As previously discussed, when assets in a portfolio exhibit normally distributed returns, the standard deviation and covariance with the market returns are sufficient in estimating the assets' expected return; which is the basis on which the standard two-moment CAPM is built. However, when the return distributions exhibit significant non-normality, such as negative skewness (fat left tails) or excessive kurtosis, then the standard two-moment CAPM model will be insufficient in estimating the asset returns and risk. Hwang and Satchell (1999) state that hedge funds and emerging markets typically exhibit returns that possess significant skewness and kurtosis. They further state that models utilising higher order moments would be necessary for modelling hedge funds and emerging market returns.

Four Moment CAPM

Martellini and Ziemann (2007) propose a four moment CAPM and incorporate it with the Black-Litterman model in order to express views on hedge funds present in a portfolio. In addition to accounting for the covariance of the asset returns with the market using the portfolio beta, the four moment CAPM accounts for the co-skewness and co-kurtosis of the returns with the market portfolio. In the same way in which the standard CAPM returns were used to estimate the implied equilibrium returns (Π), the neutral starting point for the Black-Litterman model, the four moment CAPM are used. The four moment CAPM is a simple extension of the standard CAPM and can be expressed as:

$$E[R] - r_f = \alpha_1 \beta^{(2)} + \alpha_2 \beta^{(3)} + \alpha_3 \beta^{(4)} \tag{4.3.1}$$

Where $\beta^{(i)}$ are the portfolio beta vectors defined as[27]:

$$\beta^{(2)} = \frac{\Sigma w}{\mu^{(2)}(R_p)} \tag{4.3.2}$$

$$\beta^{(3)} = \frac{\Omega}{\mu^{(3)}(R_p)}$$

$$\beta^{(4)} = \frac{\Psi}{\mu^{(4)}(R_p)}$$

Where Σ, Ω and Ψ are the variance-covariance matrix, vector of co-skewness parameters and the vector of co-kurtosis parameters for the weight vector w respectively. Each portfolio beta accounts for the respective co-moments present in the data. There are numerous methodologies available for estimating the α_i values; such as estimating the partial derivatives of the expected terminal wealth with respect to the respective co-moment parameter[28] (Hwang & Satchell, 1999) or via a generalised least squares regression (Martellini &

[27] Computation for the values of Σ, Ω and Ψ would ideally be done in terms of a shrinkage estimator for the moment tensors using a single factor estimate for the structured target estimator.

[28] Hwang and Satchell (1999) suggest the following methodology to estimate the α_i parameters:

Ziemann, 2010). It is important to note that the portfolio beta vectors are functions of the weight vector w, and therefore, equation 4.3.1 provides a deterministic relation between the benchmark weights and the implied equilibrium returns.

As a result of utilising the four-moment CAPM, a fairer representation of the implied equilibrium returns can be achieved and the assumption of normality is no longer required. In addition to obtaining implied equilibrium returns that exhibit significant higher order moments, the investor is now able to express utility functions for expected returns that incorporate particular preferences for higher order moments.

Higher Moment Reverse Optimisation

As an alternative to the four moment CAPM, the implied equilibrium returns for return data containing higher moments can be reverse optimised in a similar manner in which the two moment implied equilibrium returns were obtained. Recall that implied equilibrium returns are the returns that are optimal given that the benchmark weights are optimal.

In order to obtain the implied equilibrium returns, we assume that the benchmark portfolio is in a state of equilibrium. This implies that the benchmark weights (w_b) are the optimal weights. Therefore, the aim is to find the implied returns that correspond to the benchmark weights. This process is referred to as reverse optimisation.

For the three moment case, assume the following three moment maximum utility function[29]:

$$U_3 = w\Pi_3 - \frac{\lambda_2}{2}w'M_2w + \frac{\lambda_3}{6}w'M_3(w\otimes w) \tag{4.3.3}$$

In order to maximise the above three moment utility function portfolio with regards to the portfolio weights, it is necessary to solve for the partial derivative of the above function[30] with respect to the portfolio weights, such that $\frac{\partial U}{\partial w} = 0$. Jondeau and Rockinger (2004) illustrate that the partial derivative of the third and fourth co-moments, with respect to the portfolio weights are as follows:

$$\frac{\partial}{\partial w}\left(w'M_3(w\otimes w)\right) = 3M_3(w\otimes w)$$

$$\frac{\partial}{\partial w}\left(w'M_4(w\otimes w\otimes w)\right) = 4M_3(w\otimes w\otimes w)$$

$$\alpha_1 = \frac{dE[W]}{d\sigma(W)}\sigma(r_m) \qquad \alpha_2 = \frac{dE[W]}{d\gamma(W)}\gamma(r_m) \qquad \alpha_3 = \frac{dE[W]}{d\theta(W)}\theta(r_m)$$

[29] U_3 refers to the three moment utility function and Π_3 refers to the implied equilibrium returns corresponding to the data that exhibits three moments. M_2 and M_3 are the second and third moment tensor matrices obtained in section 4.2.1 respectively. λ_2, λ_3 and λ_4 are the risk aversion parameters to variance, co-skewness and co-kurtosis respectively.

[30] In order to compute the derivative of the third moment component, $\frac{\lambda_3}{6}w'M_3(w\otimes w)$, it is necessary to utilise the chain rule, such that $\frac{\partial}{\partial w}\left(\frac{\lambda_3}{6}w'M_3(w\otimes w)\right) = \frac{\lambda_3}{6}M_3(w\otimes w) + \frac{\lambda_3}{6}w'M_3\left(\frac{\partial}{\partial w}w\otimes w\right) + \frac{\lambda_3}{6}w'M_3\left(w\otimes\frac{\partial}{\partial w}w\right)$.

Therefore, the partial derivative of the three moment utility function with respect to the portfolio weights is as follows:

$$\frac{\partial U_3}{\partial w} = \Pi_3 - \lambda_2 M_2 w + \frac{\lambda_3}{2} M_3 (w \otimes w)$$

Given the assumption of equilibrium for the benchmark portfolio, the benchmark weights (w_b) are the optimal portfolio weights which maximise the above utility function, such that:

$$\Pi_3 - \lambda_2 M_2 w_b + \frac{\lambda_3}{2} M_3 (w_b \otimes w_b) = 0$$

Therefore, the analytical solution for the three moment implied equilibrium returns (Π_3) can be expressed as:

$$\Pi_3 = \lambda_2 M_2 w - \frac{\lambda_3}{2} M_3 (w \otimes w) \tag{4.3.4}$$

The four moment case is a simple extension of the above three moment case. The four moment utility function (U_4) and its partial derivative with respect to the portfolio weights are given as follows:

$$U_4 = w\Pi_4 - \frac{\lambda_2}{2} w' M_2 w + \frac{\lambda_3}{6} w' M_3 (w \otimes w) - \frac{\lambda_4}{24} w' M_4 (w \otimes w \otimes w)$$
$$\frac{\partial U_4}{\partial w} = \Pi_4 - \lambda_2 M_2 w + \frac{\lambda_3}{2} M_3 (w \otimes w) - \frac{\lambda_4}{6} M_4 (w \otimes w \otimes w)$$

Since the benchmark weights maximise the above equations, the four moment implied equilibrium returns (Π_4) is given as follows:

$$\Pi_4 - \lambda_2 M_2 w_b + \frac{\lambda_3}{2} M_3 (w_b \otimes w_b) - \frac{\lambda_4}{6} M_4 (w_b \otimes w_b \otimes w_b) = 0$$
$$\Pi_4 = \lambda_2 M_2 w_b - \frac{\lambda_3}{2} M_3 (w_b \otimes w_b) + \frac{\lambda_4}{6} M_4 (w_b \otimes w_b \otimes w_b) \tag{4.3.5}$$

In summary of the above, the implied equilibrium returns for the two, three and four moment cases can be expressed as follows:

$$\Pi_2 = \lambda_2 M_2 w_b \tag{4.3.6}$$
$$\Pi_3 = \lambda_2 M_2 w - \frac{\lambda_3}{2} M_3 (w \otimes w)$$
$$\Pi_4 = \lambda_2 M_2 w_b - \frac{\lambda_3}{2} M_3 (w_b \otimes w_b) + \frac{\lambda_4}{6} M_4 (w_b \otimes w_b \otimes w_b)$$

As shown above, using higher moment utility functions, analytical solutions for the implied equilibrium returns that take into account the higher moments present in the data have been created. Substituting the above expressions for the standard implied equilibrium return vector, investors can safely apply the Black-Litterman model to data that is not necessarily normally distributed. However, as discussed in section 2.2.3 (The Black-Litterman Model Applied to Active Management), in order to avoid the phenomenon of information-less trades, the implied equilibrium objective function and the posterior weight vector objective functions need to be consistent[31]. As a result of incorporating higher moments into the derivation of the implied equilibrium returns, the final posterior distribution of Black-Litterman returns will therefore be non-normal. In order to solve for the posterior Black-Litterman weights, an optimiser utilising numerical methods is required to evaluate the unknown distribution. The resulting vector of Black-Litterman weights therefore incorporates a mix of the views of the investor as well as the investor's respective aversions to higher moments present in the data.

[31] For example, if a four-moment implied equilibrium return vector (Π_4) is used, a four-moment objective function needs to be used when solving for the final Black-Litterman weights, as follows:

$$\max_{w}\left(w\mathrm{E}[\mathrm{R}]-\frac{\lambda_2}{2}w'M_2w+\frac{\lambda_3}{6}w'M_3(w\otimes w)+\frac{\lambda_4}{24}w'M_4(w\otimes w\otimes w)\right)$$

Chapter 5: Portfolio and Constraint Diagnostic Tools

For any portfolio construction process to be effective, it is essential that the investor is aware of the potential risk and return effects of a particular set of constraints or views that are imposed (Schneeweis *et al.,* 2010). When implementing views on the portfolio of assets, an iterative procedure should be followed whereby the investor expresses the views and reviews the risk and return effects that arise as a direct result of such a view or set of constraints imposed on the portfolio. Upon review of the risk implications of the particular view, the views can then be adjusted and calibrated in order to enhance the alignment between the risks associated with each view and the corresponding conviction. This particular approach to analysing the effects of every user input in the portfolio inevitably leads to a more consistent implementation of the investor views and to a more balanced portfolio that better reflects the risk and return characteristics that are expected by the investor.

In addition to analysing the risk implications of each investor view expressed on the portfolio, there are many regulatory constraints that are required to be adhered to. Examples of regulatory constraints[32], in addition to tracking error and total portfolio risk constraints, include the long-only constraint, a maximum allocation to hedge funds, a minimum allocation to cash as well as various neutrality constraints[33]. Regulatory constraints are imposed on a portfolio in order to prevent the possibility of excess risk taking as well as controlling the overall portfolio risk. However, as noted by Scherer (2007), many regulatory constraints lack the insights into portfolio theory and are not sufficient for controlling risk. In fact, instead of reducing portfolio risk, poorly specified constraints may increase portfolio risk as a result of reducing the asset managers' breadth in making active decisions, thereby impairing the diversification of active bets and forcing risk-taking in positions where the manager does not exhibit skill (Scherer, 2007). It is therefore recommended that the investor investigates the relative effects of constraints on the optimal portfolio and identifies how the respective constraints are influencing the portfolio characteristics. This is in order to determine whether particular constraints are contributing or hindering the investment objectives set by the investor.

The following sections outline a set of recommended diagnostic tools that are aimed at appropriately examining the impacts of constraints and investor views on the optimal portfolio. The first section deals with describing how the effects of constraints can be mathematically quantified in terms of the effects on the optimal portfolio weights, return, volatility and utility. The second section deals with describing how constraints, as well as views, can be considered as Bayesian priors on the optimal portfolio. In addition, the effects of constraints competing against the effects of views on the optimal portfolio are described. It must

[32] Regulation 28 that gives effect to section 36(1)(BB) of the South African Pension Funds 1956, 2010.

[33] Examples of neutrality constraints are the net zero active allocation constraint whereby the active allocations sum to zero. An additional example is the beta neutrality constraint whereby the sum of the individual beta exposures sums to zero.

be noted that the diagnostic tools discussed are developed with a particular focus on Bayesian posterior returns recommended by the Black-Litterman model.

5.1 Analysing the Impacts of Constraints on the Optimal Portfolio

Asset managers primarily use investment constraints in order to impose structure on a portfolio or to reflect a particular kind of investment strategy (Schneeweis, Crowder, & Kazemi, 2010). When no information is available, portfolio managers find it useful to impose quality and risk controls on the portfolio management process in order to avoid unintended risk exposures. However, managers need to exercise caution when imposing portfolio constraints, as over-constrained portfolios may be substantially riskier out-of-sample (Fabozzi *et al.,* 2007). As described by Scherer (2007), constraints distort the valuations made by asset allocation models, thereby significantly reducing the utility of the investor.

The following sections illustrate how investment constraints provide additional information for the portfolio allocation as well as how the incorporation of this additional information inevitably distorts the return expectations generated by the valuation model. Lagrange multipliers are used to quantify the relative effects of constraints on the optimal portfolio. In particular, the effects of constraints on investor utility, information ratio shrinkage, alpha and portfolio weights are illustrated.

5.1.1 Mathematical Decomposition of Constraints

The aim of this section is to mathematically decompose the respective impacts of portfolio constraints on the optimal portfolio. In this manner, the efficiency costs associated with each constraint can be examined and reviewed.

Consider a typical mean variance utility function:

$$\max_{w}\left(w'\alpha - \frac{\lambda}{2}w'\Sigma w\right)$$

Subject to a set of linear constraints that can be written in the form:

$$Aw = c$$

Where A is a k x n matrix for *k* constraints on *n* assets, *w* is the n x 1 weight vector and c is the *k*x1 vector of equality constraints. Included in the above constraint matrix is the long-only constraint whereby $w \geq -w_B$, where w_B is the benchmark weight vector.

As illustrated by Scherer (2007), the above can be summarised as the following Lagrange function:

$$L = w'\alpha - \frac{\lambda}{2}w'\Sigma w + \gamma'(w + w_B - \delta^2) + \theta'Aw \qquad (5.1.1)$$

Where δ^2 represents a deviance variable and γ (nx1 vector) and θ (kx1 vector) represent the Lagrange multipliers for the long-only and the equality constraints. These multipliers on the long-only constraint are always positive whereas the multipliers on the equality constraints may be either positive or negative (Scherer, 2007). The Lagrange multipliers play an important role in describing the respective costs associated with the particular constraints. In particular, for binding constraints, these multipliers describe the shadow price[34] of the constraint. Following the method of Lagrangian multipliers, solving for the partial derivatives of the above equation 5.1.1 with respect to each of the variables yields final values for each of the Lagrange multipliers.

Given that the Lagrange multipliers have been solved, the optimal solution to the Lagrange function, for the optimal *constrained* weight vector can therefore be written in closed form as:

$$w^c = \frac{1}{\lambda}\Sigma^{-1}(\alpha + \gamma + A'\theta) \quad (5.1.2)$$

The constrained portfolio implied alphas can therefore be written as[35]:

$$\alpha^c = \alpha + \gamma + A'\theta$$

The above expression therefore allows one to decompose the constrained portfolio alpha into its respective constituents. As shown above, the relative impacts of the investment constraints are captured in the Lagrange multipliers $(\gamma + A'\theta)$, which are added directly to the expected return of the unconstrained solution. In other words, the Lagrange multipliers can be considered as *adjustments* to the unconstrained expected return solution. These adjustments are made in order to ensure that the optimal portfolio weights satisfy the expressed constraints.

Given that all the information regarding effects of the investment constraints are summarised by the Lagrange multipliers, we are therefore provided with the opportunity of observing the magnitude of the inefficiencies that occur as a result of the constraints by analysing how the Lagrange multipliers are utilised in the optimal portfolio. The following sections make use of the Lagrange multipliers in order to create a set of diagnostic tools that analyse the decrease in portfolio optimality as a result of portfolio constraints.

[34] The shadow price describes the gain in utility for each unit of relaxation of a particular constraint.
[35] This can be seen from substituting the closed form solution for w^c into the following equation for the implied constrained alpha: $\alpha^c = \lambda\Sigma w_\alpha^c$

5.1.2 Constraint Diagnostic Tools

The following sections describe the tools aimed at analysing the respective impacts that constraints have on the optimal portfolio allocations. The three diagnostic tools introduced are the transfer coefficient, impact on the portfolio alpha and the impact on the investor's utility. Each of the above mentioned diagnostic tools utilise the values contained in the Lagrange multipliers as a manner in which to describe the difference in optimality between the constrained and unconstrained portfolio.

The Transfer Coefficient

The information ratio is an overall measure of the active risk adjusted return of the portfolio. The transfer coefficient is a measure of the degree of information ratio shrinkage that occurs as a result of imposing constraints on the portfolio. The transfer coefficient provides a simple top line measure of the impact of constraints and how these constraints negatively impact the portfolio risk adjusted returns. Mathematically, it is the ratio of the constrained portfolio information ratio (IR^C) to the unconstrained portfolio information ratio (IR), expressed as:

$$\mathrm{T} = \frac{IR^C}{IR} \leq 1 \tag{5.1.3}$$

The above expression illustrates that an unconstrained portfolio, in sample, will always perform at least as well as a constrained portfolio. In other words, all *ex ante* portfolio constraint are value detracting. In practise, the transfer coefficient is used to measure the impact of constraints on the *ex post* information ratio (Scherer, 2007). A small value of the transfer coefficient indicates that a significant decrease in the portfolio information ratio has occurred and therefore the constraints exert a significant influence on the optimality of the portfolio.

Impact on Alpha

Using the Lagrange multiplier expression for the portfolio alpha, the alpha differential (D_α), used to illustrate the loss in alpha (excess return above the benchmark) due to constraints, can be expressed as:

$$\begin{aligned} D_\alpha &= \alpha - \alpha^c \\ &= (w - w^c)'\alpha \\ &= (\frac{1}{\lambda}\Sigma^{-1}\alpha - \frac{1}{\lambda}\Sigma^{-1}(\alpha + \gamma + A'\theta))'\alpha \\ &= -(\gamma' w + \theta' A w) \end{aligned} \tag{5.1.4}$$

The alpha differential therefore provides a simple measure of the loss in excess return resulting directly from investment constraints.

Impact on Utility

It must be noted that the ultimate objective of any analysis of the impacts of constraints on the portfolio is to quantify the loss in investor utility. The utility differential, following the methodology of Scherer (2007) is represented by the following[36]:

$$U - U^c = \frac{1}{2\lambda}\left(\gamma'\Sigma^{-1}(\gamma + A'\theta) + \theta' A\Sigma^{-1}(\gamma + A'\theta)\right) \tag{5.1.5}$$

The first term within the brackets represents the contribution of the long-only constraint and the second term represents the contribution of the equality constraints. The above equation 5.1.5 can therefore be used to analyse the change in utility as a direct result of the constraints. Using the above equation, investors can review and calibrate particular "weak" investment constraints in order to obtain marginal utility increases resulting in improvements in portfolio efficiency by loosening particular constraints that the investor deems to be of lower importance.

5.2 The Bayesian Perspective of Constraints

When constructing a portfolio, investors often express a priori that the final portfolio weights should satisfy particular constraints. Scherer (2007) states that constraints can be thought of as prior views on the final portfolio and, in particular, the long-only constraint expresses the prior view that the final portfolio weights should be positive.

Upon analysing the effects of constraints on the final portfolio utility, risk and return characteristics, it becomes evident that the final constrained portfolio is a weighted average of the unconstrained portfolio and the imbedded Lagrange multipliers. The Lagrange multipliers that arise as a result of the constraints supplies additional informational adjustments on the portfolio return, given the imposed constraints. This apparent weighting between the unconstrained portfolio return and the Lagrange multipliers strongly signifies Bayesian information shrinkage. The result of the combination of the two sources of information yields the return vector that incorporates both the unconstrained portfolio return information as well as the information pertaining to the imposed constraints.

It can be noted that when investors impose constraints on a portfolio, they are implicitly implying that they have full confidence that the final optimal portfolio should possess the desired characteristics as described by the constraints. However, when deviance variables are included in the imposed constraint (soft

[36] The equation below assumes a utility function of the following form: $\max_w \left(w'\alpha - \frac{\lambda}{2}w'\Sigma w\right)$ subject to constraints. Applying the methodology to alternative utility functions would be straightforward if the Lagrange multipliers for the constraints are achievable.

constraint expression), the investor is implicitly implying a less than 100% confidence in the given soft constraints.

In addition, the extent of the investor's confidence in the constraints can be best described by the decrease in the utility function value relative to the utility of the unconstrained portfolio. For example, if the investor is willing to sacrifice a 30% decrease in expected utility, as a result of the imposed constraints, then the value of the utility sacrificed can be considered as a proxy for the level of conviction the investor has in the imposed constraints. However, if some of the constraints are soft and if the investor is unhappy with the proposed loss in expected utility as well as the loss in information efficiency, then full or partial relaxation of the constraints is necessary.

In order to quantify the effects of constraints on the Bayesian portfolio, it is necessary to analyse the magnitude of the effects of the constraints relative to the magnitude of the effects of the prior and frequency distributions of the Bayesian portfolio. The following sections aim to quantify the magnitude of the effects of constraints on the portfolio characteristics relative to the respective influences of the prior and frequency distributed components of the Black-Litterman model.

5.2.1 The Impact of Constraints on Black-Litterman Portfolio

The following sections mathematically describe how the inclusion of investment constraints influences the characteristics of a Black-Litterman optimised portfolio. It must be noted that while the theory to be discussed relates to the Black-Litterman model, it has the flexibility to be adapted to suit almost any Bayesian portfolio optimisation model for which a variety of sources of information are being combined.

The first section deals with the respective components of the Black-Litterman model and introduces the notion that the final portfolio can be decomposed into three components, namely, a benchmark, view and constraint component. For each of these three components, a methodology analysing their respective contributions to the portfolio return, weights, variance and beta coefficient is given. It must be noted that when constraints are implemented, the constraint component of the portfolio consumes a portion of the portfolio influenced by the investor views and the benchmark. Naturally, a favourable constrained portfolio is one for which the constraint component exerts the smallest possible influence on the final portfolio. This is due to the fact that the investor desires a portfolio that not only satisfies the investment constraints, but fully incorporates all investment information (benchmark returns and views) in the final optimal portfolio.

The second section proceeds an additional step further and analyses the respective contributions of each of the views to the final portfolio. Strong attention is drawn to how the investment constraints distort the original expression of the views. Views that have been excessively distorted by the constraints might as well have not been expressed at all since their original expression forms an insignificant portion of the optimal portfolio. In addition, in order to determine whether the expressed views adequately satisfy the risk objectives of the investor, the percentage contribution of each view to the portfolio tracking error is calculated. For the constrained case, the portion of the expressed view influenced by constraints is also examined for its contribution to tracking error and the active beta. In addition, in order to determine how each of the asset's active weight in the final optimal portfolio is influenced, the HHI[37] measure for concentration is used. In other words, the HHI measure is used to calculate which active weights are influenced by many views or are dominated by a single view.

The final section extends the methodology for the inclusion of higher moments in the portfolio optimisation. Section 5.3 deals with the issue of marginally relaxing the long-only constraint in small increments and observing at each increment how each of the above mentioned components are influenced. In other words, the change in the proportions of the three portfolio components as well as the level of view distortion is observed when the long-only constraint is marginally relaxed.

5.2.2 Black-Litterman Portfolio Analysis

It must be noted that constraints imposed on a portfolio can be considered of as prior views on the portfolio final composition. For example, the long-only constraint expresses the prior view that portfolio allocations must be positive (Scherer, 2007). As previously illustrated, constraints implicitly change the forecasts made by the valuation models and therefore create inefficiencies in expected risk and return forecasts (Da Silva et al., 2009). Within the Black-Litterman framework, the Bayesian approach is used to impose views on the implied equilibrium returns. These views are intuitively expressed as active weight positions imposed on the portfolio. However, in the presence of constraints, the impact of these views is somewhat limited and distorted due to the optimiser only identifying and considering solutions that satisfy the constraints. This therefore raises the question of whether the final recommended constrained portfolio and its respective risk and return characteristics, is driven by the investor views *or* if these views are drowned-out and dominated by the constraints.

The following methodology aims to decompose the recommended portfolio into its respective components in order to identify the proportions of the characteristics of the portfolio that are attributable to the views as well as the proportions that are attributable to the constraints.

[37] The HHI measure is a value that lies between the range of zero and one. A value close to one indicates a high concentration, whereby the object measured is influenced by a few sources. Conversely, a value close to zero indicates that there are many sources that each exerts a small amount of influence on the object measured. Traditionally, the HHI measure is used to analyse the concentration of the weights of a portfolio of assets. In this manner, the extent to which an asset's active weight is influenced by all the views, or a small number of views, can be observed.

In order to measure the impact of the investor views in the presence of constraints, the final portfolio will be analysed according to its three respective components:

- Benchmark component
- View component
- Constraint component

All three of the above components will be analysed according to their return, portfolio weight allocations and risk contributions within the mean-variance setting. In addition, the respective portfolio tracking error compositions will be analysed.

Expected Return

In order to appropriately decompose the Black-Litterman return vector into its respective components, it must be noted that Da Silva *et al.,* (2009) provide an alternative expression for the Black-Litterman expected return vector which separates Black-Litterman return vector into the implied equilibrium return and view returns components as follows:

$$\begin{aligned} E[R_{BL}] &= \Pi + \Sigma P' \left[\frac{\Omega}{\tau} + P\Sigma P'\right]^{-1} (Q - P\Pi) \\ &= \Pi + V \end{aligned} \tag{5.2.1}$$

Substituting the above Black-Litterman return decomposition into the constrained expected return vector and using Lagrange multipliers, the final *constrained* portfolio return vector can be expressed as follows:

$$E[R_{BL}^{con}] = \Pi + V + \gamma + \theta' A \tag{5.2.2}$$

The above expression for the constrained portfolio return illustrates the three components, namely: benchmark returns(Π), active view returns (V) and the Lagrange multiplier returns adjustment components, $(\gamma + \theta' A)$[38], due to the constraints.

[38] It is important to note that the benchmark implied equilibrium returns are independently estimated using the CAPM. Conversely, the expected return attributable to the views (V) are dependent on the difference between the view vector and the implied equilibrium returns $(Q - P\Pi)$. The Lagrange multipliers that arise as a result of the constraints are also dependent on the views and implied equilibrium returns. It is therefore easy to see that each of the components are not fully independent and that a change in Π will undoubtedly result in a change V and $(\gamma + \theta' A)$. However, these three components are still able to be expressed as separate components provided that they are analysed in the static portfolio case with all the respective variables remaining constant.

Portfolio Weights

Substituting the above expression 5.2.2 for the constrained Black-Litterman return vector into the analytical expression for the constrained portfolio weights, the resulting expression for the optimal constrained mean-variance portfolio weights can be expressed as follows:

$$\begin{aligned} w^c &= \frac{1}{\lambda}\Sigma^{-1}(\Pi + V + \gamma + \theta' A) \\ &= \frac{1}{\lambda}\Sigma^{-1}(\Pi) + \frac{1}{\lambda}\Sigma^{-1}(\mathrm{V}) + \frac{1}{\lambda}\Sigma^{-1}(\gamma + \theta' A) \\ &= w_{\Pi} + w_{\mathrm{V}} + w_{\mathrm{con}} \end{aligned} \tag{5.2.3}$$

The above expression splits the final optimal constrained mean-variance portfolio weight allocations into the three components;

- The weight attributable to the benchmark/implied equilibrium returns (w_{Π}) in absence of the unique investor views.
- The active weight allocations (w_{V}) arriving as a result of the unique investor views. The optimal portfolio in light of the investor views and in the *absence* of investment constraints is simply: $w_{\Pi} + w_{\mathrm{V}}$.
- The weight adjustments (w_{con}) resulting from the Lagrange multipliers in order to adjust the final portfolio weights so that the respective constraints may be satisfied.

In order to determine if the constraints dampen, distort or amplify the expressed views, it becomes necessary to observe how the addition of w_{con} influences the overall portfolio weight composition as well as observe how the active weights resulting from the views is being altered as a result of the implementation of constraints.

The relative proportional contributions of the three components of the overall constrained portfolio weight allocations can be obtained by calculating the respective norms of each of the components and dividing them by the sum of all the component norms. The proportional contribution of each component to the overall constrained portfolio weight can be written as follows:

$$\begin{aligned} \mathit{Benchmark\ Proportion} &= \frac{|w_{\Pi}|}{|w_{\Pi}| + |w_{\mathrm{V}}| + |w_{\mathrm{con}}|} \\ \mathit{View\ Proportion} &= \frac{|w_{\mathrm{V}}|}{|w_{\Pi}| + |w_{\mathrm{V}}| + |w_{\mathrm{con}}|} \\ \mathit{Constraint\ Proportion} &= \frac{|w_{\mathrm{con}}|}{|w_{\Pi}| + |w_{\mathrm{V}}| + |w_{\mathrm{con}}|} \end{aligned} \tag{5.2.4}$$

Typically, within the Black-Litterman framework, the relative proportions between the benchmark portfolio and the view portfolio remain constant regardless of the expressed constraints. When the magnitude of the weight vector corresponding to the investment constraints increases or decreases, depending on how binding the constraints are, the cost associated with an increase in the constraint portion results in an equally distributed decrease in the benchmark and view portions influencing the portfolio.

Impact on Exposure to the Benchmark

The exposure of a portfolio to a particular benchmark is measured by the beta coefficient. In order to analyse the effects that the constraints and the view active weights have on the overall portfolio, it is straightforward procedure to decompose the portfolio beta into its respective components. The portfolio beta can be decomposed as follows:

$$\begin{aligned}\beta &= \frac{w^{c\prime}\Sigma w_{\Pi}}{w_{\Pi}{}'\Sigma w_{\Pi}} \\ &= \frac{(w_{\Pi} + w_{\mathrm{V}} + w_{\mathrm{con}})'\Sigma w_{\Pi}}{w_{\Pi}{}'\Sigma w_{\Pi}} \\ &= \frac{w_{\Pi}{}'\Sigma w_{\Pi}}{w_{\Pi}{}'\Sigma w_{\Pi}} + \frac{w_{\mathrm{V}}{}'\Sigma w_{\Pi}}{w_{\Pi}{}'\Sigma w_{\Pi}} + \frac{w_{\mathrm{con}}{}'\Sigma w_{\Pi}}{w_{\Pi}{}'\Sigma w_{\Pi}} \\ &= 1 + \beta_V + \beta_{con}\end{aligned} \tag{5.2.5}$$

Where β_V and β_{con} are the view and constraint component betas that contribute to the overall portfolio beta. Negative values of either component indicate decreasing exposure to the benchmark, whereas positive values elevate the portfolio's sensitivity to variations in the benchmark. In analysing these respective beta components, it allows for a more transparent manner in which the investor can be fully aware of the impacts that the views and constraints have on the overall portfolio.

Portfolio Variance

Given the decomposition of the portfolio weights, the overall portfolio variance can be decomposed into its respective components. Each component can be analysed and assessed for its contribution to the overall portfolio variance. The portfolio variance can be expressed as follows:

$$\begin{aligned}\sigma_p^2 &= w'\Sigma w \\ &= (w_{\Pi} + w_{\mathrm{V}} + w_{\mathrm{con}})'\Sigma(w_{\Pi} + w_{\mathrm{V}} + w_{\mathrm{con}}) \\ &= \sigma_{\Pi}^2 + \sigma_V^2 + \sigma_{con}^2 + 2w_{\mathrm{V}}'\Sigma w_{\Pi} + 2w_{\mathrm{con}}'\Sigma w_{\Pi} + 2w_{\mathrm{con}}'\Sigma w_{\mathrm{V}}\end{aligned} \tag{5.2.6}$$

It is evident that the inclusion of portfolio constraints (as opposed to an unconstrained portfolio) is composed of three components, namely the constraint portfolio variance and the covariance of the constraint portfolio with the view and benchmark portfolios. The overall contribution to portfolio variance, as well as the overall proportion of the total variance that constraints contribute towards is therefore summarised by the following:

$$Constraint\ Contribution\ to\ Variance = \sigma_{con}^2 + 2w'_{\text{con}}\Sigma w_{\Pi} + 2w'_{\text{con}}\Sigma w_{\text{V}}$$

and

$$Constraint\ Variance\ Portion = \frac{\sigma_{con}^2 + 2w'_{\text{con}}\Sigma w_{\Pi} + 2w'_{\text{con}}\Sigma w_{\text{V}}}{\sigma_p^2} \quad (5.2.7)$$

It can be noted that if constraints reduce the portfolio risk, then the contribution of the constraints to the portfolio variance will be negative. However, if the contribution to the portfolio variance is positive, then the implementation of constraints has in fact *increased* the portfolio risk instead of reducing it as intended.

Tracking Error Variance Decomposition

Following the same methodology as above, the active risk component, namely the tracking error variance, is represented as follows:

$$\begin{aligned} TE^2 &= (w_{\text{V}} + w_{\text{con}})'\Sigma(w_{\text{V}} + w_{\text{con}}) \\ &= \sigma_V^2 + \sigma_{con}^2 + 2w'_{\text{con}}\Sigma w_{\text{V}} \end{aligned} \quad (5.2.8)$$

It is straight forward to see that the inclusion of constraints on a portfolio adds the following to portfolio tracking error relative to the unconstrained portfolio case: $\sigma_{con}^2 + 2w'_{\text{con}}\Sigma w_{\text{V}}$. Therefore, the proportion of the tracking error variance attributable to imposed constraints can be expressed as:

$$Constraint\ TE\ Proportion = \frac{\sigma_{con}^2 + 2w'_{\text{con}}\Sigma w_{\text{V}}}{TE^2} \quad (5.2.9)$$

It must be noted that when an investor expresses views that differ from that implied by the benchmark, active risk is incurred. The investor therefore aims to take active risks in positions that exhibit the greatest convictions. Given the imposed constraints, any portion of the active risk budget (tracking error variance) occupied by the effects of constraints[39], is a strong indication of the level of conviction that the investor has in the imposed constraints. If the investor finds that the imposed constraints occupy too great a portion of the overall portfolio tracking error, then a marginal relaxation of the constraints is recommended to the extent that the marginal risk due to constraints is appropriate given the level of conviction the investor has in the imposed constraints.

[39] Assuming that the contribution of the constraints to the portfolio tracking error variance is positive, $\sigma_{con}^2 + 2w'_{\text{con}}\Sigma w_{\text{V}} > 0$, whereby the inclusion of constraints effectively increases tracking error rather than reducing it.

5.2.3 The Impacts of Constraints on the Views of the Black-Litterman Model

In practise, investors know qualitatively how extreme their views are with respect to the market portfolio. However, it would be helpful for the manager to posses various quantitative measures that indicate which views are the most extreme relative to the market portfolio and also to provide a basis for which the investor is able to re-evaluate the expressed views and implement them in a consistent manner.

In the previous section, the overall contribution of the investor views and constraints to the portfolio weight, variance, return and beta were identified. The purpose of this section is to take one step further and analyse the individual views that make up the view vector as well as describe in detail how the imposed constraints influence or distort the views of the investor. The following sections demonstrate the respective tools that aid the investor in determining the level of impact each particular view will have on the final portfolio characteristics.

Distortion of the Views

When constraints are expressed, a set of Lagrange multipliers are added to the expected Black-Litterman return vector, in order to adjust the expected return so that an optimal portfolio that satisfies the constraints can be obtained. Therefore, the final constrained Black-Litterman return vector ($E[R_{BL}^{con}]$) can be expressed as the addition of the *unconstrained* Black-Litterman return vector ($E[R_{BL}]$) and the return adjustment due to the Lagrange multipliers ($R_{con} = \gamma + \theta' A$) as follows:

$$E[R_{BL}^{con}] = E[R_{BL}] + R_{con} \quad (5.2.10)$$

Assuming that the final return vector, given the views *and* constraints, is given by$E[R_{BL}^{con}]$, it is of interest to determine what particular views, Q_{con}, lead to the constrained return vector[40]. It can be noted that the value of Q_{con} contains both the explicit views expressed by the investor *and* the implicit views implied by the constraints[41]. Therefore, any difference between Q_{con} and the unconstrained investor view vector Q, is a measure of the distortion of the original investor views that occurs as a result of constraints imposed on the portfolio.

Solving for Q_{con} involves "backing-out" the constrained views (Q_{con}) implied by the *constrained* Black-Litterman return vector. Assuming that the Black-Litterman view uncertainty matrix (Ω) remains unchanged, the constrained Black-Litterman return equation can therefore be expressed as follows:

$$E[R_{BL}^{con}] = [(\tau\Sigma)^{-1} + P^T\Omega^{-1}P]^{-1}[(\tau\Sigma)^{-1}\Pi + P^T\Omega^{-1}Q_{con}]$$

[40] The assumption is that the convictions of the respective views, contained in Ω, remain constant and that all the effects of constraints on the constrained expected Black-Litterman return vector translate into the value of Q_{con}.

[41] The value of Q_{con} can be interpreted as the investor views ***given*** the portfolio constraints. It is an overall measure of the views of the investor, composed of the explicit views and the views implied by the constraints. The effects of the constraints on the views are captured by $(\delta + I)R_{con}$ in equation 5.2.13.

Assuming that the views expressed are *absolute* views on the expected return vector, the view identification matrix P becomes the identity matrix. This simplification yields the following equation:

$$E[R_{BL}^{con}] = [(\tau\Sigma)^{-1} + \Omega^{-1}]^{-1}[(\tau\Sigma)^{-1}\Pi + \Omega^{-1}Q_{con}] \quad (5.2.11)$$

Solving for Q_{con} in the above equation yields the following:

$$Q_{con} = \Omega\{[(\tau\Sigma)^{-1} + \Omega^{-1}]E[R_{BL}^{con}] - (\tau\Sigma)^{-1}\Pi\}$$

Let $\delta = \Omega(\tau\Sigma)^{-1}$ and simplifying the above, the equation becomes:

$$Q_{con} = \delta E[R_{BL}^{con}] + E[R_{BL}^{con}] - \delta\Pi$$

Substitute $E[R_{BL}^{con}] = E[R_{BL}] + R_{con}$ as in equation 5.2.10, gives:

$$Q_{con} = \delta E[R_{BL}] + E[R_{BL}] + \delta R_{con} + R_{con} - \delta\Pi$$
$$Q_{con} = (\delta + I)(E[R_{BL}] + R_{con} - \Pi) + \Pi$$
$$Q_{con} = (\delta + I)(E[R_{BL}] - \Pi) + (\delta + I)R_{con} + \Pi$$

Similarly, it can be noted that the expression for the unconstrained view vector can be stated as:

$$Q = (\delta + I)(E[R_{BL}] - \Pi) + \Pi \quad (5.2.12)$$

Therefore, the final constrained view vector can be expressed as the addition of the unconstrained view vector and the view adjustment due to constraints as follows:

$$Q_{con} = Q + (\delta + I)R_{con} \quad (5.2.13)$$

Expression 5.2.13 above gives the value of the view vector for views that are implied by the resulting constrained expected return. In other words, taking into account the relative distortion of the Black-Litterman return vector as a result of the imposed constraints, the value of the vector Q_{con} gives the implied views corresponding to the constrained solution.

As previously stated, the magnitude of these constrained views differs from the values of the originally expressed views, the difference for which is a measure of the level of view distortion that has occurred as a result of the imposed constraints. In order to appropriately measure the level of view distortion that has occurred, the percentage change between the constrained views and the original views gives an indication of whether a particular view has been suppressed or amplified[42]. The view distortion percentage can be written as:

[42] The terms "suppressed" and "reduced" as well as "expanded" and "amplified" will be used interchangeably throughout the thesis. A suppressed (or reduced) view distortion refers to a view that has been made smaller in magnitude as a result of the constraints. In contrast, an amplified (or expanded) view distortion refers to a view that has been made greater in magnitude as a result of the imposed constraints.

$$K = \frac{Q_{con} - Q}{Q} \tag{5.2.14}$$

Where positive percentage values of the vector K represent view expansion. Negative percentage values between 0% and -100% represent view suppression, where values near 0% indicate very little view distortion. Percentage values less for $K < -100\%$ indicate that the view has changed sign and is of a magnitude greater than the original view, thereby indicating overwhelming view distortion. Large percentage distortion values imply that the constraints have a significantly impact on the optimal portfolio and further illustrate that the original investor views have been effectively “drowned-out” by the constraints. This occurrence is of great concern to the investor as the originally expressed views are no longer exerting a significant impact on the final optimal portfolio, despite having spent valuable institutional resources on analysing markets in order to attain these particular views. In addition to wasting valuable resources on forming the views that inevitably becoming severely distorted, the constraints force the active allocations into positions that would contradict the views and convictions of the investor, thereby resulting in an increase in the risk of the portfolio.

Impact of Views on the Expected Return and Weight Vectors

For this particular section and the sections to follow, a distinction is made between constraint-implied views and the originally expressed unconstrained views. For simplicity, each of the diagnostic tools described below are developed for the unconstrained portfolio case and adaptation for the constrained portfolio case is a straightforward case of substituting Q_{con} and $E[R_{BL}^{con}]$ for Q and $E[R_{BL}]$ respectively.

Since the values of Q are interpreted as the views of the expected returns of the assets in the portfolio, it follows that these views are composed of the implied equilibrium returns (Π) as well as the *raw views* v_i for which the investor view differs from the implied equilibrium return. In particular, the view vector Q can be expressed as $Q = \Pi + \sum_i^k v_i$ where the raw views are vectors that contain single entries in the positions pertaining to the particular view with zeros elsewhere. For example, the *k* raw views can be expressed as:

$$v_1 = \begin{bmatrix} v_1 \\ 0 \\ \vdots \\ 0 \end{bmatrix};\ \dots;\ v_k = \begin{bmatrix} 0 \\ 0 \\ \vdots \\ v_k \end{bmatrix}$$

Using the expression for the decomposition of the view vector in equation 5.2.12 and solving for $E[R_{BL}]$, the following is obtained:

$$\begin{aligned}
E[R_{BL}] &= (\delta + I)^{-1}(Q + \delta\Pi) \\
&= (\delta + I)^{-1}\delta\Pi + (\delta + I)^{-1}Q \\
&= (\delta + I)^{-1}\delta\Pi + (\delta + I)^{-1}(\Pi + \sum_i^k v_i) \\
&= (\delta + I)^{-1}\delta\Pi + (\delta + I)^{-1}\Pi + (\delta + I)^{-1}\sum_i^k v_i \\
&= (\delta + I)^{-1}\left[\delta\Pi + \Pi + \sum_i^k v_i\right] \\
&= (\delta + I)^{-1}\left[(\delta + I)\Pi + \sum_i^k v_i\right] \\
&= \Pi + (\delta + I)^{-1}\sum_i^k v_i
\end{aligned} \tag{5.2.15}$$

Therefore, as illustrated above, the contribution of each investor view to the final unconstrained Black-Litterman return is given by[43]:

$$E[R_{v_i}] = (\delta + I)^{-1}v_i \tag{5.2.16}$$

Since the contribution of each investor view to the final Black-Litterman return can be separately identified by equation 5.2.16, the corresponding two moment optimised portfolio weights corresponding to each investor view, can also be identified by substituting equation 5.2.16 into equation 2.2.3 as follows:

$$w_{v_i} = \frac{1}{\lambda}\Sigma^{-1}(\delta + I)^{-1}v_i \tag{5.2.17}$$

[43] The contribution of each distorted constrained view, Q_{con}, is given by $(\delta + I)^{-1}v_i^{con}$ where v_i^{con} is the value of the raw constrained view for asset *i* given by $Q_{con} = \Pi + \sum_i^k v_i^{con}$

It can also be noted that the active weight vector is the sum of the individual raw view weight vectors[44], where:

$$w_V = \sum_i^k w_{v_i}$$
$$= \frac{1}{\lambda}\Sigma^{-1}(\delta + I)^{-1}\sum_i^k v_i \qquad (5.2.18)$$

As a result of being able to separately identify the weight vectors attributable to each view, it is now possible to calculate the extent of which each view contributes to the portfolio tracking error and active beta coefficient.

Impact of the Views on the Portfolio Tracking Error

It would be useful for an investor to be able to identify which particular views contribute the most to the overall portfolio tracking error. Tracking error is a measure of active risk and can be defined as the variance of the portfolio alpha. Alternatively, tracking error can be thought of as a measure of how closely the portfolio tracks the benchmark portfolio. A small tracking error implies a low level of deviance from the benchmark, while a high tracking error implies a high level of deviance from the benchmark portfolio and therefore a high level of active risk.

The following sections analyse the components of the portfolio tracking error in order to identify the most active positions in the portfolio. Identifying the most active positions provides the asset manager with the necessary insight in order to appropriately calibrate the views in order to result in views that are consistent with their associated view confidence levels (Scherer, 2007). Views expressed with high levels of confidence should be allocated appropriately high active risk budgets and views with low levels of confidence should be allocated with lower risk budgets. In this section, the derivation of the ***percentage contibution of the views to tracking error (PCVTE)*** will be illustrated and discussed. The PCVTE is a useful active risk diagnostic tool which is used to identify the views that are the most extreme in terms of contributing to the overall portfolio active risk. Asset managers therefore have the opportunity to examine and revisit any active decision in light of the active risk impacts of each particular view. Since the risk impacts of each particular view can be identified, managers responsible for those particular views can also be held accountable for the performance of the portfolio whereby their incentive fees can be matched to the contribution of their respective views to tracking error (Schneeweis *et al.,* 2010).

[44] It can be noted that the sum of each of the weights attributable to the constrained raw views, v_i^{con}, results in the following equality: $\sum_i^k w_{v_i^{con}} = w_V + w_{con}$

Recall that the unconstrained portfolio tracking error squared is calculated via the following formula:

$$TE^2 = w_{\mathrm{V}}{}'\Sigma w_{\mathrm{V}}$$

In order to identify the contribution of each view to the portfolio tracking error, it is necessary to expand the overall view weight vector, w_{V}, into its respective components attributable to each view. This can be accomplished by creating a $k \times k$ matrix, W_q, composed of column vectors equal to the weight vectors attributable to each view, w_{v_i} obtained from equation 5.2.17. The summation along the horizontal elements of W_q yields the overall view weight vector. Therefore, the following formula for calculating the ***contribution of each view to the portfolio tracking error*** squared can be expressed as:

$$TEQ^2 = W_q{}'\Sigma w_{\mathrm{V}} \tag{5.2.19}$$

The above result is a vector of tracking error variances that add up to the total portfolio tracking error, with each element equal to the contribution of the expressed view to the overall tracking error variance. Dividing each element in the above vector by the overall portfolio tracking error variance returns a vector of relative percentage contributions of the views to the tracking error[45]:

$$PCVTE = \frac{W_q{}'\Sigma w_{\mathrm{V}}}{w_{\mathrm{V}}{}'\Sigma w_{\mathrm{V}}} \tag{5.2.20}$$

Analysing the Contribution of the Views to the Portfolio Beta

Recall that the unconstrained portfolio *active* beta coefficient[46] is calculated via the following formula:

$$\beta_V = \frac{w_{\mathrm{V}}{}'\Sigma w_{\Pi}}{w_{\Pi}{}'\Sigma w_{\Pi}}$$

As previously illustrated, expanding the vector w_{V} by including W_q yields the following formula for calculating the ***contribution of each view to the active portfolio beta***:

$$\beta Q_V = \frac{W_q{}'\Sigma w_{\Pi}}{w_{\Pi}{}'\Sigma w_{\Pi}} \tag{5.2.21}$$

---

[45] For the constrained portfolio case, the columns of W_q^{con} contain the constrained view weight vectors ($w_{v_i^{con}}$) and the overall portfolio active weight is composed of w_{V} *and* w_{con}. Therefore, the constrained PCVTE is written as: $PCVTE^{con} = \frac{W_q^{con\prime}\Sigma(w_{\mathrm{V}}+w_{\mathrm{con}})}{(w_{\mathrm{V}}+w_{\mathrm{con}})'\Sigma(w_{\mathrm{V}}+w_{\mathrm{con}})}$

[46] Recall that the overall portfolio beta is composed of the following: $\beta = 1 + \beta_V + \beta_{con}$, where β_V is the beta component that differs from the benchmark as a result of the views. If constraints are introduced, the active component is composed of $\beta_V + \beta_{con}$

The result is a vector of beta coefficients attributable to each view, the sum of which equals the value of β_V. Dividing each of the elements by the overall active beta coefficient yields the *percentage contribution of each view to the active beta coefficient*[47]:

$$PCV\beta_V = \frac{\frac{W_q{}'\Sigma w_\Pi}{w_\Pi{}'\Sigma w_\Pi}}{\beta_V} \tag{5.2.22}$$

The above result is useful in that it provides a vector of percentage contributions of the views to the sensitivity of the portfolio to respond to changes in the market portfolio. The percentage values may be either positive or negative depending on whether the view increases or reduces the portfolio Beta. It can be noted that views which provide negative contributions to the portfolio beta strongly contradict that of the market. The investor is therefore able to identify the views that impact the portfolio favourably or unfavourably and has the opportunity to calibrate the views in order to obtain a portfolio exposure that is more suitable according to the objectives of the portfolio.

Concentration of the Views in the Black-Litterman Active Weights

As previously discussed, the active asset weights are influenced by the expressed views. Due to the fact that particular assets are correlated with each other, a view on the return for one particular asset will indirectly affect all the assets that are highly correlated with it. Therefore, the extent to which a particular active asset weighting is influenced by a single view or by a number of views is uncertain. If an asset weight is highly concentrated, then it is dominated by a few views. In contrast, if an asset active weighting is not concentrated, then it is influenced by multiple views and no single view exerts a significant dominance.

In order to measure the level of concentration of the views in each of the active portfolio weights, the HHI index measure is used. Note that the final active weight due to the Black-Litterman views can be expanded and represented as the addition of the individual *view* weight vectors as follows:

$$w_V = \sum_{i}^{k} w_{v_i}$$

$$\begin{bmatrix} w_{V_1} \\ \vdots \\ w_{V_k} \end{bmatrix} = \begin{bmatrix} w_{v_{11}} \\ \vdots \\ w_{v_{1k}} \end{bmatrix} + \cdots + \begin{bmatrix} w_{v_{k1}} \\ \vdots \\ w_{v_{kk}} \end{bmatrix}$$

---

[47] For the constrained portfolio case with views retrieved from Q_{con}, the constrained version can be written as:

$$PCV\beta_V^{con} = \frac{\frac{W_q^{con\prime}\Sigma w_\Pi}{w_\Pi{}'\Sigma w_\Pi}}{\beta_V + \beta_{con}}$$

Applying the HHI measure of concentration to the portfolio active weights at a view level yields the following vector of HHI measures:

$$\begin{bmatrix} HHI_{w_{V_1}} \\ \vdots \\ HHI_{w_{V_k}} \end{bmatrix} = \begin{bmatrix} {w_{v_{11}}}^2 \\ \vdots \\ {w_{v_{1k}}}^2 \end{bmatrix} + \cdots + \begin{bmatrix} {w_{v_{k1}}}^2 \\ \vdots \\ {w_{v_{kk}}}^2 \end{bmatrix}$$

$$= \begin{bmatrix} \sum_{i=1}^{k} {w_{v_{11}}}^2 \\ \vdots \\ \sum_{i=1}^{k} {w_{v_{ik}}}^2 \end{bmatrix} \tag{5.2.23}$$

Given that the basic assumption behind deriving the individual view weight vectors is that the views are restricted to be absolute views[48] only, it is obvious that each active weight will be dominated by the particular view for which it is directly associated. Any additional influence exerted by other indirect views is as a result of the correlation structure of the particular asset with the other assets in the portfolio. In other words, one can expect that assets which are largely independent of other assets in the portfolio will be highly concentrated. In contrast, assets that are highly correlated with other assets are expected to exhibit lower levels of view concentration.

5.2.4 Contribution of the Portfolio Components to the Higher Moments

The inclusion of higher moments in the investor utility function significantly complicates the view and constraint diagnostic tools above. In particular, obtaining an analytical expression for the portfolio weight vector decomposition is not straightforward. However, the use of numerical methods significantly simplifies the problem. Given that we know the values of the vectors Π and V, obtaining the individual weight vectors w_{Π}, w_{V} and w_{con} are obtained via separate higher moment utility optimisations.

Once the individual weight vector components corresponding to the higher moment utility function have been obtained, they can simply be substituted into the above equations (for tracking error, beta, etc), as per the mean-variance case.

Given a higher moment utility function, whereby the investor typically expresses preference for positive portfolio co-skewness and lower co-kurtosis, analysing the various sources of information that contribute to the overall portfolio higher moments becomes practically important. As a result of being able to identify the factors that contribute most significantly to the optimal portfolio higher moments, those particular factors can be re-assessed and calibrated in order to obtain a more favourable portfolio distribution. In

[48] Recall that within the Black-Litterman framework, views can be either absolute or relative. Absolute views are views on single assets in isolation, for example "Asset A is expected to have a return of *x*%". In contrast, relative views compare assets with each other, such as "Asset A will outperform Asset B by y%". The assumption of absolute views only ensures that the view participation matrix P is the identity matrix.

particular, the relative effects of the relaxation of particular constraints can be observed and it can be determined whether certain constraints contribute favourably to the portfolio higher moments.

The following sections outline how the three components (benchmark, views and constraints) contribute to the portfolio higher moments.

Portfolio Co-Skewness

In order to estimate the marginal effects on the portfolio co-skewness as a result of including constraints and investor views, it is necessary to estimate the portfolio co-skewness coefficients in a stepwise manner and analyse the resulting marginal changes that occur when views and secondly, the constraints are added to the portfolio optimisation.

The benchmark portfolio co-skewness (third moment) can be expressed as:

$$S_{\Pi} = w_{\Pi}{}' M_3 (w_{\Pi} \otimes w_{\Pi})$$

The marginal co-skewness[49] that arises as a result of the inclusion of views as well as the inclusion of constraints *and* views on the optimal portfolio can be expressed as:

$$Marginal\ S_{\mathrm{V}} = (w_{\Pi} + w_{\mathrm{V}})' M_3 ((w_{\Pi} + w_{\mathrm{V}}) \otimes (w_{\Pi} + w_{\mathrm{V}})) - w_{\Pi}{}' M_3 (w_{\Pi} \otimes w_{\Pi}) \qquad (5.2.24)$$

$$\begin{aligned} Marginal\ S_{\mathrm{con}} = {} & (w_{\Pi} + w_{\mathrm{V}} + w_{\mathrm{con}})' M_3 ((w_{\Pi} + w_{\mathrm{V}} + w_{\mathrm{con}}) \otimes (w_{\Pi} + w_{\mathrm{V}} + w_{\mathrm{con}})) \\ & - (w_{\Pi} + w_{\mathrm{V}})' M_3 ((w_{\Pi} + w_{\mathrm{V}}) \otimes (w_{\Pi} + w_{\mathrm{V}})) \end{aligned} \qquad (5.2.25)$$

The proportion of the overall contribution to portfolio co-skewness for each component can be expressed as[50]:

$$\begin{aligned} Portfolio\ CoSkewness &= \frac{|w_{\Pi}{}' M_3 (w_{\Pi} \otimes w_{\Pi})| + |Marginal\ S_{\mathrm{V}}| + |Marginal\ S_{\mathrm{con}}|}{(w_{\Pi} + w_{\mathrm{V}} + w_{\mathrm{con}})' M_3 ((w_{\Pi} + w_{\mathrm{V}} + w_{\mathrm{con}}) \otimes (w_{\Pi} + w_{\mathrm{V}} + w_{\mathrm{con}}))} \\ &= \frac{|w_{\Pi}{}' M_3 (w_{\Pi} \otimes w_{\Pi})|}{w' M_3 (w \otimes w)} + \frac{|Marginal\ S_{\mathrm{V}}|}{w' M_3 (w \otimes w)} + \frac{|Marginal\ S_{\mathrm{con}}|}{w' M_3 (w \otimes w)} \end{aligned} \qquad (5.2.26)$$

If any particular component contributes negatively to the investor utility, the investor may consider adjusting the magnitude of which each of the individual components influences the portfolio. For example, if the view component has a significant negative contribution to the portfolio co-skewness, the relative weighting between the view portfolio and benchmark portfolio can then be adjusted (via the calibration of

[49] The marginal skewness of the views was calculated by subtracting the skewness of the benchmark portfolio from the overall (view and benchmark) unconstrained portfolio skewness. Positive marginal contributions to the portfolio co-skewness are favourable as it increases the probability of positive returns. Conversely, negative marginal contributions to co-skewness are unfavourable as the probably of negative returns are increased.

[50] The numerators are composed of the respective norms of the individual marginal contributions of the views, benchmark and constraints.

τ)[51]. A mild relaxation of the portfolio constraints may also result in favourable portfolio co-skewness movements.

Portfolio Co-Kurtosis

In the same manner in which the portfolio co-skewness can be analysed, the portfolio co-kurtosis can be analysed according to its respective contributing components. The relative effects of three individual components on the overall portfolio co-kurtosis can be identified numerically via separate optimisations.

By running separate optimisations, the benchmark implied equilibrium return co-kurtosis can be expressed as follows:

$$K_{\Pi} = w_{\Pi}{}' M_4 (w_{\Pi} \otimes w_{\Pi} \otimes w_{\Pi})$$

The marginal contribution of the views and constraints can be expressed as follows[52]:

$$\begin{aligned} Marginal\ K_{\mathrm{V}} &= (w_{\Pi} + w_{\mathrm{V}})' M_4((w_{\Pi} + w_{\mathrm{V}}) \otimes (w_{\Pi} + w_{\mathrm{V}}) \otimes (w_{\Pi} + w_{\mathrm{V}})) \\ &\quad - w_{\Pi}{}' M_4(w_{\Pi} \otimes w_{\Pi} \otimes w_{\Pi}) \end{aligned} \tag{5.2.27}$$

$$\begin{aligned} Marginal\ K_{\mathrm{con}} &= (w_{\Pi} + w_{\mathrm{V}} + w_{\mathrm{con}})' M_4((w_{\Pi} + w_{\mathrm{V}} + w_{\mathrm{con}}) \otimes (w_{\Pi} + w_{\mathrm{V}} + w_{\mathrm{con}}) \otimes (w_{\Pi} + w_{\mathrm{V}} \\ &\quad + w_{\mathrm{con}})) - (w_{\Pi} + w_{\mathrm{V}})' M_4((w_{\Pi} + w_{\mathrm{V}}) \otimes (w_{\Pi} + w_{\mathrm{V}}) \otimes (w_{\Pi} + w_{\mathrm{V}})) \end{aligned} \tag{5.2.28}$$

As described above, in order to obtain a favourable co-kurtosis, each of the above components can be calibrated according to their relative contributions via calibration of the covariance shrinkage factor and relaxation of particular constraints.

[51] The value of τ (covariance shrinkage factor) determines the relative weighting between the benchmark portfolio and view portfolio in the Black-Litterman model. If the investor wishes to increase the view component contribution, the value of τ should be increased. Within the Black-Litterman framework, an increase in τ increases the uncertainty in the benchmark returns, which in turn tilts the portfolio towards the view portfolio.

[52] In the same manner in which the marginal skewnesses were obtained, the marginal kurtosis due to the views is obtained by calculating the kurtosis of the view + benchmark portfolio and subtracting the skewness of the benchmark portfolio.

5.3 Relaxing the Long-Only Constraint

The long-only constraint has traditionally been observed as a risk management tool devised with the aim of protecting the portfolio from severe loss due to excessive risk taking (Scherer, 2007). In other words, asset managers have an incentive to increase risk taking in order to achieve superior returns. The long-only constraint, as well as other constraints, are therefore enforced in order to protect the client's portfolio from such excessive risk-taking. Even though constraints are enforced with the good intensions of protecting the investor, it can be noted that constraints may effectively *increase* risk taking, whereby particular constraints effectively impair diversification benefits as a result of reducing the breadth of active decisions by forcing asset managers to take risks in areas wherby the manager exhibits insufficient skills (Scherer, 2007).

Scherer (2007) notes that most investment constraints impose a severely outdated risk management approach, typical of the 1950s, due to the lack of the insights of portfolio theory. He further notes that only risk measures can appropriately control portfolio risk and conversely, the use of static constraints are unable to effectively control risk as they do not take into account changing volatilities, correlations as well as the infinite number of possible scenarios that may occur.

Clarke *et al.,* (2004) state that in terms of efficiency loss, the long-only constraint is the most often the most significant contributor and that relaxing the long-only constraint may effectively improve the information transfer, therefore resulting in superior risk adjusted returns. Portfolio constraints limit the asset manager's ability to fully capitalise on valuable information, thereby resulting in a significant loss in efficiency. Clarke *et al.,* (2004) illustrate that even a modest relaxation of the long-only constraint leads to a significant improvement in the transfer of information efficiency[53].

It must be noted that when the long-only constraint is marginally relaxed, in order to offset the short positions, excess long active positions may be taken, the net result being a portfolio of asset weights that still add up to 100%. Clarke *et al.,* (2004) find that the formation of a 120/20 portfolio[54] leads to a more informationally efficient transfer of investment information into the final portfolio active weights.

[53] Clarke et al., (2004) measure the improvement in information efficiency via the transfer coefficient (TC), which is calculated as the correlation between risk-adjusted expected returns and the weighted risk exposures of assets in the portfolio. In short, the TC measures the extent of information transfer from an asset ranking signal into active portfolio weights. The higher the value of the TC; the better the transfer of valuable signal information into asset weightings.

[54] A 130/30 portfolio is a portfolio composed of 20% short positions and 120% weighted in long positions. The net result is a portfolio of assets with weights added up to one.

5.3.1 The Benefits of Relaxing the Long-Only Constraint

Traditional long-only strategies limit the amount of underweighting of stocks in a portfolio, since the most the portfolio manager can underweight a particular stock is given by its weight in the benchmark. In other words, a portfolio manager marginally decreases the stock's weight in the portfolio up until the point where the stock is not held at all, given by a zero weighting. Therefore, stocks that exhibit a lower level of market capitalisation possess a low level of underweighting potential in the portfolio (Clarke *et al.,* 2004).

In a long-only portfolio, the investor is not able to fully reflect negative expectations of the return of particular assets in the portfolio since the underweighting ability is severely restricted. Therefore, by relaxing the long-only constraint by allowing short sales, the investor is then able to capitalise on valuable information regarding poorly ranked stocks in the portfolio. In this way, stocks that are expected to achieve a negative return are also able to contribute to the performance of the overall portfolio since they form a greater proportion of the entire portfolio. In addition, when short positions on unattractive stocks are utilised, additional attractive stocks can be further overweighted in order to offset the short positions[55]. This therefore provides greater diversification among the top ranked stocks and a greater level of exposure to the positive alpha opportunites given by the benchmark (Johnson *et al.,* 2007).

5.3.2 The Relationship between Active Risk and the Long-Only Constraint

The systematic overweighting and underweighting of asset weights within a portfolio refers to active asset allocation. For each active asset weighting, active risk is incurred as a result of diverting away from the benchmark allocations. Active risk is measured by the portfolio tracking error and, depending on the investment objectives of the investor, a target level of active risk (tracking error) is typically specified in order to take advantage of particular alpha opportunities that may arise. It can be noted that increasing the target tracking error (the active risk budget) of a portfolio while still enforcing the long-only constraint, naturally forces the portfolio manager to take greater active positions in assets that exhibit larger market capitalisation weights (Clarke *et al.,* 2004). This is due to the fact that smaller capitalisation weighted assets can only be underweighted by as much as their market capitalisation weighting until a zero allocation is reached. Therefore, active allocations are not efficiently utilised as they are forced into assets of the portfolio that typically contain greater market capitalisation weightings. It can therefore be noted that the greater the desired level of active investment, the greater the possibility of the long-only constraint being a binding factor in the optimal asset allocation. Therefore, for more active investment strategies, the long-only constraint will have a greater negative impact on the information transfer and it is for these particular portfolios where a mild relaxation of the long-only constraint will have the greatest positive influence on

[55] Since short sales finance buys, the underweights and overweights of assets in the portfolio must sum to zero. The portfolio manager can therefore make better use of negative alpha information as well as positive alpha information.

the optimal portfolio. All things being equal, higher tracking error investment strategies require higher levels of short sales (Johnson *et al.,* 2007).

As described by Clarke *et al.,* (2004), the formation of long-short portfolios can substantially add significant value over tradiational long-only constrained portfolios. Depending on the target level of active risk, significantly positive and negative alpha opportunities can be exploited without additing too much additional risk via marginally relaxing the long-only constraint.

Chapter 6: Application

The aim of this section is to apply the concepts discussed in the previous sections to practical examples in the South African context. In particular, section 6.1 deals with a simple example of applying the robust estimates of higher moments to a Domestic Balanced Fund and observing how the portfolio allocations change for increasing the risk aversions to the portfolio skewness and kurtosis. Section 6.2 deals with analysing the respective components of the Black-Litterman model as well as how a mild relaxation of the long-only constraint can improve the overall efficiency of the portfolio and the expression of the views of the model.

The aim of the following two applications is therefore to provide the investor with the necessary tools and insight in order to be aware of the practical implications of higher moments and constraints on the optimal portfolio.

6.1 Asset Allocation Under Higher Moments: An Example

In this section, a portfolio consisting of 6 asset classes will be analysed via a series of optimisations incorporating higher moments. More specifically, using a generic higher moment utility function, the optimal portfolio will be analysed according to varying magnitudes and combinations of the risk aversion parameters to portfolio variance, co-skewness and co-kurtosis.

In order to generate the optimal portfolios, the following generic utility function will be used[56]:

$$\max U = E(\mu' w) - \lambda_2 w' M_2 w + \lambda_3 w' M_3 (w \otimes w) - \lambda_4 w' M_4 (w \otimes w \otimes w) \tag{6.1.1}$$

where λ_2, λ_3 and λ_4 are the risk aversion parameters for variance, co-skewness and co-kurtosis respectively. Using the above utility function, it can be noted that the optimal portfolio depends on the investor's preferences[57] for return relative to variance, co-skewness and co-kurtosis. An investor who values positive portfolio skewness will therefore be inclined to express high values of λ_3 relative to λ_2 and λ_4. Therefore, as the values of the risk aversion parameters change, the respective risk objectives of the portfolio change.

[56] Robust estimates for μ, M_2, M_3 and M_4 were obtained using the methodology outlined in section 4.2.2
[57] As discussed in section 4.2, investors typically prefer lower variance, positive skewness and lower kurtosis.

As a result of varying the respective risk aversion parameters and observing the resulting portfolio, the following questions will be answered:

- How does the inclusion of higher moments (in addition to portfolio variance) in the utility function influence the optimal asset allocations?
- More specifically:
    - How do the asset allocations change for varying degrees of risk aversion?
    - How does the portfolio variance, skewness and kurtosis change?
- What are the observed relationships between the portfolio moments?

The following sections describe the data used as well as the results obtained when optimising the portfolio when higher moments are included.

6.1.1 Data Description

Assuming that investors exhibit preferences for higher moments, it can be noted that data which is significantly non-normal (possesses significant skewness and excess kurtosis) will be the most effective in illustrating how changes in the asset allocation can affect the portfolio's higher moment characteristics.

For this particular example, a Domestic Balanced Fund consisting of six asset classes will be analysed. Analysing a Domestic Balanced Fund with respect to its higher order co-moments has a significant practical element in that pension fund investors, especially those nearing retirement, will intuitively be highly risk averse towards excess portfolio kurtosis, thereby resulting in a portfolio return distribution that exhibits fat left tails – indicating the increased probability of a catastrophic loss to the portfolio. As described by Fabozzi *et al.,* (2009), investors are also willing to forgo additional return in preference of a more positively skewed return distribution, thereby indicating a higher frequency of positive returns.

The domestic balanced fund to be examined is composed of the following indices: Domestic Equities (J203), Domestic Listed Property (J253), Domestic Preference Shares (J251), Domestic Fixed Income (ALBI), Domestic Inflation-Linked (IFL) Bonds (BSAGI) and Domestic Cash (STFIND). The frequency and period of observation of the data will be monthly returns on the respective indices from 31 March 2002 to 29 February 2012 (120 observations for each asset)[58].

---

[58] In order exercise prudence when estimating the higher order co-moments, a large dataset was used despite the fact that structured estimates were used. It can be noted that for the sample estimates, for a portfolio of $n = 6$ assets, the following number of parameters are required to be estimated:

$$M_2: \frac{1}{2}n(n+1) = 21 \qquad M_3: \frac{1}{6}n(n+1)(n+2) = 56 \qquad M_4: \frac{1}{24}n(n+1)(n+2)(n+3) = 126$$

Given that the robust methodology discussed in section 4.2.2 was used to estimate the higher order moments, the dataset of 120 observations for each asset is sufficient.

As shown in table 1 on the following page, Domestic Equity and Domestic Listed Property earned a significantly greater annualised monthly return of 15.1% and 15.5% compared to the remaining assets in the portfolio. It should be noted that while their respective mean returns are highly favourable, they both exhibit significantly larger (annualised monthly) return standard deviations (17.9% and 16.5%) as well as significant negative skewness (-0.48 and -0.15) in conjunction with severe excess kurtosis coefficients, thereby resulting in a return distribution that exhibits “fat left tails”; thereby exposing the investor to an increased probability of significantly large and devastating losses.

The Domestic Preference Shares asset exhibits the lowest annualised monthly return in the portfolio (1.7%). However, it exhibits a favourable positive skewness coefficient (0.07) as well as a comparatively low kurtosis coefficient (3.33); the combination of which is very attractive to any investor and its inclusion in the portfolio could show to posses significant risk diversification properties (Fabozzi *et al.,* 2007).

Domestic Fixed Income and Domestic IFL Bonds possess similar characteristics. Both assets provide a moderately high annualised monthly return of 10.8% and 11.5% with comparatively low annualised monthly standard deviations of 6.7% and 4.2% respectively. Due to the low variances, fixed income assets are renowned for providing relatively stable income growth (Hwang & Satchell, 1999). However, despite the favourable return and variance properties, the Domestic Fixed Income asset exhibits unfavourable kurtosis (5.35) and a significantly large negative skewness (-0.25). Conversely, Domestic IFL Bonds exhibits a significantly lower kurtosis (3.98) and a large positive skewness, thereby making it an extremely attractive asset in the portfolio.

Domestic Cash, being a riskless asset with a standard deviation near zero, possesses a comparatively low annualised return. In addition to risk free nature, cash also has a positive skewness (0.52) and extremely low kurtosis (2.12), thereby making it an ideal asset to hold in the portfolio when risk aversion to any moment (variance, skewness and kurtosis) significantly increases.

In addition to the above summary statistics, the Jarque-Bera test was used to test each asset for normality. Under the null hypothesis that the distribution follows a normal distribution, it was concluded that every null hypothesis was rejected, therefore inferring that none of the assets follow a normal distribution. In addition, the individual p-values were very small, with Domestic Preference Shares having the greatest p-value of 0.009, which was significantly lower than the critical 5% level.

Table 1: Summary Statistics for the Six Asset Classes of a Domestic Balanced Fund

	Domestic Equity	Domestic Listed Property	Domestic Preference Shares	Domestic Fixed Income	Domestic IFL Bonds	Domestic Cash
Mean	15.1%	15.5%	1.7%	10.8%	11.5%	8.5%
Median	17.8%	16.6%	-0.6%	9.3%	10.2%	7.8%
Standard Deviation	17.9%	16.5%	8.4%	6.7%	4.2%	0.6%
Sample Variance	0.032	0.027	0.007	0.005	0.002	0.000
Standardised Kurtosis	5.57	6.57	3.33	5.35	3.98	2.12
Standardised Skewness	-0.48	-0.15	0.07	-0.25	0.64	0.52
Jarque-Bera Test	Reject (p=0.002)	Reject (p=0.001)	Reject (p=0.009)	Reject (p=0.001)	Reject (p=0.007)	Reject (p=0.009)

6.1.2 Relationship between Variance, Skewness and Kurtosis

In order to appropriately interpret the asset allocation changes that are a result of changes to the risk aversion parameters, it is worthwhile considering the relationships between the respective moments of the distribution.

The variance of the portfolio determines the level of spread of the asset returns, while the skewness describes the asymmetry of the return distribution. The kurtosis of the distribution describes the height or peak of the distribution. It can be noted that when the risk aversion to variance increases, one would expect the distribution of returns to become narrower and more centred about the mean. A higher frequency of returns about the mean inevitably results in an increased kurtosis. Conversely, as the risk aversion to kurtosis increases, the distribution of returns flattens thereby resulting in an increased variance.

The following simple example below best illustrates how the portfolio skewness, kurtosis and standard deviation changes for varying weights in a two asset portfolio. For the figures below, a two asset portfolio containing Domestic Equities and Domestic Fixed Income is examined according to changes in the respective asset weightings. The weight in Domestic Equities is given by w and similarly, the weight given to Domestic Fixed Income is given by $(1-w)$, with the value of w varying between zero and one on the horizontal axis. The figures below illustrate the portfolio standard deviation, skewness and kurtosis for increasing the weight in Domestic Equities relative to Domestic Fixed Income[59].

[59] For purposes of illustration, data between 31 March 2010 and 29 February 2012 was used to construct figures 2, 3 and 4.

Figure 2: Analysing the Portfolio Standard Deviation for an Increase in the Weight in Domestic Equities

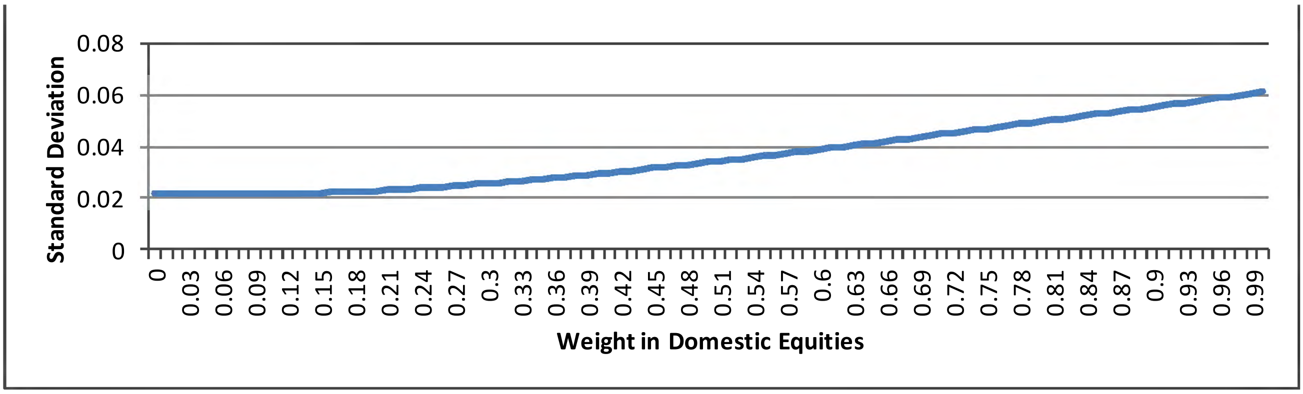

Figure 3: Analysing the Portfolio Skewness for an Increase in the Weight in Domestic Equities

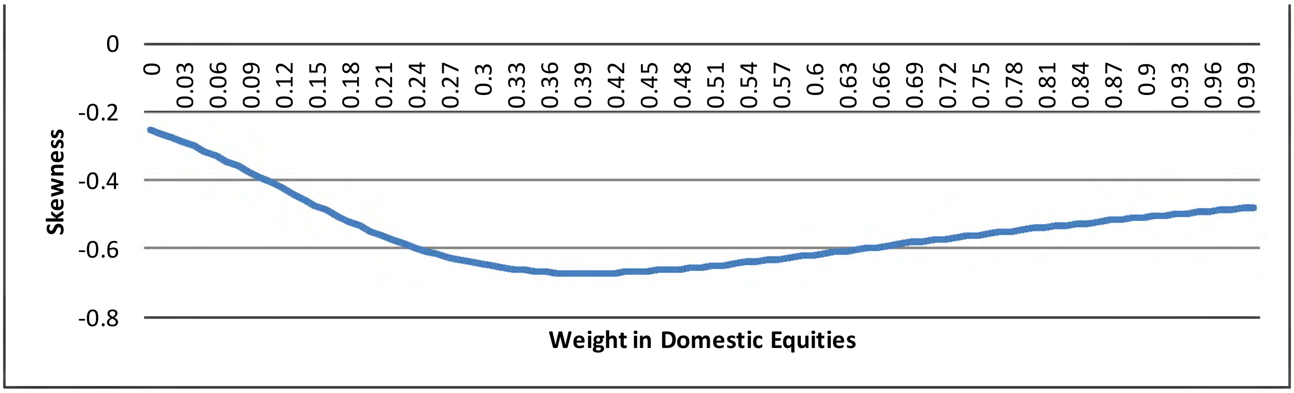

Figure 4: Analysing the Portfolio Kurtosis for an Increase in the Weight in Domestic Equities

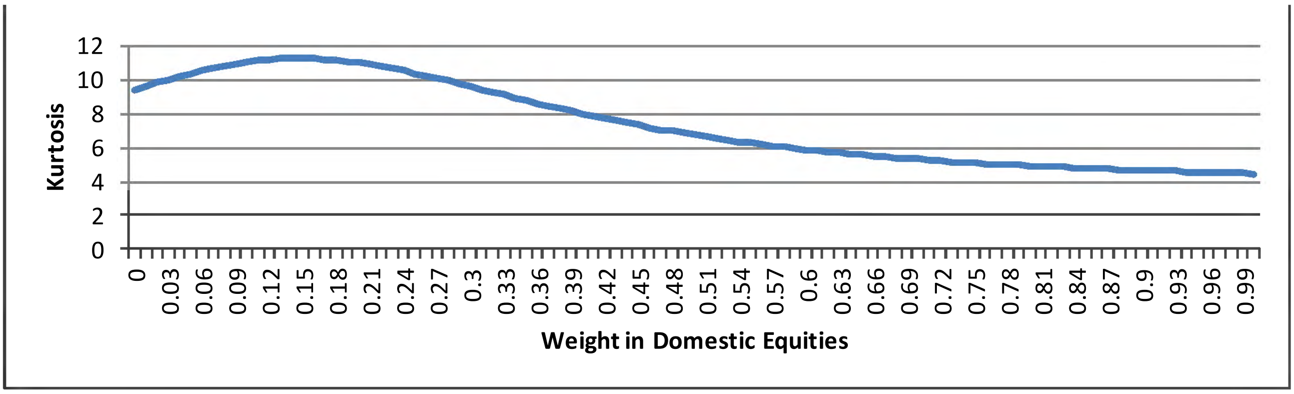

As shown by the figures above, the minimum standard deviation portfolio occurs for when $w = 0.08$, thereby resulting in a portfolio that is composed of 8% in Domestic Equities and the majority of 92% invested in Domestic Fixed Income. It is also clear that the behaviour of the standard deviation for varying values of w is convex while it can also be noted that the skewness and kurtosis of the portfolio are clearly nonlinear functions of w and exhibit multiple points of maxima and minima.

For this particular portfolio, it is clear that any investor exhibiting strong preferences for higher moments, the minimum variance portfolio will be suboptimal. As shown above, while the minimum variance portfolio (8% Domestic Equities and 92% Domestic Fixed Income) also approximately minimises the portfolio skewness, it also results in an approximate maximisation of the portfolio kurtosis. Assuming that the investor exhibits strong preferences for a lower kurtosis[60], it is clear that in this particular case, the minimum variance portfolio will be undesirable. Furthermore, it is clear that the classical mean-variance approach would inevitably recommend a suboptimal portfolio and that a higher moment utility function be used in order to capture the particular higher moment preferences of the investor as well as the respective trade-offs when optimising a portfolio exhibiting higher moments.

It must be noted that the above two asset portfolio is a significant simplification of the dynamics between variance, skewness and kurtosis of the portfolio for varying asset weights. For an increasing number of assets involved in the portfolio, one has to consider the co-moments that exist between the assets. As a result of the diversifying dynamics between the respective co-moments, there is a significant opportunity for the assets to result in more favourable portfolio moments.

The following sections illustrate how the portfolio moments (variance, skewness and kurtosis) change for increasing risk aversions to the respective moments of the return distribution for portfolios that include six asset classes that exhibit typically non-normal return behaviour.

6.1.3 The Impact of Skewness on the Optimal Portfolio

The aim of this section is to analyse the optimal portfolio for changes to the skewness risk aversion coefficient (λ_3). The effects of changing λ_3 will be observed for its effect on the asset allocations as well as the respective portfolio moments, such as variance, skewness and kurtosis. In addition, the effects of changing λ_3 on the optimal portfolio will be observed for its inclusion as the only risk measure as well as its inclusion along with the portfolio variance.

Using the generic utility function below, the effects of changing λ_2 and λ_3 will be observed for the optimal portfolio.

$$\max U = E(\mu' w) - \lambda_2 w' M_2 w + \lambda_3 w' M_3 (w \otimes w) \tag{6.1.2}$$

Subject to the following long-only and equality constraints:

$$w \geq 0$$

$$\sum_{i=1}^{6} w_i = 1$$

It must be noted that increasing the value of λ_3 results in an increase in the risk aversion towards negative portfolio skewness. Therefore, taking into account the joint co-skewness between the respective assets,

[60] It is rational to assume that an investor's preference would be to achieve high odd moments (mean and skewness), as this would decrease the possibility of extreme values on the side of losses and increase the probability for increases in gains. In addition, investors prefer low even moments (variance and kurtosis), which results in decreased dispersion and uncertainty of the returns.

one would expect the portfolio to tilt away from strongly negatively skewed assets in favour of positively skewed assets.

Portfolio Skewness for Simultaneously Changing Risk Aversions

Figure 5 below plots the portfolio's third moment[61] (z-axis) for simultaneous changes in the risk aversion to portfolio variance (x-axis) and aversion to negative portfolio skewness (y-axis). The risk aversion to variance ranges between 0.5 and 30 in steps of 0.5 ($0.5 \leq \lambda_2 \leq 30$), while the risk aversion to skewness ranges between 0 and 30 ($0 \leq \lambda_3 \leq 30$). Therefore, each combination of λ_2 and λ_3 represents a coordinate on the surface corresponding to a particular portfolio skewness.

Figure 5: Analysing the Third Portfolio Moment for Simultaneous Changes in the Risk Aversions to Skewness and Variance

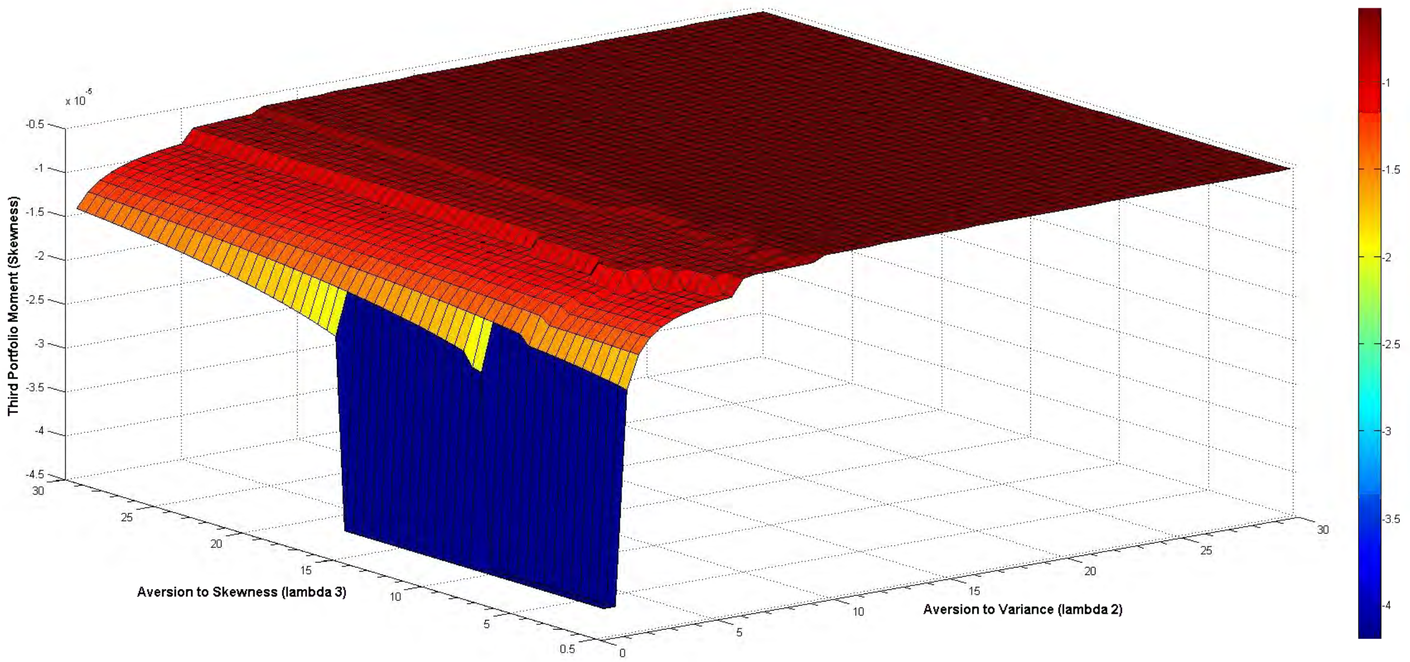

As shown in figure 5 above, when the risk aversion to variance is zero ($\lambda_2 = 0$), increases in the risk aversion to portfolio skewness do not result in any significant changes to the portfolio skewness for values of λ_3 between 0.5 and 12. However, for values of $\lambda_3 > 15.5$, marginal decreases in the portfolio skewness are observed. It can be noted that for infinitely increasing values of λ_3 (with $\lambda_2 = 0$), the asset allocations will progressively move away from higher return assets and towards a portfolio consisting of positively skewed assets, namely Fixed Income, IFL Bonds, Preference Shares and Cash.

It is interesting to note that the portfolio skewness sharply decreases (becomes closer to zero) for small increases in the value of λ_2, regardless of the risk aversion towards negative skewness. It can be observed that values of $\lambda_2 \geq 2$, it becomes optimal for the portfolio to switch to positions held in assets that are not

[61] It can be noted that portfolio skewness is synonymous with the standardised third moment. Therefore, analysing changes to the third moment provides a clear picture of how the aversion to portfolio skewness is affected.

only exhibit low variances, but also possess favourable diversification properties which facilitate the achievement of an overall lower level of negative portfolio skewness.

In order to provide further insight into how increasing the risk aversion to negative skewness and variance influences the optimal portfolio, the optimal asset allocations for increasing risk aversion to negative skewness (λ_3), are illustrated in figures 6 and 7 for values of $\lambda_2 = 0$ and $\lambda_2 = 2$ respectively. In figure 6 below, when preferences for variance are excluded from the optimisation problem ($\lambda_2 = 0$), the optimal portfolio allocations remain fairly constant for low levels of λ_3, thereby indicating that, given the co-skewness structure between the respective assets, there is very little opportunity to decrease the overall skewness of the portfolio without forgoing a considerable amount of return. However, for $\lambda_3 \geq 15$, one observes a abrupt shift in the asset allocations: Equity drops to 40.8% (from 55%), Property increases to 3.2% (from 2%), Preference Shares increase to 5.4% (from 3%), Fixed Income remains fairly stationery at 30.4% while IFL Bonds and Cash increase to 11.9% and 8.2% (from 7% and 3%). We therefore see the portfolio shift away from the most negatively skewed asset (Equity) in favour of the most positively skewed assets (Cash and IFL Bonds). A possible justification for the increase in allocation for Property, Preference Shares and Fixed Income, which exhibit negative skewnesses, could be due to their attractive mean return and diversity enabling co-skewness properties.

Figure 6: Analysing the Changes of the Asset Allocations for an Increase in the Risk Aversion to Skewness for a Zero Risk Aversion to Variance

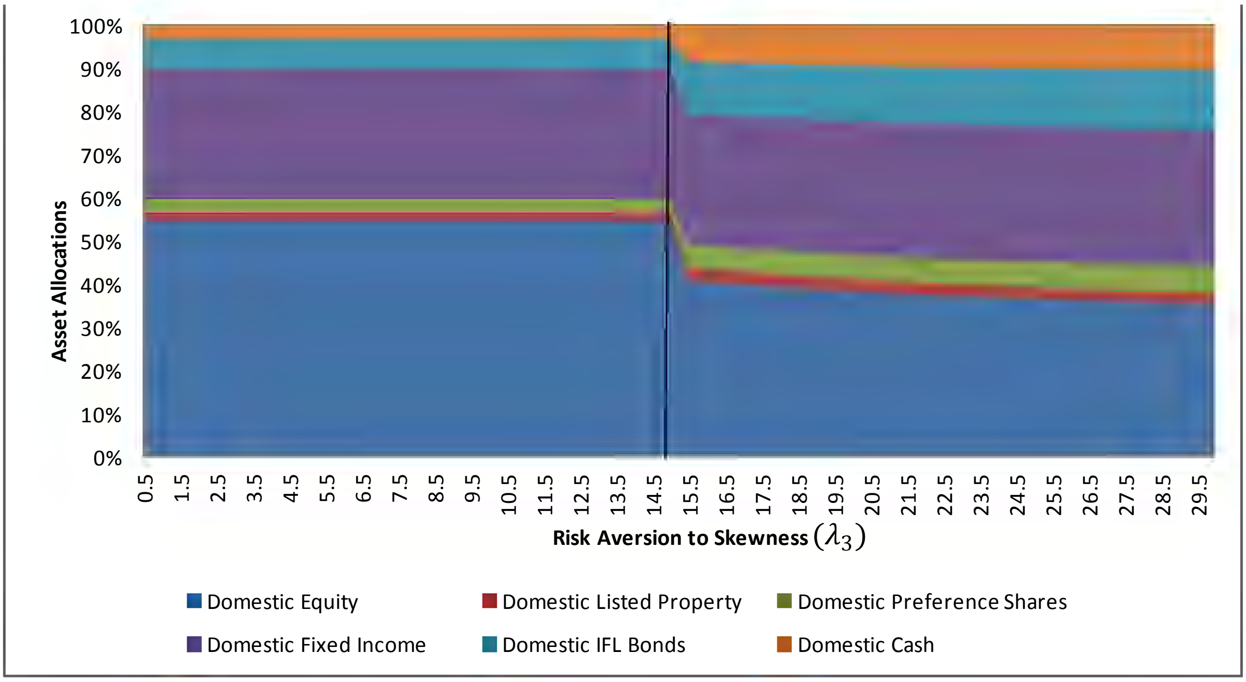

Figure 7: Analysing the Changes in the Asset Allocations for an Increase in the Risk Aversion to Skewness when the Risk Aversion to Variance is fixed at 2

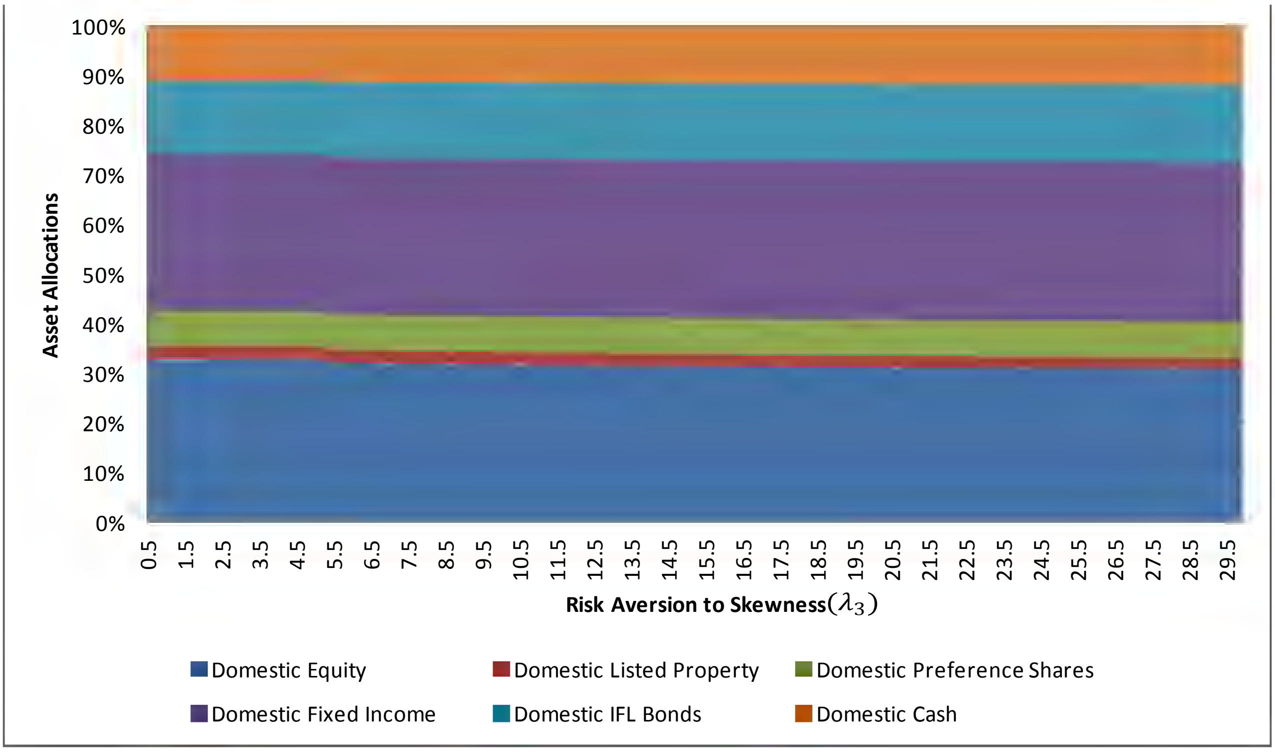

As shown in figure 7 above and, in comparison to figure 6, when a low level of risk aversion to portfolio variance ($\lambda_2 = 2$) is introduced, the asset allocations dramatically change. However, while there is a significant reaction of the asset allocations to the introduction of risk aversion to variance, the allocations remain insensitive to changes in the risk aversion to skewness, as shown by the *relatively* horizontal lines for increasing values of λ_3.

In order to better interpret figures 6 and 7 above, table 2 below numerically illustrates how the asset allocations differ for values of λ_2 and λ_3. As previously stated, for $\lambda_2 = 0$, one observes significant changes to the asset allocations for changes in λ_3. For $\lambda_2 = 2$, in the second column of table 2, with the exception for the decrease in allocation for Domestic Equity, there is no significantly large change (greater than 1%) for any of the other asset classes for changes in the value of λ_3 between 0 and 30. More specifically, for $\lambda_2 = 2$, very small changes to the asset allocation are observed for large values λ_3. Ironically, the lowest (most favourable) third portfolio third moment (-0.103×10^{-4}), is achieved for $\lambda_2 = 2$ and $\lambda_3 = 30$, which corresponds to a standardised skewness of -0.92. It can be noted that this figure is a slight improvement over the case where the risk aversion to skewness is ignored for $\lambda_2 = 2$ and $\lambda_3 = 0$.

Table 2: Summary of the Asset Allocations and Portfolio Statistics for Various Combinations of Risk Aversions to Skewness and Variance

	$\lambda_2 = 0$			$\lambda_2 = 2$		
	$\lambda_3 = 15$	$\lambda_3 = 15.5$	$\lambda_3 = 30$	$\lambda_3 = 0$	$\lambda_3 = 15$	$\lambda_3 = 30$
Domestic Equity	55%	40.8%	35.6%	32.7%	31.48%	30.9%
Domestic Listed Property	2%	3.2%	2.3%	2.8%	2.32%	2.05%
Domestic Preference Shares	3%	5.4%	6.5%	7.2%	7.42%	7.5%
Domestic Fixed Income	30%	30.4%	31.5%	31.7%	32.01%	32.2%
Domestic IFL Bonds	7%	11.9%	13.9%	14.7%	15.28	15.6%
Domestic Cash	3%	8.2%	10.2%	10.9%	11.49%	11.8%
Third Moment	-0.419×10^{-4}	-0.2×10^{-4}	-0.143×10^{-4}	-0.119×10^{-4}	-0.108×10^{-4}	-0.103×10^{-4}
Standardised Skewness	-0.98	-0.97	-0.95	-0.95	-0.93	-0.92

The above tables and figures have illustrated that for a small 6 asset domestic balanced fund portfolio, the introduction of portfolio skewness into the generic utility function results in significant changes to the optimal portfolio only when the investor expresses strong preferences for positive portfolio skewness (values of $\lambda_3 \geq 15$). However, the figures have also illustrated that when the investor expresses preferences for lower portfolio variance, the asset allocations become fairly insensitive to changes in λ_3. Ironically, as shown by figure 5, small increases in the risk aversion to variance (λ_2), regardless of the expressed value of λ_3, results in a more significant reduction in negative portfolio skewness. However, this could be due to the fact that the low variance assets (Preference Shares, IFL Bonds and Cash) coincidently exhibit positive skewness coefficients.

Given the above results, it can be concluded that, for a generic utility function, based on this particular 6 asset Domestic Balanced Fund Portfolio, the inclusion of risk aversion to portfolio skewness (λ_3) becomes somewhat redundant when implemented in conjunction with the risk aversion to portfolio variance. This is due to the fact that the portfolio skewness was approximately minimised (became less negative) when the portfolio variance was minimised, thereby making the inclusion of the third moment somewhat redundant. However, despite the above obtained results, an asset allocation for a different set of assets, a larger number of assets, a different investor utility function and relaxing the imposed constraints, may result in dramatically different asset allocations[62].

[62] Fabozzi et al., (2007) describe and illustrate how the overall portfolio skewness is a non-linear function of the individual asset skewnesses. The non-linear nature of the portfolio skewness as well as the structure of the co-skewness matrix between the assets, best explains why the overall portfolio skewness can be more extreme than the

6.1.4 Impact of Portfolio Kurtosis on the Optimal Portfolio

Keeping consistent with the methodology described in the previous section, the attention will now shift to including the portfolio kurtosis in the generic utility function optimisation.

In order to analyse the effects of changing risk aversion to variance and kurtosis, the following generic utility function was used:

$$\max U = E(\mu' w) - \lambda_2 w' M_2 w - \lambda_4 w' M_4 (w \otimes w \otimes w) \quad (6.1.3)$$

Subject to the following long-only and equality constraints:

$$w \geq 0$$

$$\sum_{i=1}^{6} w_i = 1$$

As per the previous section, figure 8 below plots the portfolio fourth moment[63] (z-axis) for simultaneous changes to the risk aversion to portfolio variance (λ_2) (x-axis) and aversion to portfolio kurtosis (λ_4), (y-axis). The risk aversion to variance ranges between 0 and 30 in steps of 0.5 ($0 \leq \lambda_2 \leq 30$), while the risk aversion to kurtosis (λ_4) ranges between 0.5 and 30 ($0.5 \leq \lambda_4 \leq 30$). Therefore, each combination of λ_2 and λ_4 represents a coordinate on the surface corresponding to a particular portfolio fourth moment. It can be noted that rational investors prefer lower portfolio kurtosis and therefore, as the value of λ_4 increases, one would expect the fourth moment to decrease.

Figure 8: Analysing the Fourth Portfolio Moment for Simultaneous Changes in the Risk Aversions to Kurtosis and Variance

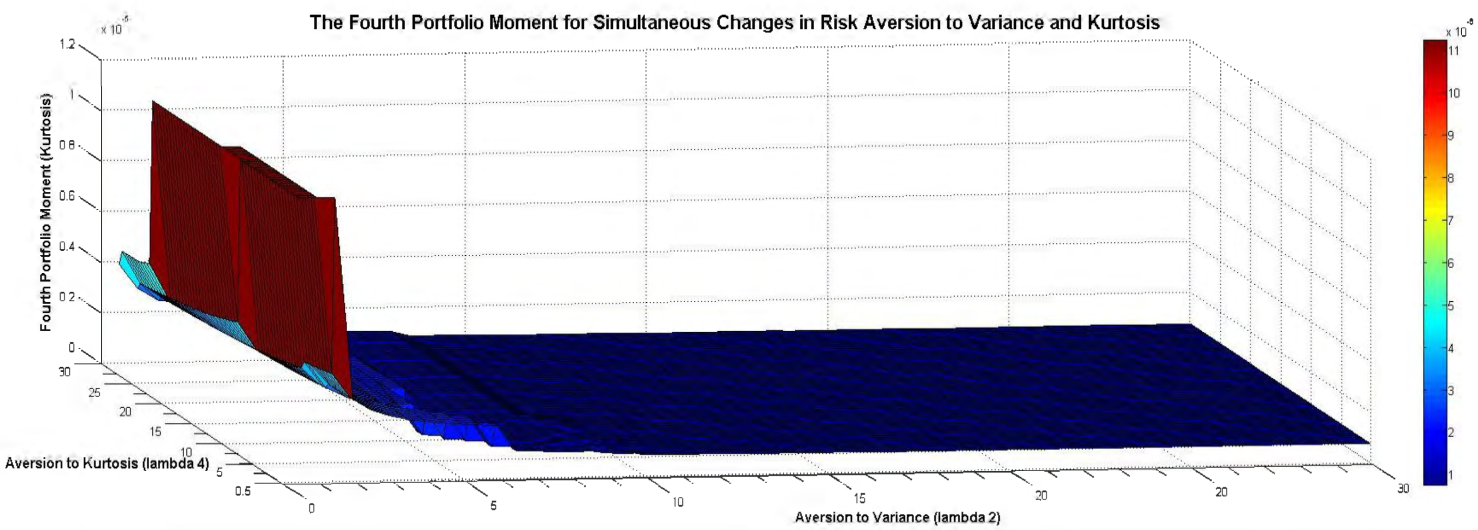

extremely skewed asset in isolation. They further show that for discontinuous or S-shaped utility functions, that there is a significant loss in utility when higher moments are ignored.

[63] It can be noted that portfolio kurtosis is synonymous with the standardised fourth moment. Therefore, analysing changes to the fourth moment provides a clear picture of how the aversion to portfolio kurtosis is affected.

As shown in figure 8 above, when the risk aversion to variance is excluded from the utility function $(\lambda_2 = 0)$, for values of $\lambda_4 < 24.5$, no reduction in the portfolio fourth moment is observed. However, for values of $\lambda_4 \geq 24.5$, the result is a sudden shift in the asset allocations, which in turn results in a sharp decrease in the portfolio fourth moment. Similarly, when risk aversion to variance is introduced $(\lambda_2 > 0)$, almost no decrease in the portfolio fourth moment is observed in figure 8 for changes in λ_4. More specifically, for a small change in $\lambda_2 > 1$, a sharp decrease in the fourth moment and an eventual steadying decline for $\lambda_2 \geq 6$ is observed.

Figure 9 below plots the asset allocations for increasing risk aversion to kurtosis for the risk aversion to variance set at zero $(\lambda_2 = 0)$. Consistent with figure 8, there is no significantly observable change in the asset allocations for $\lambda_4 < 24.5$ and thereafter, a sudden shift in the optimal asset allocation is observed. For $\lambda_4 \geq 24.5$, the optimal allocation for the two asset classes with the greatest kurtosis drops significantly. As further illustrated in table 3, Equity decreases from 55% to 43.1% and Domestic Listed Property falls to 0.12%, almost being excluded from the optimal allocation completely. Conversely, the optimal allocation for Cash, Fixed Income and IFL Bonds increase significantly while Domestic Preference Shares increase slightly. In addition to the changes in asset allocations, the fourth moment decreases significantly, from 0.112×10^{-4} to 0.472×10^{-5}. However, despite the decrease in the fourth moment, the standardised kurtosis[64] increased from 7.49 to 7.82. This is a strong indication that minimising the portfolio kurtosis also minimises the portfolio variance. In order to further illustrate how the asset allocation change for increasing λ_4 when $\lambda_2 = 0$, the asset allocations for $\lambda_4 = 30$ are included in table 3. The result of an increase from $\lambda_4 = 25$ to $\lambda_4 = 30$, while not as prominent as the change from $\lambda_4 = 0.5$ to $\lambda_4 = 25$, makes it clear that as λ_4 increases, the portfolio continues to shift away from assets with significantly larger kurtosis parameters in favour of assets that result in an overall lower portfolio fourth moment and lower variance.

[64] This is due to the fact that, as the asset allocations changed, switching to positions aimed at minimising the fourth moment, the portfolio variance was also decreased, therefore resulting in a greater ratio of the fourth moment to the square of the variance; given as: $Standardised\ K = \frac{w' M_4 (w \otimes w \otimes w)}{(w' M_2 w)^2}$

Figure 9: Analysing the Changes in the Asset Allocations for an Increase in the Risk Aversion to Kurtosis when the Risk Aversion to Variance is Zero

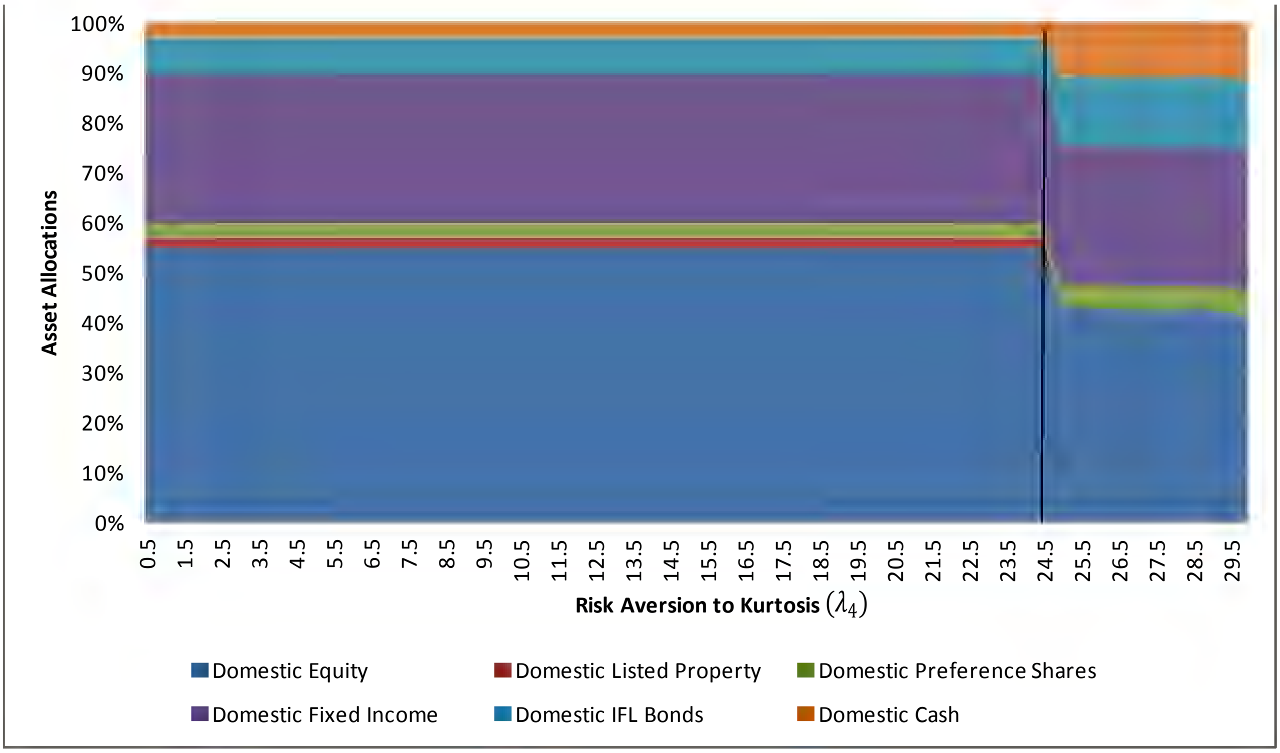

Figure 10 below plots the asset allocations for increasing risk aversion to kurtosis (λ_4) for a low risk aversion to variance, with $\lambda_2 = 2$. In addition to figure 10, the second half of table 3 further illustrates how the asset allocations change for increasing λ_4 with $\lambda_2 = 2$.

As illustrated in figure 10 below, it can be observed that, when variance is included in the utility optimisation, the optimal asset allocations become fairly insensitive to changes in λ_4. For increases in λ_4, the result can be interpreted as an overall (minimal) decline in the allocation to equity, from 32.8% to 31.5% for $\lambda_4 = 0.5$ and $\lambda_4 = 30$ respectively. In addition, Domestic Listed Property and Fixed Income both decline from 2.8% to 0.9% and from 31.6% to 30.8%. The decline in allocation Equity, Fixed Income and Domestic Listed Property is offset by a corresponding increase in Domestic IFL Bonds and Domestic Cash (from 14.7% to 16.5% and from 10.9% to 13%). It can be noted that an overall lower level for the fourth portfolio moment is achieved when variance is included in the utility function; whereby for each value of λ_4, a lower fourth moment is achieved when $\lambda_2 = 2$ as compared to the case when $\lambda_2 = 0$. However, despite the fact that a lower fourth moment is achieved when variance is included in the optimisation, the standardised kurtosis increases significantly. This is due to the relative scaling of the variance versus the fourth portfolio moment.

Figure 10: Analysing the Changes in the Asset Allocations for an Increase in the Risk Aversion to Kurtosis when the Risk Aversion to Variance is fixed at 2

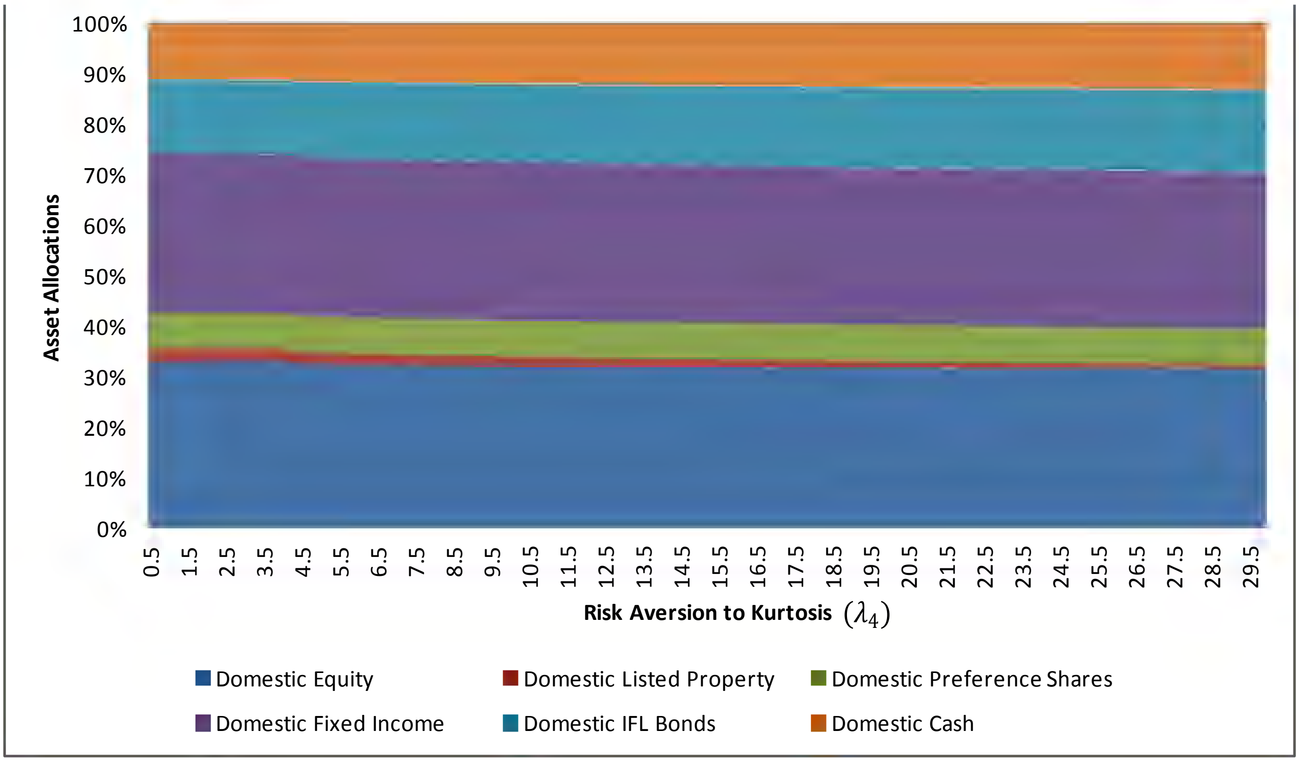

Table 3: Summary of the Asset Allocations and Portfolio Statistics for Various Combinations of Risk Aversions to Kurtosis and Variance

	$\lambda_2 = 0$			$\lambda_2 = 2$		
	$\lambda_4 = 0.5$	$\lambda_4 = 25$	$\lambda_4 = 30$	$\lambda_4 = 0.5$	$\lambda_4 = 25$	$\lambda_4 = 30$
Domestic Equity	55%	43.1%	40%	32.8%	31.6%	31.5%
Domestic Listed Property	2%	0.12%	0.013%	2.8%	1.09%	0.9%
Domestic Preference Shares	3%	4.8%	5.2%	7.2%	7.3%	7.3%
Domestic Fixed Income	30%	27.5%	27.5%	31.6%	30.9%	30.8%
Domestic IFL Bonds	7%	13.7%	14.8%	14.7%	16.3%	16.5%
Domestic Cash	3%	10.7%	11.9%	10.9%	12.8%	13%
Fourth Moment	0.112×10^{-4}	0.472×10^{-5}	0.389×10^{-5}	0.280×10^{-5}	0.232×10^{-5}	0.227×10^{-5}
Standardised Kurtosis	7.49	7.82	7.97	9.44	9.25	9.21

As illustrated above, inclusion of the fourth portfolio moment in the generic utility function does not result in any significant changes to the portfolio allocations for low values of λ_4. In particular, changes to the asset allocations, with respect to increasing λ_4, were only observed for significantly large values of $\lambda_4 > 25$. In addition, consistent with the results observed in section 6.1.3, when the risk aversion to variance (λ_2) was included, the effect of increasing λ_4 on the optimal asset allocations was significantly reduced.

6.1.5 Conclusions

The higher moment generic utility function was applied to a 6 asset domestic balanced fund and observations were made in order to determine whether the inclusion of the third and fourth moment, in addition to the portfolio variance, effectively made a difference to the optimal asset allocations as well as the overall moments of the optimal portfolio distribution.

It was determined that, for a generic utility function and in the absence of risk aversion to variance, the inclusion of higher moments only made a significant difference to the optimal allocations (and portfolio moments) when a strong aversion to portfolio skewness and *very* strong aversion to kurtosis was expressed. More specifically, the risk aversion to skewness (λ_3) only made a significant difference to the asset allocations for values of $\lambda_3 \geq 15$. For the portfolio kurtosis, the risk aversion to kurtosis only made a difference to the asset allocations for $\lambda_4 \geq 24.5$. In both cases, a sudden shift in the asset allocations was observed. For the aversion to skewness, the optimal portfolio shifted away from highly negatively skewed assets, such as Equity, in favour of IFL Bonds and Cash. For the aversion to kurtosis, the optimal portfolio shifted away from assets that exhibited high kurtosis (Equity and Listed Property) in favour of assets that exhibited a lower level of kurtosis, such as IFL Bonds, Cash and Preference Shares.

For the case when a low level of risk aversion to variance was included for $\lambda_2 = 2$, it was found that the aversion to variance dominated the asset allocations and that the allocations were fairly insensitive to changes to the risk aversion to skewness and kurtosis. In particular, when aversion to variance was expressed, fairly large values of λ_3 and λ_4 were required in order to result in even a minimal change to the asset allocations. It was therefore determined that the aversion to variance aversion dominates the portfolio allocations such that the risk aversion to negative skewness and higher kurtosis components of the generic utility function effectively become irrelevant.

In addition to noting how the introduction of higher portfolio moments influences the optimal asset allocation, it was observed that increasing the risk aversion to variance, generally resulted in a decrease in the standardised skewness and a general increase in the standardised kurtosis. This result was expected as a decrease in the dispersion of returns results in a distribution of returns further centred about the mean, which in turn results in a lower level of skewness as the amount of observations situated in the tails of the distribution decreases. As for kurtosis, the returns are situated more closely about the mean and therefore a higher peak (kurtosis) for the distribution is observed.

It therefore follows that, for a generic utility function and for *this particular six-asset Domestic Balanced Fund*, the mean-variance optimisation suitably approximates the solution and that the addition of higher moments does not produce any significant/notable results, unless significantly high preferences for positive skewness and lower kurtosis are expressed. However, as previously discussed, a larger portfolio with different assets may yield significantly contrasting results and therefore, it is important that the investor considers the relative effects of higher moments in the portfolio before naively applying a mean-variance optimisation to the portfolio.

Given the obtained results of the Domestic Balanced Fund, the practical implication of the obtained results is that, if an investor does not possess considerably high preferences for higher moments, then the mean-variance optimisation is effective in efficiently determining the optimal portfolio.

6.2 Analysing the Views and Components of the Black-Litterman Model

The aim of this section is to apply the concepts discussed in section 5.2.2. Section 6.2.2 will deal with identifying and illustrating the three influential components of the Black-Litterman model, namely the benchmark, view and constraints components. Section 6.2.4 will specifically deal with analysing the effects of the views on the final portfolio by directly analysing the active weight contributions and tracking error that arises as a direct result of the expressed views. Section 6.3 deals with analysing how constraints distort the original views expressed by the investor as well as how a mild relaxation of the long only constraint can result in a significant reduction in the distorting effects of the constraints.

The following section describes the data to be used in this example as well as how the views are formed.

6.2.1 Description of the Data and the Expressed Views

A portfolio of 40 stocks was selected from the FTSE JSE ALSI top 40 for the period between February 2002 and February 2012. A series of 120 observations for each asset was used to estimate a robust estimate of the covariance matrix between the assets. The market capitalisation weights as of February 2012 were used to calculate the portfolio weights. Using the covariance matrix, the benchmark weights and a risk aversion coefficient of $\lambda_2 = 2.5$, the implied equilibrium returns were calculated[65].

Following the Black-Litterman methodology for expressing *absolute* views, a series of views and associated confidence levels were expressed for the respective implied equilibrium returns for each asset. The following sections describe how the views were expressed as well as a table summarising the assets and associated views is provided.

Expressing Views and Confidence Levels

Within the framework of the Black-Litterman model, investors are able to express views in the form of absolute and relative views. For this particular application, the views expressed will be absolute views only. In order to simplify the complexity of generating meaningful views, the following expression, as described by Meucci (2009), will be used:

$$Q_k = (P\Pi)_k + \frac{\eta_k}{5}\sqrt{(P\Sigma P')_{kk}} \tag{6.2.1}$$

In the above expression, the value of the view for asset k (Q_k) is expressed as a function of the benchmark implied equilibrium return (Π) and the standard deviation of the particular asset ($\sqrt{(P\Sigma P')_{kk}}$).

For each asset in the portfolio, investors are required to provide a **view score** value, η_k. The value of η_k ranges between -10 and 10 and determines the extent to which the investor is positive (or negative) on a

[65] The value of λ_2 was calibrated in order to ensure a tracking error of 3% was maintained.

particular stock. For example, a value of $\eta_k = -10$ is interpreted as an extremely negative view and, in contrast, a value of $\eta_k = 10$ refers to an extremely positive view. A value of $\eta_k = 0$ represents a neutral view score, resulting in the investor view Q_k equalling the implied equilibrium return Π_k for asset *k*. For example, if the investor is extremely negative on a particular asset, the expectation of the return on that particular share would be two standard deviations lower that the return implied by the equilibrium returns. Therefore, the degree of the magnitude of the view ranges between -2 and 2 standard deviations of the particular asset's return.

As shown in table 4 below, the values of the view scores provided by the investor are split into eight respective categories, ranging from extreme negative $(-10 \leq \eta < -7.5)$ to extreme positive $(7.5 < \eta \leq 10)$. Each respective category is colour coded and represented by alphabetical characters for simple representation.

Table 4: Legend for Determining the Respective View Score Categories

Extreme Negative (A)	Strong Negative (B)	Weak Negative (C)	Mild Negative (D)	Mild Positive (E)	Weak Positive (F)	Strong Positive (G)	Extreme Positive (H)
$(-10 \leq \eta < -7.5)$	$(-7.5 \leq \eta < -5)$	$(-5 \leq \eta < -2.5)$	$(-2.5 \leq \eta < 0)$	$(0 < \eta \leq 2.5)$	$(2.5 < \eta \leq 5)$	$(5 < \eta \leq 7.5)$	$(7.5 < \eta \leq 10)$

The original Black-Litterman model assumes that investor views are normally distributed with a mean and variance, $Q \sim N(Q, P\Sigma P)$. The investor is therefore required to input a particular variance for each view in order to approximate the degree of confidence of the expressed view. It can be noted that many practitioners experience difficulty in expressing confidence levels in terms of variances (Idzorek, 2005). Therefore, in order to simplify the specification of the view confidence levels, it is practical to work in terms of confidence *intervals* for each of the views rather than calculating individual view variances.

The 95% confidence interval for the expressed view can be expressed as: $Q_k \pm 1.96\sigma_v$. Therefore, the size (S) of the confidence interval is given as $S = 2 \times 1.96\sigma_v$. If the investor specifies the size of the confidence interval, the variance of the expressed view can be obtained by $\sigma_v = \frac{S}{2\text{x}1.96}$. The variances of the individual views form the main diagonal of the view uncertainty matrix Ω. The greater the size of the 95% confidence interval, the lower the degree of confidence in the particular view. As a result of providing a confidence level, the variance of the expressed view can be obtained and input into the Black-Litterman model formulation.

Table 5: Implied Equilibrium Returns, Variances, Benchmark Weights and View Data for the Top 40 Assets on the FTSE/JSE ALSI

Asset	Implied Equilibrium Return	Asset Variance	Benchmark Weight	Views: View Expected Return	Views: View Variance	Views: View Score	Views: View Category
BHP BILLITON	0.097	0.206	0.129	0.357	0.053	1.261	E
SABMILLER	0.061	0.146	0.119	-0.512	0.062	-3.908	C
SASOL	0.084	0.192	0.063	1.063	0.057	5.090	G
MTN GROUP	0.068	0.213	0.063	0.798	0.057	3.422	F
RICHEMONT	0.078	0.197	0.059	1.809	0.064	8.797	H
STANDARD BK.	0.066	0.169	0.043	0.016	0.052	-0.297	D
NASPERS	0.081	0.225	0.042	0.162	0.060	0.362	E
ANGLO PLATINUM	0.120	0.276	0.039	1.716	0.060	5.784	G
FIRSTRAND	0.068	0.183	0.033	0.329	0.054	1.424	E
ANGLOGOLD ASHANTI	0.063	0.246	0.031	1.374	0.057	5.330	G
ABSA GROUP	0.057	0.178	0.028	1.553	0.048	8.404	H
OLD MUTUAL	0.077	0.200	0.026	-1.084	0.048	-5.794	B
IMPALA PLATINUM	0.117	0.270	0.026	0.701	0.056	2.163	E
GOLD FIELDS	0.062	0.273	0.021	-0.125	0.061	-0.686	D
NEDBANK GROUP	0.050	0.176	0.020	0.462	0.064	2.334	E
EXXARO RESOURCES	0.100	0.285	0.018	-1.594	0.049	-5.938	B
SHOPRITE	0.034	0.172	0.018	0.644	0.057	3.552	F
SANLAM	0.059	0.158	0.016	-1.471	0.055	-9.654	A
REMGRO	0.041	0.127	0.016	-0.651	0.047	-5.468	B
BIDVEST GROUP	0.052	0.148	0.014	1.278	0.053	8.290	H
TIGER BRANDS	0.045	0.139	0.012	0.790	0.050	5.369	G
STEINHOFF INTL.	0.082	0.215	0.012	0.225	0.061	0.667	E
ASPEN PHMCR.HDG.	0.039	0.198	0.012	-1.674	0.052	-8.633	A
RMB	0.064	0.177	0.011	-0.999	0.056	-6.009	B
HARMONY GOLD	0.081	0.350	0.010	-3.337	0.050	-9.760	A
AFN.RAINBOW MRLS.	0.100	0.263	0.010	2.556	0.058	9.332	H
MASSMART	0.042	0.193	0.009	1.326	0.052	6.644	G
WOOLWORTHS	0.053	0.182	0.009	-0.146	0.059	-1.094	D
TRUWORTHS INTL.	0.040	0.187	0.009	1.870	0.059	9.810	H
ASSORE	0.065	0.257	0.009	1.576	0.060	5.886	G
GROWTHPOINT PROPS.	0.028	0.144	0.009	-0.005	0.055	-0.228	D
CAPITAL SHOPCTS.GP.	0.046	0.178	0.008	-1.269	0.048	-7.388	B
AFRICAN BANK INVS.	0.070	0.227	0.008	-1.236	0.051	-5.755	B
IMPERIAL	0.068	0.212	0.008	1.471	0.063	6.622	G
ARCELORMITTAL	0.100	0.297	0.007	-0.948	0.050	-3.525	C
DISCOVERY	0.037	0.166	0.007	0.832	0.062	4.776	F
MMI HOLDINGS	0.052	0.180	0.007	-1.174	0.056	-6.803	B
FOSCHINI GROUP	0.057	0.194	0.007	0.954	0.065	4.630	F
LONMIN	0.111	0.301	0.007	-2.573	0.048	-8.926	A
MEDICLINIC INTERNATIONAL	0.027	0.126306	0.006	0.569	0.055	4.288	F

Table 5 above briefly describes the data used as well as the expressed views and implied equilibrium returns for each asset in the portfolio. The first three columns describe the data *before* the views have been expressed and the last four columns describe the data *after* the views have been expressed (the data given the views). The assets are listed in order of the benchmark weightings. It can be noted that BHP Billiton and SAB Miller occupy the greatest benchmark weightings of 12.9% and 11.9% respectively, whereas, in contrast, Mediclinic International exhibits the lowest benchmark weighting of 0.6%.

Given the implied equilibrium returns, variances and view expected returns, one can compare and contrast the differences that occur as a result of the expressed view scores. For example, considering the extreme view scores expressed for Lonmin (extreme negative score: -8.926) and Richmont (extreme positive score: 8.797), the extremely large view scores result in correspondingly extremely large negative and positive view expected returns of -257.3% and 180.9% respectively. Understandably, these annualised monthly returns seem unrealistic, however, a possible reason for an investor expressing such extreme views could be due to a prediction of an upcoming stock market shock whereby the respective assets are about to experience a severe drop (increase) in price. It also follows that if the investor is not satisfied with the given view expected returns, the respective view scores can be adjusted accordingly in order to result in more suitable view implied returns.

The following sections describe the above views and analyse the portfolio according to the three respective components (benchmark, views and constraints) and measure each of the respective contributions to the portfolio statistics. In particular, strong emphasis is expressed on how constraints distort the view scores in table 5 above. In addition, the improvements in the levels of view distortion are observed for marginally relaxing the long-only constraint.

6.2.2 Analysing Influential Components of the Black-Litterman Model

The aim of this section is to illustrate how the final Black-Litterman portfolio is composed of three respective components, namely the benchmark, view and constraint components. Each of these three components possesses an influence on the characteristics of the final optimal portfolio and it is therefore useful for the investor to be able to identify the sources of risk and return.

The aim of this section is to expose the relative contributions of each of these components to the respective statistics of the final portfolio. It must be noted that an attractive Black-Litterman portfolio is one for which the investment constraints are satisfied whereby there is minimal distortion of the expressed investor views.

It is important to note that, when analysing the contribution of constraints to the portfolio statistics, the Lagrange multipliers are constantly changing as the asset returns and views change. The Lagrange multipliers change in order for the portfolio to satisfy the constraints recommended by the investor. It is therefore important to keep in mind that, when analysing the respective contribution of constraints to the portfolio, the contribution of the constraints is not static and is likely to change in a non-linear manner depending on how binding each constraint is at a particular time. It must therefore be noted that the example below is a simple illustration of how constraints may influence the characteristics that drive the portfolio's risk and return statistics.

This section will therefore illustrate how each component contributes to the statistics of the portfolio. In addition, this section will analyse the impact of constraints on the optimal portfolio in terms of the reduction in expected utility and, more specifically, how each of the individual views are distorted. In particular, specific attention will be drawn to how the respective view scores, provided by the investor, change as a result of the constraints, therefore providing an indication of the magnitude of which constraints uncontrollably distort the optimal portfolio specified by the investor.

For analysing the respective components of the final optimal portfolio, the following portfolio is constructed:

$$Max\left\{w'E[R_{BL}]-\frac{\lambda_2}{2}w'\Sigma w\right\}$$

Subject to the following constraints:

$(w-w_b)'\Sigma(w-w_b)\leq 0.03^2$	3% Tracking Error Constraint
$w\geq 0$	Long-Only Constraint
$\bar{1}(w-w_b)=0$	Zero Net Active Position Constraint
$\frac{w'\Sigma w_b}{w_b{}'\Sigma w_b}=1$	Beta neutrality constraint (Beta equal to one)

Applying the methodology outlined in section 5.2.2 to the above portfolio optimisation, the relative contributions of the three components to the portfolio characteristics are illustrated below:

Figure 11: Bar Charts Representing the Benchmark, View and Constraint Component Proportions for the Optimal Portfolio Statistics

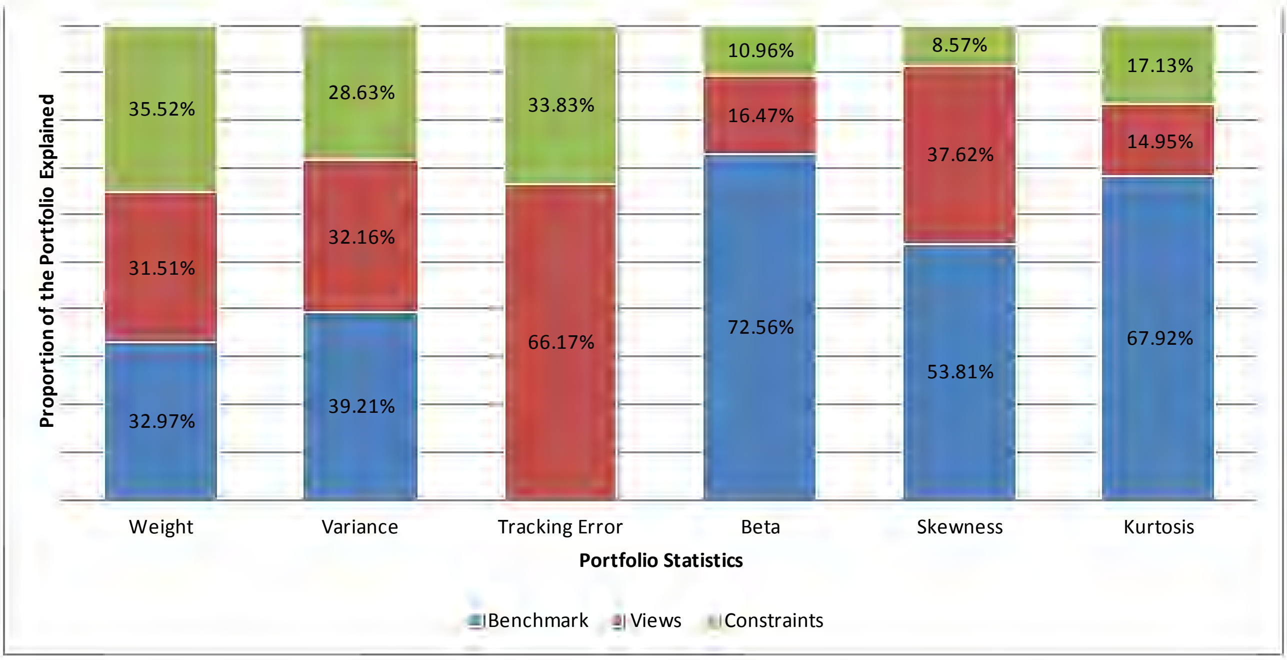

Each of the three components for the six portfolio statistics illustrated above, are calculated as the norm of the portion due to constraints, views or benchmark, divided by the sum of the individual component norms, therefore giving the percentage of the particular statistic explained by each respective component. As illustrated in table 5 and figure 11 above, one can see that constraints influence a significant percentage of the portfolio statistics[66]. In addition, it can be noted that when investors express views and associated confidence levels on particular assets, they are implicitly providing information on where they would prefer the active risk to be allocated. For example, if a strongly positive view (with a high level of confidence) for MTN is expressed, the investor is implicitly inferring that a significant proportion of the active risk budget must be allocated towards overweighting MTN in the portfolio. However, when constraints are enforced, the active weight allocated to MTN will be offset by the weights inferred by the constraints, thereby "redistributing" the active risk budget in a non-transparent manner towards other assets in the portfolio in order to satisfy the imposed constraints.

The most noteworthy observation is that, of the three respective components, the constraint component occupies over a third (35.52%) of the overall recommended portfolio weights. More specifically, of the entire portfolio, 32.97% is derived from the benchmark portfolio, 31.51% from the expressed views and 35.52% is as a result of the imposed constraints. Therefore, since a significant proportion of the overall portfolio weights are driven by constraints, once can deduce that an overwhelming degree of investor information contained in the benchmark and view portfolios have been severely distorted and will not be accurately represented in the final portfolio. In addition, it can be noted that since a significant proportion of the portfolio weights are explained by constraints, the relative effects of the constraints inevitably permeate into the other statistics of the optimal portfolio.

Observing the respective component distributions of the variance and tracking error, it can be inferred that approximately a third (33.83%) of the active risk budget and 28.63% of the total risk of the portfolio arises as a result of the expressed constraints. This is very disturbing as, in practise, a significant amount of funding and research is directed at forming the investment views as well as deciding where the active risk needs to be situated – which is then subjected to constraints and is distorted to the extent that approximately a third of the risk budget is decided by the constraints and not by the investor. As a result of the constraints, the investor has little control over where the active risk budget can be allocated.

Due to the non-linear relationship between portfolio weights and the computation of the portfolio beta, skewness and kurtosis, it is difficult to appropriately interpret the respective contributions of the three components. It can be noted above that the constraint portions occupy comparatively smaller proportions

[66] It must be noted that the ratio of the relative contributions of the benchmark and view components remain constant for both the constrained and unconstrained portfolios. Therefore, when constraints occupy a proportion of the respective portfolio characteristics, both the benchmark and view component proportions decrease in equal amount.

of 10.96%, 8.57% and 17.13% respectively. Due to the beta neutrality constraint, it is surprising that the constraints explain such a small portion of the overall portfolio beta. However, as noted above, this could be due to the non-linear estimation of the portfolio beta. The same holds true for the portfolio skewness and kurtosis. A plausible interpretation for the lower constraint proportions could be due to the constraint implied active weights being allocated to positions that exhibit significantly lower covariance, co-skewness and co-kurtosis parameters with the other assets in the portfolio.

In addition to the above, it can be noted that as a result of the constraints, a 30.15% reduction in the investor's utility was observed. Section 6.3.5 further describes and illustrates how the utility of the investor can be improved as a result of marginally relaxing the long-only constraint.

As illustrated above, it is apparent that constraints influence a significant proportion of the optimal portfolio statistics. The following sections describe in further detail how the constraints influence the optimal portfolio as well as the expression of the investor views.

6.2.3 Analysing the Views and Constraints of the Black-Litterman Model

The aim of this section is to analyse the views expressed by the investor as well as how the enforcement of constraints distorts the expressed views. In particular, the impact of constraints on the portfolio can be measured by the degree to which the expressed views are distorted. In addition, the originally expressed views and constraint distorted views will be analysed and compared according to their respective contributions to the portfolio tracking error.

Analysing the Impacts of Constraints on the Views of the Black-Litterman Model

As described in section 5.2.3, enforcing constraints on the Black-Litterman portfolio severely impacts the views expressed by the investor. Assuming that an investor has a particular benchmark portfolio and then alters the benchmark portfolio by adding views and constraints, the net effect is given by the interaction of the two respective contributions. The impact that the constraints have on the views is observed by how the final constrained portfolio differs from the unconstrained portfolio.

Following the methodology described in section 5.2.3, one[67] of the methods to be used to describe the level of view distortion is the percentage change of the *raw*[68] constraint-implied views from the original *raw* views, $K = \frac{Q_{con} - Q}{Q}$. The values of the elements in K can be interpreted as follows:

[67] In addition to using the percentage change in the view scores, the difference between the original and constraint-implied view scores as well as the number of view category changes as a result of the constraints will be analysed.
[68] These raw views refer to the extent to which the view vector Q differs from the implied equilibrium returns Π. The raw views are thus given by $Q - \Pi$.

- $K = 0$: Values equal to zero indicate that no view distortion has occurred as a result of the constraints. Therefore, the constraints do not impede the expression of the views. Values of K close to zero are attractive as a rational investor would prefer as little view distortion as possible.
- $K > 0$: Values greater than zero percent indicate views that have been amplified or have become more extreme as a result of the constraints. In particular, negative views become increasingly negative and positive views become larger positive views.
- $-1 < K < 0$: Values between negative one and zero indicate view suppression whereby a view has been made less extreme (smaller) in magnitude relative to the original view. More specifically, the view score is closer to zero (increasingly neutral) relative to the original view.
- $K < -1$: Values less than negative one indicate that the view has changed sign. For example, original views that are positive (negative) on a specific asset are now negative (positive), as implied by the constraints. Therefore, values of $K < -1$ indicate severe view distortion as the constraint implied view directly contrasts that of the original view.

Following from the previous example in section 6.2.2, views were expressed on the assets in the portfolio and the same constraints were enforced (long only, zero net active position and beta neutrality constraints). In addition, the views were subjected to the same 3% tracking error.

Table 6 and table 7 on the following pages illustrate how the respective views have changed for each particular view category. Table 6 counts the number of views that have changed from the original views to the constrained portfolio implied views. In interpreting table 6, the original views are expressed on the horizontal rows, with the total count for each original view category listed in the final column. In contrast, the number of constraint-implied views is listed as columns, with the totals for each constrained view category listed in the final row. The numbers in the centre of the table indicate the number of view changes from each original view category to each constrained view category.

Table 6: Summary of the Number of View Score Changes that Occurred for Each Category

		Constraint-Implied Views								
		Extreme Negative (A)	Strong Negative (B)	Weak Negative (C)	Mild Negative (D)	Mild Positive (E)	Weak Positive (F)	Strong Positive (G)	Extreme Positive (H)	**Total Original**
Original Views	Extreme Negative (A)				**4**					**4**
	Strong Negative (B)				**7**					**7**
	Weak Negative (C)	**1**			**1**					**2**
	Mild Negative (D)			**1**	**3**					**4**
	Mild Positive (E)	**1**			**5**					**6**
	Weak Positive (F)				**1**	**1**	**3**			**5**
	Strong Positive (G)			**1**	**2**		**1**	**2**	**1**	**7**
	Extreme Positive (H)			**1**		**1**	**2**	**1**		**5**
Total Constrained		**2**	**0**	**3**	**23**	**2**	**6**	**3**	**1**	**40**

Extreme Negative	Strong Negative	Weak Negative	Mild Negative	Mild Positive	Weak Positive	Strong Positive	Extreme Positive
(A)	(B)	(C)	(D)	(E)	(F)	(G)	(H)
$(-10 \leq \eta < -7.5)$	$(-7.5 \leq \eta < -5)$	$(-5 \leq \eta < -2.5)$	$(-2.5 \leq \eta < 0)$	$(0 < \eta \leq 2.5)$	$(2.5 < \eta \leq 5)$	$(5 < \eta \leq 7.5)$	$(7.5 < \eta \leq 10)$

Figure 12: A Graphical Illustration of the Components of the Constraint-Implied View Scores in Comparison to the Original View Scores

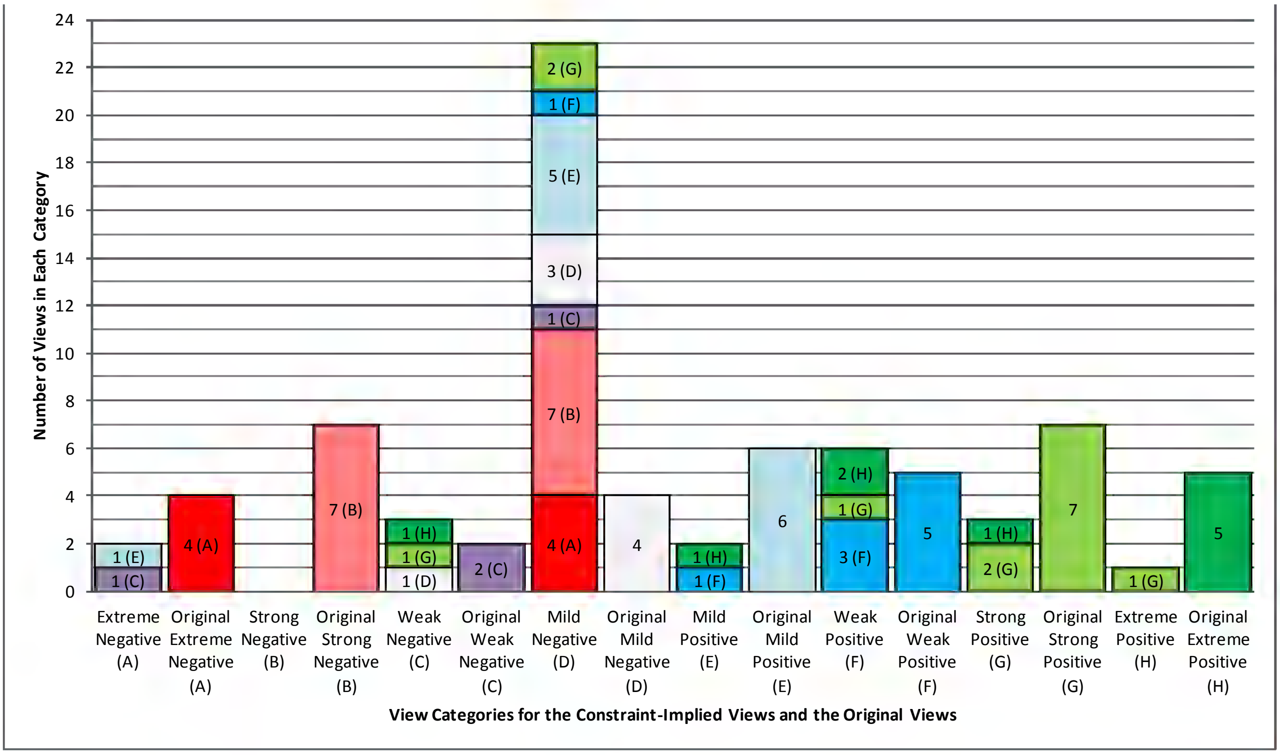

Figure 12 above illustrates the information contained in table 6 and compares the original view distribution with the constraint-implied view distribution. For each constraint-implied view bar, there is a corresponding original view bar illustrating the number of views in the particular view category. Each stacked constraint implied view bar is composed of the number of views originating from a particular original view category. For example, there are two extreme negative views implied by the constraints, of which, one was originally a weak negative view and the other view was originally a mild positive view. There were originally four extreme negative views, as originally expressed by the investor. Further, there were seven originally expressed strong negative views and there are currently zero constraint-implied strong negative views.

Observing table 6 and figure 12 above, the following can be noted:

- All of the four *original* **extreme negative views**, $\left({}^{4}/_{4}, 100\%\right)$ and all the seven *original* **strong negative views** $\left({}^{7}/_{7}, 100\%\right)$, were reduced to mild negative views. Therefore, the one can expect the significantly large active asset under-weightings (as implied by the extreme and strong negative views) to be reduced to smaller asset underweighting relative to the benchmark portfolio.
- Of the two original **weak negative views**, one view was reduced to a mild negative view $\left({}^{1}/_{2}, 50\%\right)$, therefore reducing the active under weights taken on a particular asset. The other view was amplified to become an extreme negative view $\left({}^{1}/_{2}, 50\%\right)$, therefore acquiring an additional active under-weighting as a result of the imposed constraints.
- Examining the original **mild positive views**, a disturbing observation can be noted: all of the positive views have changed sign and have become negative. The majority $\left({}^{5}/_{6}, 83.33\%\right)$ have changed sign to become mild negative views and one $\left({}^{1}/_{6}, 16.67\%\right)$ has become an extreme negative view. Therefore, while the investor originally intended for these six views to be positive in order to achieve positive active asset over-weightings, the imposed constraints distorted the views to such an extent that the views became negative, thereby resulting in asset under-weightings and even an extreme asset underweighting.
- The original **weak positive views** present a much less disturbing picture whereby the majority $\left({}^{3}/_{5}, 60\%\right)$ remain in the same view category of being weak positive views. One view $\left({}^{1}/_{5}, 20\%\right)$ was reduced to a mild positive view and the only alarming case for concern was for one $\left({}^{1}/_{5}, 20\%\right)$ view which changed sign to become a mild negative view.
- Examining the original **strong positive views**, it can be noted that two $\left({}^{2}/_{7}, 28.57\%\right)$ views remained in the same view category with one $\left({}^{1}/_{7}, 14.28\%\right)$ view being amplified to become an extreme positive view, one $\left({}^{1}/_{7}, 14.28\%\right)$ being reduced to a weak positive view and another alarming situation where three views have changed sign where two $\left({}^{2}/_{7}, 28.57\%\right)$ became mild negative and one $\left({}^{1}/_{7}, 14.28\%\right)$ became a weak negative constraint-implied view.
- It can be noted that all the originally **extreme positive views** were reduced, with one $\left({}^{1}/_{5}, 20\%\right)$ becoming a strong positive view, two $\left({}^{2}/_{5}, 20\%\right)$becoming weak positive views, one $\left({}^{1}/_{5}, 20\%\right)$ becoming a mild positive view and only one $\left({}^{1}/_{5}, 20\%\right)$ alarming case whereby the view changed sign to become a mild negative constraint-implied view.

Table 7: Summary of the Frequency of the Number of View Category Changes

	Zero Category Changes	One Category Change	Two Category Change	Three Category Change	Four Category Change	Five Category Change	Six Category Change	Seven Category Change
Frequency of Observed Changes	8	11	11	7	2	1	0	0

It must be noted that when interpreting the view distortion percentages, that each value is interpreted *relative to the original view*. Therefore, the purpose of the view distortion percentage figure is to help identify which views have been reduced, amplified and most importantly, which views have changed sign. In order to obtain a more comprehensible picture of the degree of view distortion, the view difference and the number of view category changes need to be observed. The table above illustrates the magnitude of the view distortions in terms of the number of view category changes. There are eight individual view categories, ranging from an extreme negative view to an extreme positive view. Therefore, for example, a view that has changed by 7 categories is a view that has changed from being extremely negative to extremely positive and has consequently been severely distorted as the constraint implied view is the complete opposite of the originally intended investor view. It can therefore be noted that the greater the number of category changes from the original view, the greater the level of distortion.

Examining the table 7 above, it can be observed that there were 8 views that did not change to a different view category, therefore indicating very little view distortion. Of the remaining 32 views, 11 views changed by one view category and a further 11 changed by two view categories; indicating a significant level of distortion. Of significant concern are the remaining views that changed by three or more view categories; 7 views changed by three categories, 2 views by four categories and one view that was completely distorted, therefore resulting in five view category changes. It is therefore evident that the majority of the original views were severely distorted, indicating constraint-implied views that are very different (even completely opposite) to what the investor originally intended.

View Score Changes and Distortions

Examining the views in more detail, table 8 below lists the 40 assets and the associated views. Each original view is contrasted against the corresponding constrained-implied view. The original views (O) as well as the constraint-implied views (C) are summarised for each asset. The view scores, the associated view expected returns and the resulting view type for the original and constrained portfolios are summarised next to each other for an easy comparison. For every pair of original and constraint-implied views, a view distortion percentage is given, which determines the degree of view distortion that has occurred. The last two columns of table 8 provides the actual difference in view scores[69] as well as an indication of whether the

[69] The difference in view scores is simply given by the difference of original view score and the constraint-implied view score; $\eta_{original} - \eta_{constraint}$

particular view has changed to a different view category as a result of the imposed constraints. In addition, the number of view category changes is included in parentheses.

Comparing the original view scores with the constraint-implied view scores as well as examining the view distortions in the table below, amongst others, it is apparent that the views on BHP Billiton, SASOL, Anglo Platinum and Bidvest experienced radically harsh view distortions, whereby the views have been ***reversed***, from a positive view to a negative view. In particular, the view on BHP Billiton has been reversed from a mild positive view (1.261) to an extremely negative view (-8.4), which is a -776.09% change from the originally expressed view, therefore resulting in an annualised monthly expected return view of -13.6% instead of the originally expressed expected return of 35.7%. In addition, the SASOL, Anglo Platinum and Anglo Ashanti view scores changed sign from strongly positive original views to weakly negative and mild negative constraint-implied views, with corresponding distortion percentages of -187.39%, -138.30% and -134.77% respectively. As stated earlier, these distortions are extremely influential, as the constraints have changed the signs of the originally expressed view scores, resulting in negative view scores that differ by more than 8 view score points.

In addition to the views on assets that have been reversed as described above, it can be noted that the views on SAB Miller, SANLAM, Asphen Pharmaceuticals, Harmony Gold and Lonmin possessed views that were heavily distorted, resulting in constraint-implied views that represent severe view suppression and expansion. In particular, the following can be noted:

- The view score on SAB Miller was expanded from a weakly negative view of -3.9 to an overwhelmingly extreme negative view score of –11.59, which is beyond the specified bounds of -10 and 10 for the originally expressed view scores.
- The remaining assets discussed above, experienced significant view score suppression, that have resulted in strongly negative views being suppressed into constraint-implied views that are mild negative and nearing neutral views (view scores close to zero). In all the mentioned cases below, it can be noted that the views were extreme to the extent that the view annualised monthly return expectations were highly unlikely and were consequently reduced to less extreme views as a result of the constraints. The following cases of view suppression were noted:
    - Aspen Pharmaceuticals: View score change from -8.633 to -0.292 (-96.62% distortion). The associated annualised monthly view return expectation of -167.4% was reduced to -0.2%.
    - Harmony Gold: View score change from -9.760 to -0.426 (-95.64% distortion). The highly unlikely view annualised monthly expected return of -333.7% was reduced to a less extreme view of -0.6% annualised monthly expected return.
    - Lonmin: View score change from -8.926 to -0.427 (-95.22%), which resulted in an annualised monthly expected return view of -257.3% to be reduced (almost neutralised) to -0.1%.

In all three cases above, the imposed constraints demonstrated the effect of reducing highly risky and extreme views to views that were progressively more neutral in nature and therefore less risky. However, one could argue that such a view reduction would be undesirable if the investor *intended* for the views to be extreme in magnitude. It therefore follows that in *some*[70] cases, constraints have the effect of uncontrollably reducing risky active positions taken on particular assets in the portfolio.

It can be concluded that, as a result of carefully examining tables 6, 7 and 8 as well as figure 12, one can obtain a clear picture of how each of the individual views were distorted as well as the overall summary statistics of how each particular constraint-implied view category is composed of the originally categorised views. It must be noted that thus far, there has been no detailed discussion of how the constraints have influenced the active risk of the portfolio as well as the overall portfolio characteristics. The following section discusses how the constraints have influenced the percentage contribution of each asset to the overall portfolio tracking error.

[70] It can be noted that, while the constraints reduced four of the originally extreme negative views to mild negative views, thereby avoiding excessive risk taking, the constraints also forced a mild positive view (BHP BILLITON) and a weak negative view (SAB Miller) to become extreme negative views, therefore unintentionally increasing active positions in the larger market capitalisation weighted assets.

Table 8: Comparison of the Original Views and Constraint-Implied Views as well as the Associated View Distortions

Asset	View Score (O)	View Score (C)	View Expected Return (O)	View Expected Return (C)	View Category (O)	View Category (C)	View Distortion: Distortion Percentage	View Distortion: View Score Difference	View Distortion: View Change
BHP BILLITON	1.261	-8.400	0.357	-0.136	E	A	-766.10%	9.661	YES (4)
SABMILLER	-3.908	-11.59	-0.512	-0.137	C	A	196.76%	7.690	YES (2)
SASOL	5.090	-4.449	1.063	-0.064	G	C	-187.40%	9.539	YES (4)
MTN GROUP	3.422	3.548	0.798	0.069	F	F	3.68%	-0.126	NO (0)
RICHEMONT	8.797	7.494	1.809	0.129	H	G	-14.82%	1.304	YES (1)
STANDARD BK.GP.	-0.297	-3.375	0.016	-0.042	D	C	1036.36%	3.078	YES (1)
NASPERS	0.362	-0.335	0.162	0.001	E	D	-192.54%	0.697	YES (1)
ANGLO PLATINUM	5.784	-2.216	1.716	-0.041	G	D	-138.30%	8.000	YES (3)
FIRSTRAND	1.424	-2.313	0.329	-0.030	E	D	-262.49%	3.737	YES (1)
ANGLOGOLD ASHANTI	5.330	-1.853	1.374	-0.033	G	D	-134.77%	7.183	YES (3)
ABSA GROUP	8.404	2.178	1.553	0.037	H	E	-74.09%	6.226	YES (3)
OLD MUTUAL	-5.794	-1.774	-1.084	-0.023	B	D	-69.39%	-4.021	YES (2)
IMPALA PLATINUM	2.163	-1.103	0.701	-0.015	E	D	-151.01%	3.266	YES (1)
GOLD FIELDS	-0.686	-0.976	-0.125	-0.017	D	D	42.38%	0.291	NO (0)
NEDBANK GROUP	2.334	-1.592	0.462	-0.019	E	D	-168.19%	3.926	YES (1)
EXXARO RESOURCES	-5.938	-0.912	-1.594	-0.013	B	D	-84.65%	-5.027	YES (2)
SHOPRITE	3.552	-1.299	0.644	-0.016	F	D	-136.56%	4.851	YES (2)
SANLAM	-9.654	-1.541	-1.471	-0.015	A	D	-84.04%	-8.113	YES (3)
REMGRO	-5.468	-2.367	-0.651	-0.022	B	D	-56.71%	-3.101	YES (2)
BIDVEST GROUP	8.290	-4.663	1.278	-0.053	H	C	-156.25%	12.953	YES (5)
TIGER BRANDS	5.369	2.631	0.790	0.034	G	F	-51.00%	2.738	YES (1)
STEINHOFF INTL.	0.667	-0.902	0.225	-0.009	E	D	-235.26%	1.568	YES (1)
ASPEN PHMCR.HDG.	-8.633	-0.292	-1.674	-0.002	A	D	-96.62%	-8.341	YES (3)
RMB	-6.009	-0.760	-0.999	-0.006	B	D	-87.35%	-5.249	YES (2)
HARMONY GOLD MNG.	-9.760	-0.426	-3.337	-0.006	A	D	-95.64%	-9.334	YES (3)
AFN.RAINBOW MRLS.	9.332	4.370	2.556	0.104	H	F	-53.17%	4.962	YES (2)
MASSMART	6.644	8.511	1.326	0.141	G	H	28.09%	-1.867	YES (1)
WOOLWORTHS HDG.	-1.094	-0.368	-0.146	-0.001	D	D	-66.41%	-0.727	NO (0)
TRUWORTHS INTL.	9.810	4.171	1.870	0.068	H	F	-57.49%	5.640	YES (2)
ASSORE	5.886	6.338	1.576	0.141	G	G	7.69%	-0.453	NO (0)
GROWTHPOINT PROPS.	-0.228	-0.279	-0.005	-0.001	D	D	22.37%	0.051	NO (0)
CAPITAL SHOPCTS.GP.	-7.388	-0.961	-1.269	-0.011	B	D	-86.99%	-6.427	YES (2)
AFRICAN BANK INVS.	-5.755	-0.417	-1.236	-0.002	B	D	-92.75%	-5.338	YES (2)
IMPERIAL	6.622	5.986	1.471	0.111	G	G	-9.60%	0.636	NO (0)
ARCELORMITTAL SA.	-3.525	-0.255	-0.948	0.002	C	D	-92.78%	-3.270	YES (1)
DISCOVERY	4.776	2.742	0.832	0.041	F	F	-42.59%	2.034	NO (0)
MMI HOLDINGS	-6.803	-0.326	-1.174	-0.001	B	D	-95.21%	-6.477	YES (2)
THE FOSCHINI GROUP	4.630	2.596	0.954	0.047	F	F	-43.92%	2.034	NO (0)
LONMIN	-8.926	-0.427	-2.573	-0.001	A	D	-95.22%	-8.499	YES (3)
MEDICLINIC INTERNATIONAL	4.288	0.194	0.569	0.004	F	E	-95.49%	4.095	YES (1)

6.2.4 Impacts of Constraints on the View Percentage Contributions to Tracking Error

The previous section discussed how the constraints impacted the views of the Black-Litterman Model in terms of the level of view distortion. The aim of this section is to analyse *how* the constraints have impacted the optimal portfolio and, more specifically, how the constraints have altered the active risk contribution of each the respective views as a result of distorting the original views of the investor. In particular, the percentage contribution of the original views and the constraint-implied views to the portfolio tracking error[71] ($PCVTE$ and $PCVTE^{con}$) will be analysed and compared. In this way, the investor is then able to analyse each view from a risk budgeting perspective in order to determine the following as a result of the portfolio constraints:

- Whether the current contributions to tracking error are in line with the original objectives of the investor.
- The most influential views can be identified in order to give the investor a clearer picture of the respective components of the active portfolio.
- The respective proportions of tracking error explained by the constraint-implied views can be compared to the original portfolio and, more specifically, analysed to determine the most influential views whereby the active risk is concentrated.

Table 9 below lists the assets according to their respective view contributions to the portfolio tracking error[72], in order from the largest contribution to the least contribution. The PCVTE provides an indication of the magnitude for which each particular view contributes to the active risk of the portfolio. It therefore allows the investor to reassess the magnitude and convictions of the particular views and re-calibrate the views in order to result in a portfolio whereby the contribution of the individual views to active risk (tracking error) is in line with the respective convictions of those views. Negative PCVTE values indicate that the given views decrease the active risk of the portfolio where as positive values indicate a positive contribution to the active risk. Table 9 below compares the originally (O) expressed views to the constraint-implied (C) views according to the respective contributions to tracking error.

[71] The PCVTE illustrates how each respective view contributes to the active risk of the portfolio. Views that possess large PCVTE values are views that result in large active weight positions, thereby resulting in significant active risk being acquired. Views that possess negative PCVTE values are views that contradict the other views expressed on the portfolio. These particular views therefore reduce the active risk of the portfolio. The fact that assets are influenced by views expressed on other highly correlated assets is an attractive quality of the Black-Litterman Model as it guards against excessive risk taking as a result of inconsistent views.

[72] The risk aversion coefficient was calibrated in order to ensure a tracking error of 3%.

Table 9: Comparison of the Original View and Constraint-Implied View Tracking Error Contributions

Asset	View Score		PCVTE		Proportion of TE Explained			
					Original View		Constrained View	
	(O)	(C)	(O)	(C)	(O)	Rank	(C)	Rank
HARMONY GOLD	-9.760	-0.426	3335.47%	172.31%	**29.12%**	1	1.61%	16
LONMIN	-8.926	-0.427	1125.99%	87.61%	**9.83%**	2	0.82%	21
AFN.RAINBOW MRLS.	9.332	4.370	728.27	204.15%	6.36%	3	1.9%	15
ASSORE	5.886	6.338	719%	1357.6%	6.27%	4	**12.66%**	2
TRUWORTHS INTL.	9.810	4.171	573.42%	562.56%	5.1%	5	5.25%	6
EXXARO RESOURCES	-5.938	-0.912	496.38%	171.14%	4.33%	6	1.59%	17
MASSMART	6.644	8.511	473.65%	1502.24%	4.13%	7	**14.014%**	1
IMPERIAL	6.622	5.986	457.92%	699.46%	3.99%	8	6.52%	5
RICHEMONT SECS.	8.797	7.494	392.4%	362.71%	3.43%	9	3.38%	11
ARCELORMITTAL SA.	-3.525	-0.255	323.59%	45.22%	2.83%	10	0.42%	25
ABSA GROUP	8.404	2.178	280.3%	89.19%	2.45%	11	0.83%	20
THE FOSCHINI GROUP	4.630	2.596	276.62%	291.11%	2.41%	12	2.71%	12
BIDVEST GROUP	8.290	-4.663	243.17%	-144.36%	2.12%	13	1.34%	19
ASPEN PHMCR.HDG.	-8.633	-0.292	239.67%	2.954%	2.09%	14	0.03%	37
MTN GROUP	3.422	3.548	223.56%	414.58%	1.95%	15	3.86%	9
ANGLOGOLD ASHANTI	5.330	-1.853	**-179.07%**	380.41%	1.56%	16	3.55%	10
DISCOVERY	4.776	2.742	176.21%	157.69%	1.54%	17	1.47%	18
CAPITAL SHOPCTS.GP.	-7.388	-0.961	149.87%	43.38%	1.31%	18	0.41%	26
TIGER BRANDS	5.369	2.631	117.76%	67.03%	1.03%	19	0.63%	23
SHOPRITE	3.552	-1.299	116.36%	-81.18%	1.02%	20	0.76%	22
AFRICAN BANK INVS.	-5.755	-0.417	-92.32%	-12.47%	0.81%	21	0.12%	32
SANLAM	-9.654	-1.541	-92.078%	-19.36%	0.8%	22	0.18%	29
RMB	-6.009	-0.760	-87.09%	-0.416%	0.76%	23	0.004%	40
GOLD FIELDS	-0.686	-0.976	74.98%	269.29%	0.65%	24	2.51%	13
IMPALA PLATINUM	2.163	-1.103	-68.18%	269.15%	0.59%	25	2.51%	14
NEDBANK GROUP	2.334	-1.592	62.16%	-50%	0.54%	26	0.47%	24
ANGLO PLATINUM	5.784	-2.216	-53.74%	556.17%	0.47%	27	5.19%	7
OLD MUTUAL	-5.794	-1.774	-47.76%	35.23%	0.42%	28	0.33%	27
MEDICLINIC INTERNATIONAL	4.288	0.194	40.33%	2.03%	0.35%	29	0.02%	39
REMGRO	-5.468	-2.367	-34.68%	3.24%	0.3%	30	0.03%	36
SASOL	5.090	-4.449	-30.37%	543.74%	0.26%	31	5.07%	8
SABMILLER	-3.908	-11.59	30.06%	813.67%	0.26%	32	7.59%	4
FIRSTRAND	1.424	-2.313	29.93%	12.51%	0.26%	33	0.12%	31
WOOLWORTHS HDG.	-1.094	-0.368	-22.05%	-12.04%	0.19%	34	0.11%	33
NASPERS	0.362	-0.335	20.4%	-21.25%	0.18%	35	0.198%	28
STEINHOFF INTL.	0.667	-0.902	19.65%	-5.39%	0.17%	36	0.05%	35
MMI HOLDINGS	-6.803	-0.326	-6.66%	-2.75%	0.058%	37	0.026%	38
STANDARD BK.GP.	-0.297	-3.375	-6.29%	13.86%	0.055%	38	0.13%	30
GROWTHPOINT	-0.228	-0.279	-5.3%	-10.77%	0.046%	39	0.1%	34
BHP BILLITON	1.261	-8.400	-1.57%	1229.82%	0.014%	40	**11.47%**	3

As illustrated in table 9 above, with respect to the *originally* expressed views, the following can be noted:

- The originally expressed view score for Harmony Gold (-9.760) contributes the greatest percentage to the portfolio tracking error, whereby the active contribution to tracking error is 3335.47% that of the overall portfolio tracking error. Analysed in isolation, the original view on Harmony Gold would result in significantly large active weight allocations, therefore resulting in a large observed percentage contribution to active risk. Observing the last section (Proportion of TE Explained) of table 9, the view on Harmony Gold explains the greatest and significantly dominating proportion (29.12%) of the overall portfolio tracking error. Therefore, the conviction of the view on Harmony Gold can be interpreted as the boldest and most controlling view influencing the portfolio's active risk.

- The originally expressed view score on AngloGold Ashanti (5.330) contributes a *negative* percentage value of -179.07% to the overall tracking error. Therefore, this indicates that the view score expressed on AngloGold Ashanti has the net effect of reducing the overall portfolio tracking error. Given that this particular view effectively reduces the portfolio tracking error, it can be identified as a view that contrasts other views on the portfolio expressed on highly correlated assets[73]. Given that the view on AngloGold Ashanti is identified as being inconsistent to the other views expressed on the portfolio, the investor therefore has the opportunity to review and calibrate the originally expressed view.
- In addition to the view on AngloGold Ashanti, the originally expressed views on African Bank, Sanlam and RMB resulted in negative percentage contributions to the portfolio tracking error. While these three views were ranked as the 21st, 22nd and 23rd most influential views respectively, they resulted in a decrease in the overall portfolio tracking error, therefore offsetting the positive risk contributions of the other remaining views on assets in the portfolio.

Observing and comparing the proportion of tracking error explained by the original and constraint-implied views, the following changes to the respective risk contributions of the views as a result of imposing constraints can be observed:

- Of the originally expressed views, 14 of the views decreased the portfolio tracking error (negative contribution) and consequently, 26 views increased the tracking error (positive contribution). In comparison to the originally expressed views, there were 13 risk contribution sign changes resulting in the constraint-implied views where 11 of the views decreased the portfolio tracking error and 29 positively contributed to the portfolio tracking error. This is of concern as 32.5% of the views in the portfolio have changed to such an extent that their respective percentage contributions to tracking error have changed sign, therefore effectively changing the risk dynamics of the portfolio. These changes to the portfolio can be significantly different to the risk objectives of the investor.
- The proportion of tracking error explained by the view on Harmony Gold, which originally contributed the most to the portfolio tracking error, with a corresponding rank number of 1, has decreased significantly from explaining 29.12% to only 1.61% (rank number of 16) of the overall portfolio tracking error. Therefore, in this case, the originally most influential view has been reduced to the extent that it contributes significantly less to the constrained portfolio's active risk.

[73] It is useful to note that when views of the Black-Litterman model are expressed on assets in the portfolio, the correlations between the assets are taken into account. Therefore, if two assets are highly correlated, views on either asset will influence both of the assets. Therefore, a view that reduces the portfolio tracking error, must contradict the views on other highly correlated assets.

- In addition, the second highest ranked original view on Lonmin decreased its tracking error proportion of 9.83% to 0.82%. This particular change resulted in the proportion of tracking error explained by the view on Lonmin decreasing to become the 21st most influential view on the portfolio tracking error, thereby significantly reducing the magnitude of the active weights resulting from the originally expressed view.
- In contrast to the above, the proportion of the tracking error explained by the original view on Massmart increased, from 4.13% (by more than three times) to explain 14.014% of the portfolio tracking error, therefore resulting in a rank change from 7th to 1st. In addition, the proportion of the tracking error explained by the original view on Assore increased from 6.27% to 12.66%, making it the second most influential constraint-implied view. For both cases, the PCVTE increased significantly, therefore indicating that as a result of the imposed constraints, views on assets that originally resulted in relatively small active risk contributions, have significantly increased to an extent that may not necessarily be in line with the objectives of the investor.
- Noteworthy changes in the PCVTE from the originally expressed views to the constraint-implied views can be observed for the views expressed on AngloGold Ashanti and Anglo Platinum. The PCVTE for AngloGold Ashanti changed from -179.07% to 380.41% and for Anglo Platinum, the PCVTE changed from -53.74% to 556.17%. For these two particular cases, the risk contributions changed from decreasing the portfolio active risk to effectively increasing the tracking error. These particular changes can be noted as a complete reverse in the risk contribution of the respective views as a result of the expressed constraints.

As described by the few examples above, the effect of imposing constraints on a portfolio significantly alters the risk contributions of the respective views. While originally extreme views that influenced a significant proportion of the portfolio tracking error existed, imposing constraints had the uncontrollable effect of reducing the PCVTE for views that dominated the active risk of the portfolio. Another disturbing aspect noted above is the fact that the imposed constraints not only changed the respective tracking error contributions of the views but effectively *reversed* the active risk contributions of 13 of the views on the assets in the portfolio. It can therefore be noted that as a result of the distorting effect of constraints, not only have the views themselves been distorted, but the corresponding risk contributions have also been distorted in an uncontrollable manner.

Portfolio View Risk Concentrations

In addition to analysing the individual risk contributions of the respective views, it is helpful to determine how concentrated the portfolio active risk is with respect to the number of views that effectively control the portfolio tracking error. For example, a portfolio tracking error that is highly concentrated will most likely be influenced by a relatively small number of views. In contrast, a scenario whereby each view equally contributes to the portfolio tracking error will result in a low level of risk concentration.

In order to measure the risk concentration of the views in the portfolio tracking error, the Herfindahl-Hirschman Index (HHI) measure will be used. Using the HHI measure for concentration, the concentration of the views in the portfolio tracking error can be computed. Table 10 below illustrates and compares the concentration of the portfolio tracking error for the original views as well as for the constraint-implied views. For the original views as well as the constraint implied-views, the cumulative proportions of tracking error explained as well as the cumulative HHI for the tracking error is given. For each column, the 40 views are analysed in percentiles in order of decreasing proportions. For example, the 25th percentile consists of the top 10 views $(^{10}/_{40})$ that contribute the most to the portfolio tracking error. The final row, 100th Percentile, gives the complete proportions and the overall HHI measures for both the original and constraint-implied view cases. Using this method of calculating the cumulative proportions as well as splitting the HHI into the respective percentiles, one obtains an improved idea of how the overall portfolio tracking error is influenced.

As illustrated in table 10 below, the HHI for the originally expressed views was computed as 0.116 and the corresponding constraint-implied HHI was computed as 0.0744. Therefore, comparing these two measures, it follows that the constraint-implied tracking error exhibits a lower level of view concentration, therefore indicating that the constraints had the effect of diversifying the active risk amongst the views expressed on the portfolio. However, despite the increased diversity (in terms of the risk budget), when analysing the respective percentiles[74] of the views, it becomes evident that there is no significant improvement in the diversification of active risk amongst the views expressed on the portfolio. More specifically, while there is a significant improvement of the constraint-implied HHI for the 25th percentile (0.0696 versus 0.112), it must be noted that the top 10 constraint-implied views still account for 75.19% (compared to 75.31%) of the portfolio tracking error. This suggests that while there is more risk diversification, the majority (75.19%) of the active risk is still accounted for by 25% of the views (i.e. 10 of the 40 views). Observing the remaining percentiles, it becomes evident that the constraint-implied tracking error becomes less diversified compared to the original view case, as seen by the 50th and 75th percentiles whereby the cumulative proportions of the constraint-implied views account for marginally more (95.07% versus 92.79% and 99.4% versus 98.49%) of the total tracking error compared to the original view case. The final row of table 10 provides the cumulative proportion of tracking error explained as well as the cumulative HHI *with the most influential view removed.* The results show that with view on Harmony Gold (for the original case) removed and the view on Massmart (for the constraint-implied case) removed, the tracking error view diversification changes significantly. In particular, the total HHI for the constraint-implied view case is significantly poorer than the original view case (0.0547 versus 0.031). This therefore suggests that, with the exception of the most influential view, the remaining constraint-implied view case results in a view portfolio that is

[74] The views were ordered in descending order in terms of their respective contributions to the portfolio tracking error. Therefore, the 25th percentile represents the top 10 views (out of 40) that contributed the most to portfolio tracking error.

significantly more concentrated whereby active risk is concentrated in fewer views on assets in the portfolio.

Table 10: Summary of the View Tracking Error Concentrations for the Percentiles of the Original and Constraint-Implied Views

	Cumulative Proportion of Tracking Error Explained		Cumulative HHI	
	Original Views	Constraint-Implied Views	Original Views	Constraint-Implied Views
25th Percentile (Top 10 Views)	75.31%	75.19%	0.112	0.0696
50th Percentile (Top 20 Views)	92.79%	95.07%	0.116	0.0741
75th Percentile (Top 30 Views)	98.49%	99.4%	0.116	0.0743
100th Percentile (All 40 Views)	100%	100%	**0.116**	**0.0744**
Most Influential View Removed	70.88%	85.99%	**0.031**	**0.0547**

6.2.5 Summary and Conclusions

The aim of this particular example was to analyse the components of the Black-Litterman model when applied to an equity portfolio of the FTSE/JSE top 40 stocks. A set of random views (in the form of view scores) were expressed on each asset and their respective contributions to the final portfolio weights were analysed. More specifically, the impact of constraints on the expressed views were analysed in terms of the change in the magnitude of the originally expressed view (measured as a view distortion percentage) as well as the change in the percentage contribution to tracking error.

The results indicated that for the particular portfolio of 40 assets and associated views and constraints, the constraints typically accounted for over a third (35.52%) of the portfolio weight and tracking error (33.83%). This disturbing result is an indication that the investor has lost over a third of the control of the portfolio. In addition, the impacts of constraints *on a view level* were analysed whereby 32 out of the 40 views changed to a different view category[75]. In particular, 11 views were noted to have changed sign (from being positive on an asset to becoming negative and vice versa), with many more views being amplified and suppressed to the extent that there was almost no resemblance between the originally expressed views and the constraint implied views. In addition, even though the most extreme view tracking error contribution (on Harmony Gold) was significantly reduced when constraints were enforced, the overall portfolio risk budget became significantly more concentrated, with fewer views occupying the majority of the portfolio active risk. This therefore indicates that instead of the constraints reducing portfolio risk as intended, the results

[75] A view score change is determined if the constraint-implied view score does not fall within the same bounds as the originally expressed view score. This differs from the criteria utilised for figure 24 whereby a view score change is recorded when the constraint-implied view differs by more than 2.5 from the originally expressed view score.

indicate that for this particular portfolio and set of constraints, the views and view risk contributions became severely distorted and more concentrated, therefore resulting in a riskier portfolio with active weights in unintended positions where the investor may not exhibit sufficient skill or convictions.

The practical implication of the obtained results is that it is imperative that investors analyse the influence of the views on the portfolio before and after constraints are imposed in order to determine whether the objectives of the investment portfolio are still being met. Given the tools in order to analyse the relationships between the views, constraints and active weights, investors are able to calibrate each component in order to result in a portfolio that accurately reflects the respective convictions and investment goals set out by the investor.

6.3 Analysing the Impact of Relaxing the Long-Only Constraint

The previous sections illustrated how constraints impact the optimal portfolio, from a view distortion as well as from a risk budgeting perspective. As illustrated in the previous sections, the constraints severely impacted the clarity of the expressed views as well as impaired the risk and return characteristics of the portfolio. The aim of this section is to examine how a mild relaxation of the long-only constraint can significantly improve the optimal constrained portfolio in terms of the consistency in which the views are represented in the final portfolio relative to the unconstrained portfolio. The positive impacts of relaxing the long-only constraint can be observed as a decrease in the overall proportion of the portfolio due to constraints, which would further imply that the portfolio is more strongly driven by the view and benchmark components implied by the investor. In addition, the benefits of marginally relaxing the long-only constraint can be observed by the reduction of the view distortions.

Given that the Lagrange multipliers arise as a result of the expressed constraints and are determined in a non-linear and indistinct manner, controlling the proportion of the portfolio statistics influenced by the constraints is not a simple task. The most apparent manner in which the impact of constraints on the portfolio can be minimised is to marginally relax some of the most binding constraints.

For each marginal relaxation of the long-only constraint, one can expect marginal improvements that increase at a decreasing rate up until a point where the improvements of relaxing the long-only constraint flatten out and no further improvement in the distortion effects is observed. It is at this point where the additional imposed constraints (zero net active position and tracking error constraints) become binding and exert a greater influence on the final portfolio weights. The following sections analyse the various characteristics of the portfolio for marginal relaxations of the long-only constraint.

6.3.1 Impact on Portfolio Weights and Tracking Error

Figure 13 below illustrates the respective proportions of the overall portfolio weights that are composed of the benchmark, views and constraint component weights[76]. It must be noted that as the long-only constraint is relaxed, the constraint weight proportion decreases, while the benchmark and view proportions increase, while maintaining a constant benchmark-view weight ratio[77]. It therefore follows that as the constraint proportion decreases, the constrained portfolio becomes increasingly similar to the unconstrained portfolio and therefore, the benchmark and view weights exert a larger influence on the

[76] Figure 13 separates the final portfolio weight vector into three respective components. For each component vector, the norm is calculates and is divided by the sum of the each respective component's norms, therefore computing the proportion of the overall portfolio weight vector attributable to each component.

[77] It is relatively easy for the investor to control and adjust the proportions of the portfolio explained by the benchmark and investor views. It is simply a matter of adjusting the covariance shrinkage factor (τ). Increasing the value of τ results in increased uncertainty of the implied equilibrium returns. The result is a shift away from the benchmark weights in the direction of the expressed views.

final portfolio weights. One can therefore expect the other portfolio characteristics (return, variance etc) to be increasingly driven by the views of the investor relative to the benchmark portfolio.

Figure 13: Proportions of the Portfolio Weight for Relaxing the Long-Only Constraint

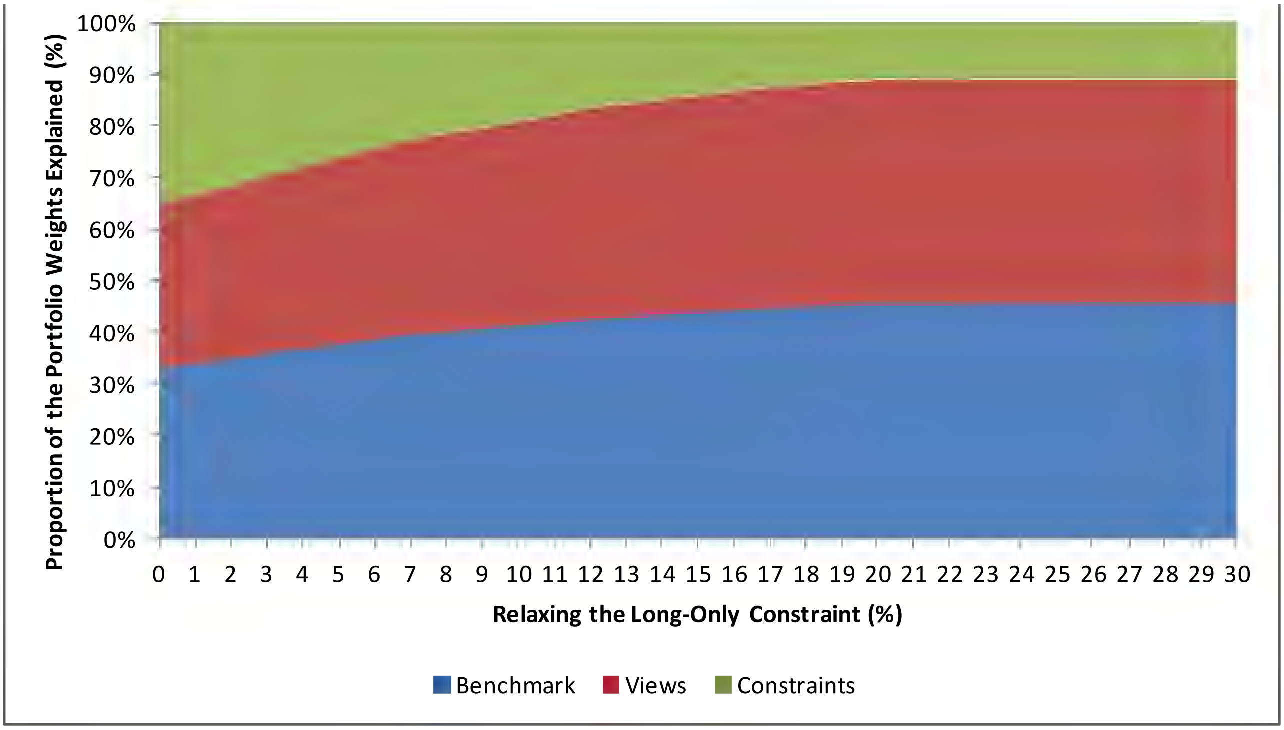

In particular, when analysing figure 13, for the fully constrained portfolio (0% relaxation of the long-only constraint), the portfolio weight proportions are as follows: benchmark (32.97%), Views (31.57%) and Constraints (35.46%). As the long-only constraint is relaxed, the following changes to the respective proportions are observed:

- For a 10% relaxation of the long only constraint, a 16.31% decrease in the constraint portion and a corresponding 8.15% increase to each of the benchmark and view proportions is observed.
- For a 15% relaxation of the long only constraint, a 21.38% decrease in the constraint portion (relative to a zero relaxation) and a corresponding 10.69% increase to each of the benchmark and view proportions is observed.
- For a 20% relaxation of the long only constraint, a 24.49% decrease in the constraint portion (relative to a zero relaxation) and a corresponding 12.25% increase to each of the benchmark and view proportions is observed.
- For a 23% relaxation of the long only constraint, a 24.70% decrease in the constraint portion (relative to a 0% relaxation) and a corresponding 12.35% increase to each of the benchmark and view proportions is observed. It can be noted that further marginal relaxations of the long-only constraint do not result in any further changes to the portfolio components proportions.

Observing figure 13, for a 23% long-only constraint relaxation, the 123/23 final portfolio proportions remain constant as follows: Benchmark (45.58%), Views (43.65%) and Constraints (10.75%). This is a significant decrease in the constraint proportion, from 35.46% to 10.75%. The following sections will discuss how this observed significant decrease in the constraint proportion results in a portfolio that exhibits a significantly lower level of view distortion and will therefore be more in line with the objectives of the investor.

Given that figure 13 illustrated how the views and benchmark component proportions increased as the long-only constraint was marginally relaxed[78], the objective is now to observe how the tracking error components change for marginal relaxations of the long-only constraint.

Figure 14: Proportions of the Portfolio Tracking Error Components for Relaxing the Long-Only Constraint

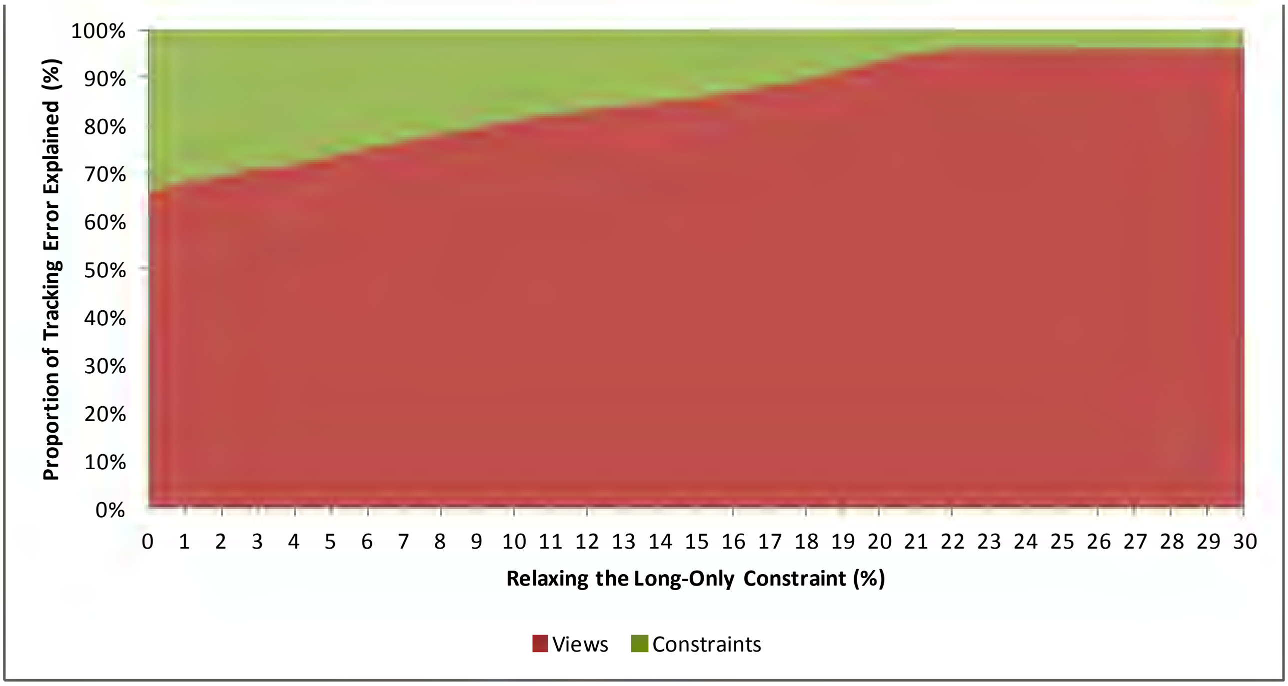

As observed by figure 14, a decreasing constraint component proportion is favourable as it provides the indication that the active weights deviating from the benchmark are progressively more as a result of the investor views as opposed to the uncontrollable effects of the imposed constraints. For marginal relaxations of the long-only constraint, the following can be noted:

- For the fully constrained portfolio (0% relaxation), the view component occupies 66.17% and the constraint proportion occupies 33.83% of the overall portfolio tracking error. Therefore, the constraints imposed on the portfolio effectively control a third of the portfolio active risk.
- For each marginal relaxation of the long-only constraint, a near linear reduction in the constraint proportion is observed. In particular, for each marginal 1% relaxation between 0% and 23% relaxation, an approximate average of a 1.29% increase in the view component proportion is observed. Therefore, for a 23% relaxation of the long-only constraint, a 29.76% reduction in the

[78] Given that the portfolio tracking error is a measure of active risk in excess of the benchmark, the components are the views and constraints respectively.

constraint component proportion is observed, thereby leading to a 95.93% view component proportion and corresponding 4.06% constraint proportion, after which no further change is observed.

The significant decrease of the constraint proportion (from 33.83% to 4.06%) for the portfolio tracking error, indicates that the investor has significantly more control over how the active risk of the portfolio is distributed.

It must be noted that analysing the relative proportions of the portfolio weight and tracking error components does not necessarily provide a complete indication of whether there is indeed an increase in the consistency between the constrained and unconstrained portfolios. More specifically, in addition to analysing the portfolio components, one needs to observe specifically *how* the portfolio is changing. In order to determine the rate at which the constrained portfolio becomes increasingly similar to the unconstrained portfolio as the long-only constraint is relaxed, the Mahalanobis distance between the vector of tracking error contributions for the constraint-implied views and the original view tracking error contributions was used. In addition, the Mahalanobis distance between the vector of constrained portfolio active weights and vector of unconstrained active weights was recorded for each marginal relaxation of the long-only constraint.

Figure 15: Mahalanobis Distances between the Unconstrained and Constrained Portfolio Active Weights and View Tracking Error Contributions for Relaxing the Long-Only Constraint

Examining figure 15 above, as the long-only constraint is relaxed, the Mahalanobis distance for the tracking error and portfolio active weights monotonically decreases significantly. This therefore suggests that for every marginal relaxation (between 0% and 23%), an increased level of consistency between the constrained and unconstrained portfolio is observed.

Observing the tracking error Mahalanobis distance, there are few improvements arising from a 1% to 2% as well as from a 3% to 4% relaxation of the long-only constraint (as observed by the horizontal “steps”). Therefore, from the perspective of a the risk budget consistency between the constrained and unconstrained portfolios[79], a 1%, 3% or a relaxation greater than 4% would be mostly beneficial to the portfolio as it provides significant breadth for the portfolio allocations to move into superior positions that would significantly improve the consistency of the tracking error risk budget.

Observing the active weights Mahalanobis distance curve, there is a relatively sharp decline for a 2% long-only constraint relaxation, after which a kink is observed and the curve gradually flattens out and the rate of decrease for the Mahalanobis distance decreases until the point (19% relaxation) at which the marginal decreases become fairly negligible.

Given the above describe sharp decreases in the Mahalanobis distances for the active weights and the tracking error contributions, it can be concluded that there is indeed an increase in the similarity between the constrained and unconstrained portfolios as the long-only constraint is marginally relaxed. Therefore, relaxing the long-only constraint has resulted in a significant reduction in the distortion of the optimal constrained portfolio in terms of the active weights and the consistency of risk budget.

As illustrated above, a 23% relaxation of the long-only constraint results in the greatest reduction in the constraint proportion of the portfolio as well as the consistency between the tracking error contribution vectors. In order to observe and compare the final portfolio weights of the fully constrained, 123/23 portfolio and unconstrained portfolio, figure 16 below plots the asset allocations for each asset for each of the three portfolios. As illustrated, it can be noted that the fully constrained portfolio is relatively concentrated in approximately 14 of the 40 assets. In contrast, the 123/23 portfolio more closely tracks the unconstrained portfolio for each asset allocation. Comparing the 123/23 portfolio and the unconstrained portfolio, it can also be noted that for each asset (with exception for Massmart), the unconstrained portfolio allocations are significantly larger in magnitude, therefore indicating an increased level of riskiness.

[79] More specifically, for the risk budgets between the constrained and unconstrained portfolios to be consistent, the tracking error contribution of each asset needs to be similar.

Figure 16: Comparing the Allocations for the Fully Constrained, 123/23 and Unconstrained Portfolios

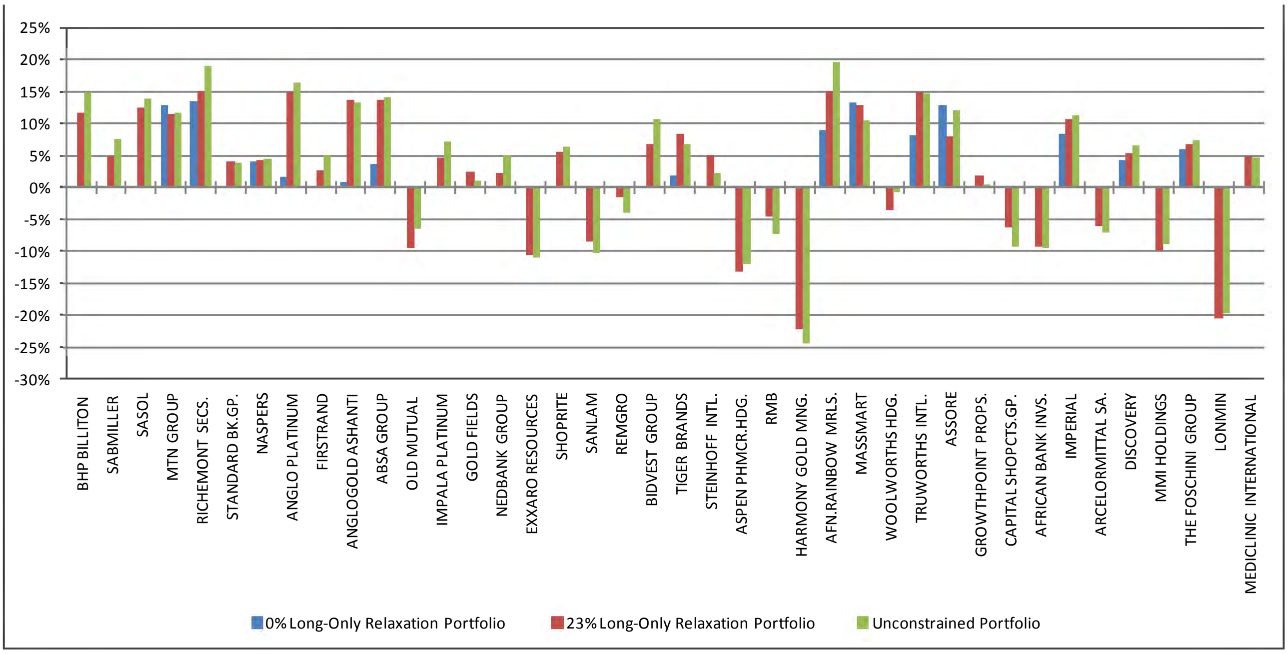

6.3.2 Impact on the Investor View Distortions

The previous section analysed how the overall constrained portfolio weights and tracking error became increasingly similar to the unconstrained portfolio in terms of an observed reduction of the constraint component as well as the decrease of Mahalanobis distance between the two portfolios. The aim of this section is to observe how the individual view distortion percentages change as the long-only constraint is relaxed. In this way, one obtains a clearer picture of the benefits of relaxing the long-only constraint on an individual view level.

In order to appropriately analyse the view distortion percentages as well as for improved readability, the view distortions for each of the 40 assets were placed in four groups and were plotted on four different figures. Figures 17 to 20 all illustrate how the respective view distortions change for marginal relaxations of the long only constraint.

It can be noted that a view distortion percentage near zero indicates a low level of distortion, with a percentage value of exactly zero indicating no distortion[80]. It must be noted that while the long-only constraint is relaxed, the zero net active position, beta neutrality and tracking error constraints still remain enforced. Therefore, as the long-only constraint is relaxed and becomes less binding, the remaining constraints become increasingly binding and exert a greater influence on the portfolio allocations.

As illustrated by figure 17 (Group 1) below, as the long-only constraint is marginally relaxed, there are many confounding effects that occur as a result of the additional constraints becoming increasingly binding. Of

[80] A comprehensive discussion and interpretation of the view distortion percentages was covered in section 6.2.3

major concern are the view distortions that exceed -100%, thereby indicating a view score that has changed sign[81]. In particular, the views on First Rand, SASOL, Anglo Platinum and Anglo Gold are of significant concern as their respective constraint-implied views are essentially a reversal, by a larger magnitude, of the originally expressed views. Observing figure 17 further, a 6% or greater long-only constraint relaxation corrects for the reversal of the above mentioned view distortions. However, it can also be noted that for a long-only constraint relaxation greater than 18%, the view distortions marginally increase, mainly as a result of the magnitude of view distortion on FirstRand becoming increasingly negative again, as well as the view distortions of Old Mutual and ABSA increasing at a steady rate.

Figure 17: Asset Group 1 View Distortion Percentages as the Long-Only Constraint is Relaxed

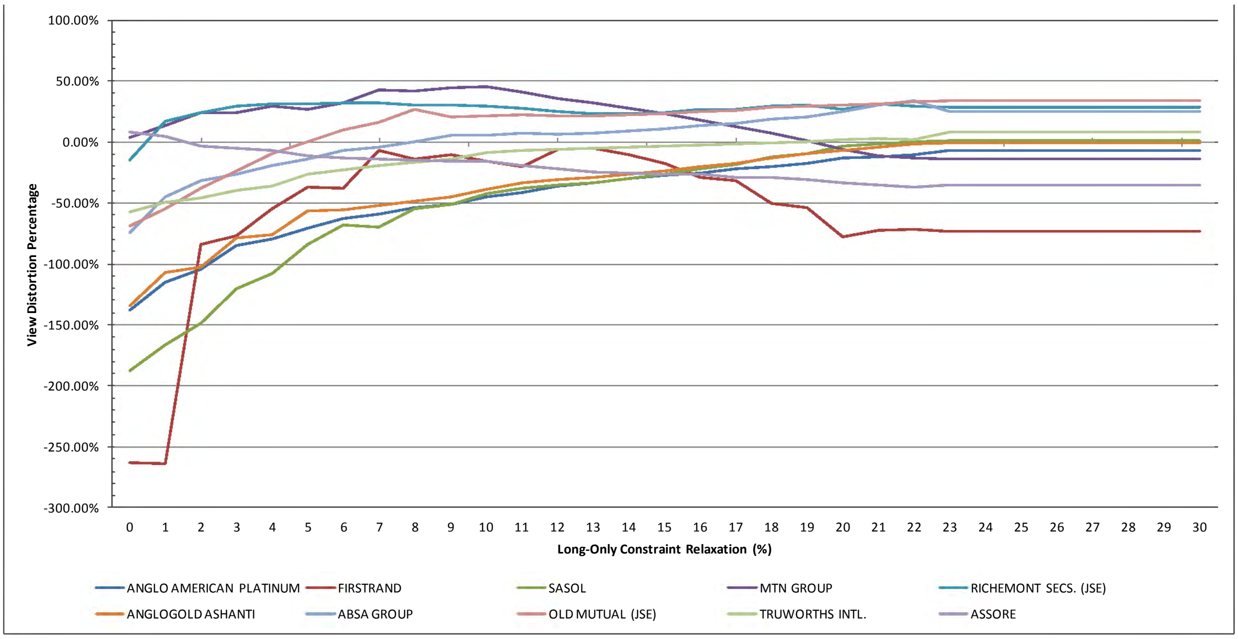

Observing figure 18 (group 2) below, it is evident that for a 0% relaxation of the long-only constraint, all the expressed views experienced a negative view distortion, with the views on Shoprite and Bidvest initially experiencing a view reversal as shown by the initial view distortions of -136.56% and -156.25% respectively. As shown by the view distortions close to -100%, it can also be noted that the views on Harmony Gold and Aspen Pharmaceuticals were effectively neutralised as a result of the constraints. Their respective view distortions of -95.63 and -96.62% effectively results in constraint-implied view scores that are near zero, therefore resulting in near neutral view scores being expressed.

Relaxing the long-only constraint by 6% results in a significant reduction in the distortion levels of the above mentioned views, whereby none of the views are reversed. In addition, while the majority of the distortion levels for the views decrease for a 6% relaxation, the views on Tiger Brands and Remgro experience a view expansion (40.59% and 54.13%) at a magnitude similar to the original (0% relaxation) view reduction values of -51.01% and -56.71% respectively. It can also be noted that the greatest reduction in the view

[81] A view score that has changed sign is one that has changed from originally being positive (negative) to becoming negative (positive) on the particular asset. A view score sign change is an indication of a severe level of distortion.

distortion percentages occurs for a 23% relaxation of the long-only constraint, where the majority of the views experience a significant decrease in distortion.

Figure 18: Asset Group 2 View Distortion Percentages as the Long-Only Constraint is Relaxed

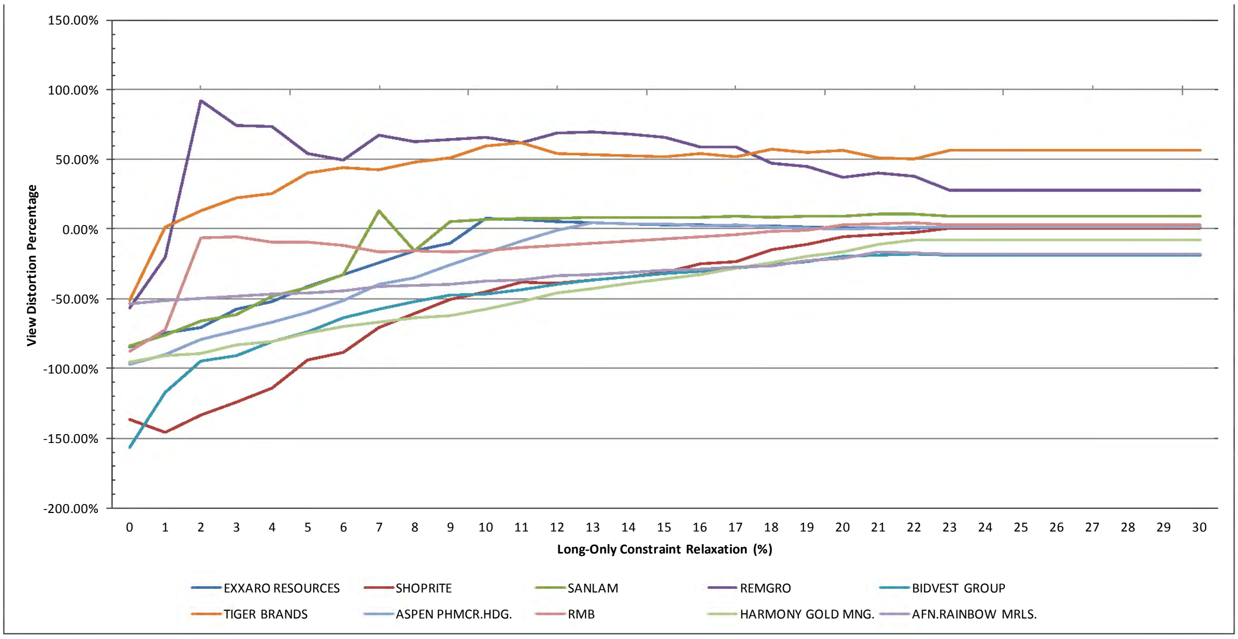

Figure 19: Asset Group 3 View Distortion Percentages as the Long-Only Constraint is Relaxed

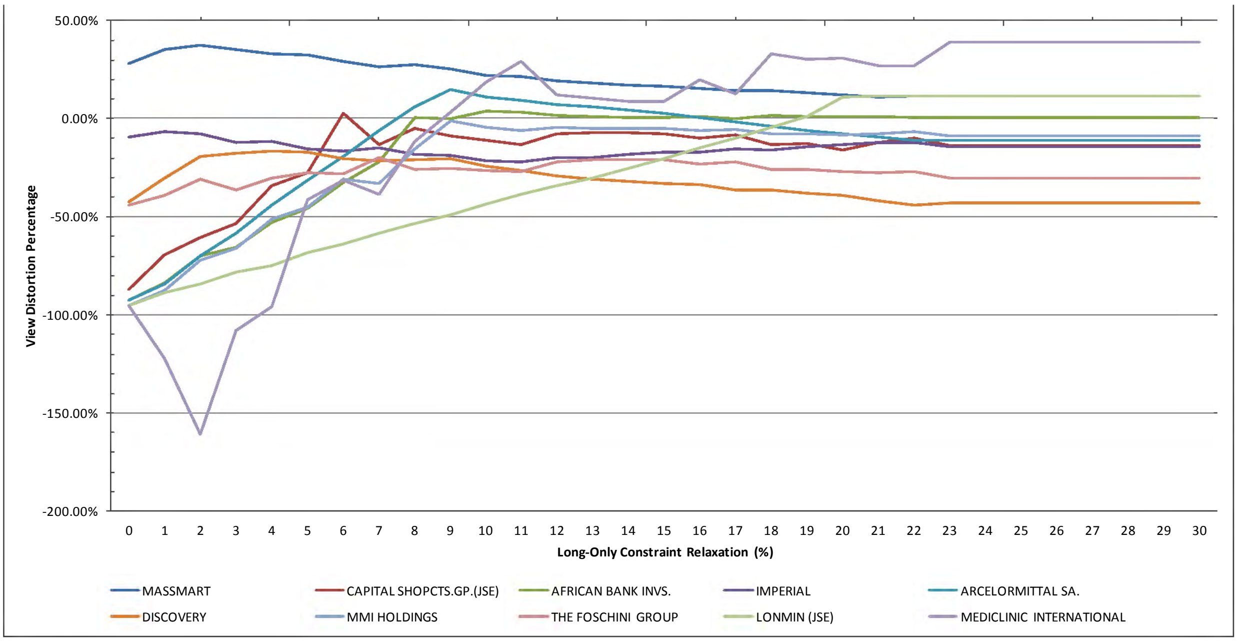

Examining figure 19 above (Group 3), as the long-only constraint is relaxed, an overall level of view distortion reduction is observed. However, while the overall level of view distortion decreases in magnitude, the view distortion of Mediclinic International significantly increases for a 3% long-only constraint relaxation, whereby the view distortion changes from being a near neutralising distortion to a -160.82% view reversing distortion. Further relaxing the long-only constraint rectifies the increased distortion for the Mediclinic view, whereby the lowest view distortion levels for the group 3 views occurs for an 18% relaxation of the long-only constraint.

Figure 20 (Group 4) below groups the views that experienced the greatest level of view distortion. It can be noted that all the views in group 4 experienced view distortions in magnitude greater than 100% and, in particular, it can be noted that the view distortions of Standard Bank and Growth Point Properties varied considerably as the long-only constraint was relaxed. The view distortion on Standard Bank varied from a 1036.72% view distortion (for a 0% relaxation) to a -371.48% view distortion (for an 8% relaxation) and reaching a minimum 0% distortion level for an approximate 17.5% long-only constraint relaxation. Of significant concern is how the view distortion percentage for Growth Point Properties effectively *increases* for a relaxation of the long-only constraint. In particular, for a 1% relaxation of the long-only constraint, the view distortion percentage dramatically increases from its original 22.47% distortion percentage to 189.10% and reaches a -843.46% view distortion percentage for an 11% relaxation. In addition, it can be further noted that an 18% long-only constraint relaxation results in the lowest level of view distortion amongst the views included in group 4 below.

Figure 20: Asset Group 4 View Distortion Percentages as the Long-Only Constraint is Relaxed

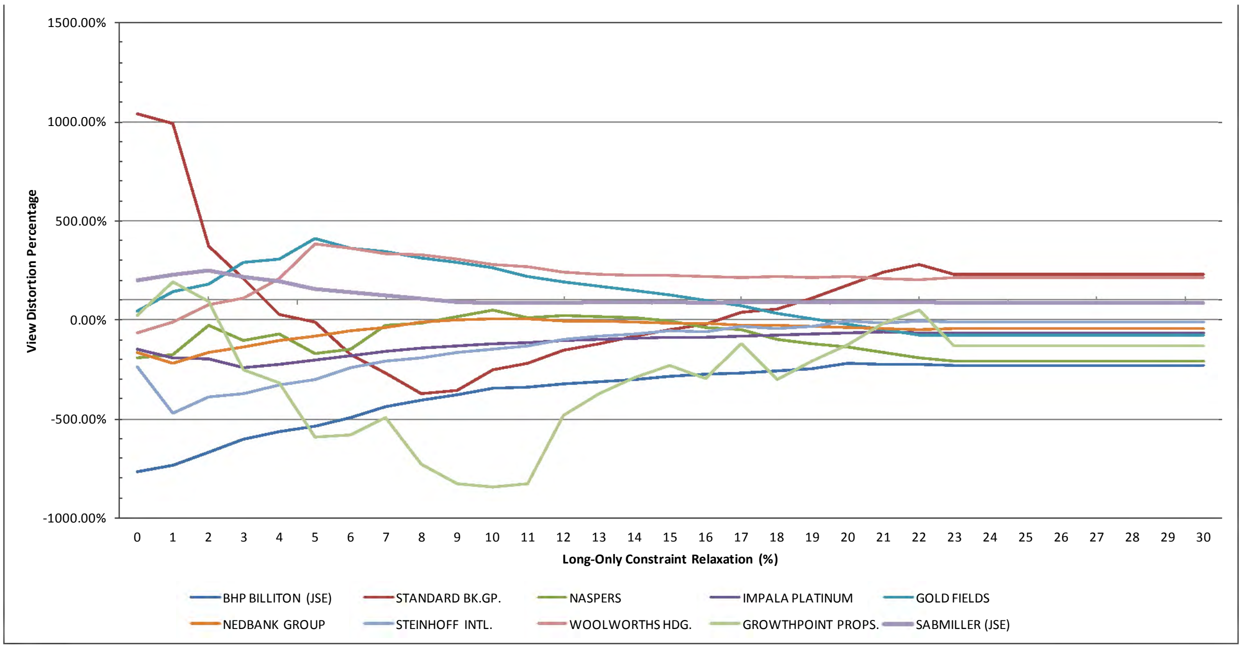

As illustrated in figures 17 to 20, it can be concluded that, on an overall view level, the level of view distortion significantly decreases when the long-only constraint is marginally relaxed. However, given the confounding effects of the additional portfolio constraints, the individual view distortion levels may increase or decrease for marginal relaxations of the long-only constraint. However, it should be noted that despite individual view distortions increasing, the overall level of view distortion decreases and, consistent with the previous results, no further improvements in the view distortions are observed for a long-only constraint relaxation in excess of 23%.

6.3.3 Tracking Error Contributions Proportions

In order to further illustrate how marginal relaxations of the long-only constraint influence the active risk contributions of the investor views, figure 21 below plots the tracking error contribution proportions for each view as the long-only constraint is relaxed[82]. In addition, figure 21 adds to figure 15 in the previous section by illustrating *how* the tracking error is composed as well as how relaxing the long-only constraint influences the contribution proportion of each respective view. Closely observing the tracking error proportions for a 0% relaxation and comparing the percentage contributions to a 23% long-only constraint relaxation, it becomes apparent that the respective contributions, amongst others, of Harmony Gold, Massmart, Lonmin, Assore and BHP Billiton experienced the most significant changes. In particular, the following can be noted:

- The tracking error contribution proportion of Harmony Gold for the fully constrained portfolio rapidly increased from 1.61% to 27.77% of the overall portfolio tracking error proportion contribution. This therefore illustrates that the implementation of constraints effectively neutralised the investor's view on Harmony Gold.
- The contribution of Lonmin increased significantly from a 0.8% contribution to a 17.65% contribution, therefore indicating an additional view that was effectively drowned out as a result of the imposed constraints.
- The view tracking error contributions of Massmart and Assore experienced similar changes, whereby the contribution of Massmart decreased from 14.01% to 4.05% and the view tracking error proportion of Assore decreased from 12.66% to 2.13%. This therefore indicates a significant level of view amplification whereby the contraints effectively *increased* the active risk contribution of Massmart and Assore.
- For a 0% relaxation, the constraint-implied view for BHP Billiton contributed 11.47% to the overall portfolio tracking error. As the long-only constraint was relaxed, the proportion decreased to 0.59%, which is consistent with the unconstrained portfolio case. Given that BHP Billiton occupies the greatest weight in the benchmark, any significant change in the weight allocated to it will most likely result in a significantly large contribution to the portfolio tracking error.

In addition to the aforementioned, analysing figure 21 further, for relatively low levels of relaxation (between 2% and 7%), the tracking error proportions for Anglo Platinum, Anglo Ashanti and SASOL become fairly eratic. In particular, the sharp decrease in the proportions for Anglo Platinum and Anglo Ashanti for a 2% relaxation occur at a point where their respective contributions change from positive to negative and continue to grow thereafter. In addition, the sharp decrease for SASOL occurs between 3% and 7% at a

[82] Caution needs be exercised when interpreting figure 21. Figure 21 plots the respective tracking error proportions for each of the view but does not indicate whether the view tracking error contributions are changing sign as the long-only constraint is relaxed. For example, as a result of the view distortions, as the long-only constraint is relaxed, a particular view which originally possesses a negative tracking error contribution may change to positively contributing to the overall portfolio tracking error.

point where the tracking error contribution changes from positive to negative. In all three cases, as the long-only constraint is relaxed, the respective contributions move in a direction that is increasing ly characteristic of the unconstrained portfolio case.

Figure 21: Constraint-Implied View Tracking Error Contribution Proportions as the Long-Only Constraint is Relaxed

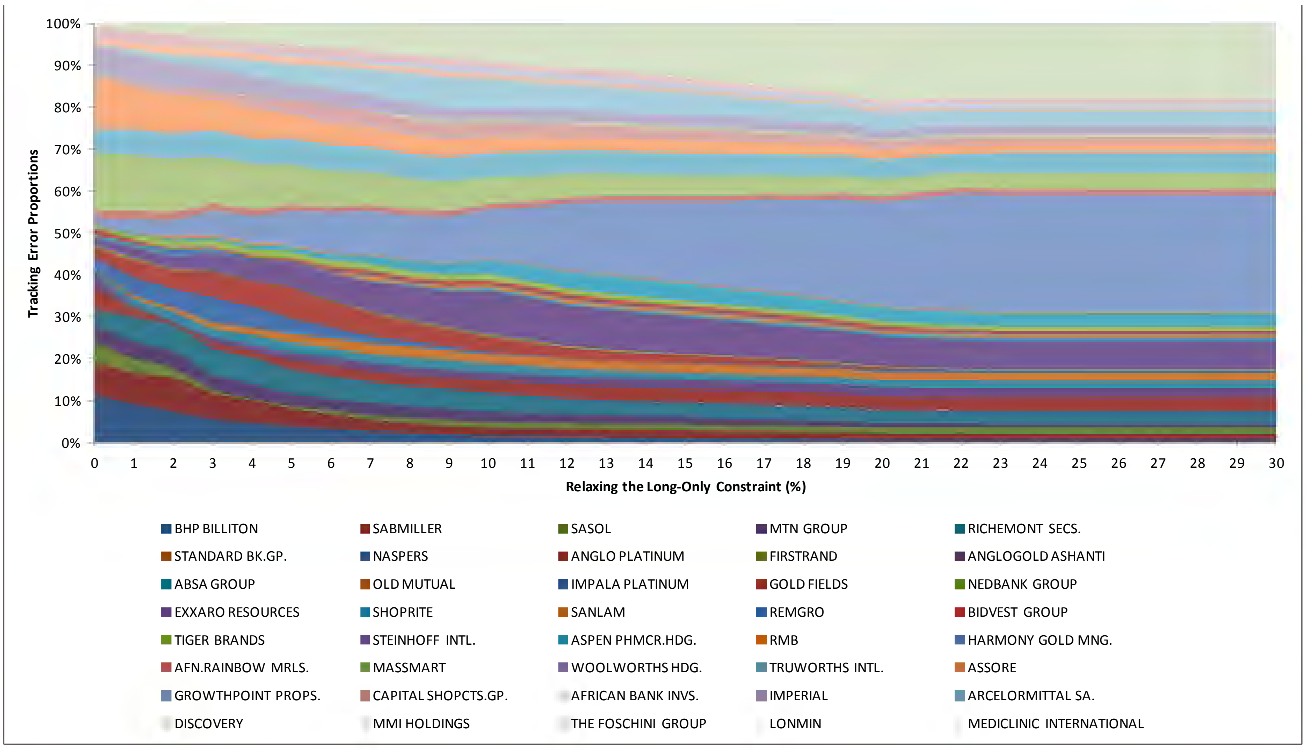

Figure 22 on the following page complements figure 21 by comparing the tracking error contribution proportions of the 23% relaxed long-only constrained portfolio (123/23 portfolio) to the unconstrained portfolio. Examining figure 22 in conjunction with figure 21, it becomes evident that the 123/23 portfolio view tracking error proportions is significantly closer to the unconstrained portfolio. In particular, it can be noted that, with exception for 5 views, the remaining 35 views differ by no more than 2% from the original tracking error contribution proportions. The 5 views for the 123/30 portfolio versus the unconstrained portfolio, that differed significantly were:

- Lonmin: 17.65% versus 9.83% (7.82% difference)
- African Rainbow: 1.28% versus 6.36% (-5.08% difference)
- Assore: 2.13% versus 6.28% (-4.14% difference)
- Anglo Platinum: 3.33% versus 0.47% (2.86% difference)
- Exxaro Resources: 7.03% versus 4.33% (2.70% difference)

As illustrated above, it therefore follows that a mild 23% relaxation of the long-only constraint has resulted in an increased level of consistency between the constraint-implied view tracking error contributions and the unconstrained portfolio view tracking error contributions. This therefore further results in a clearer expression of the investor views in the presence of remaining imposed constraints.

Figure 22: Pie Charts Comparing the View Tracking Error Contribution Proportions for the Unconstrained Portfolio and the 123/23 Portfolio

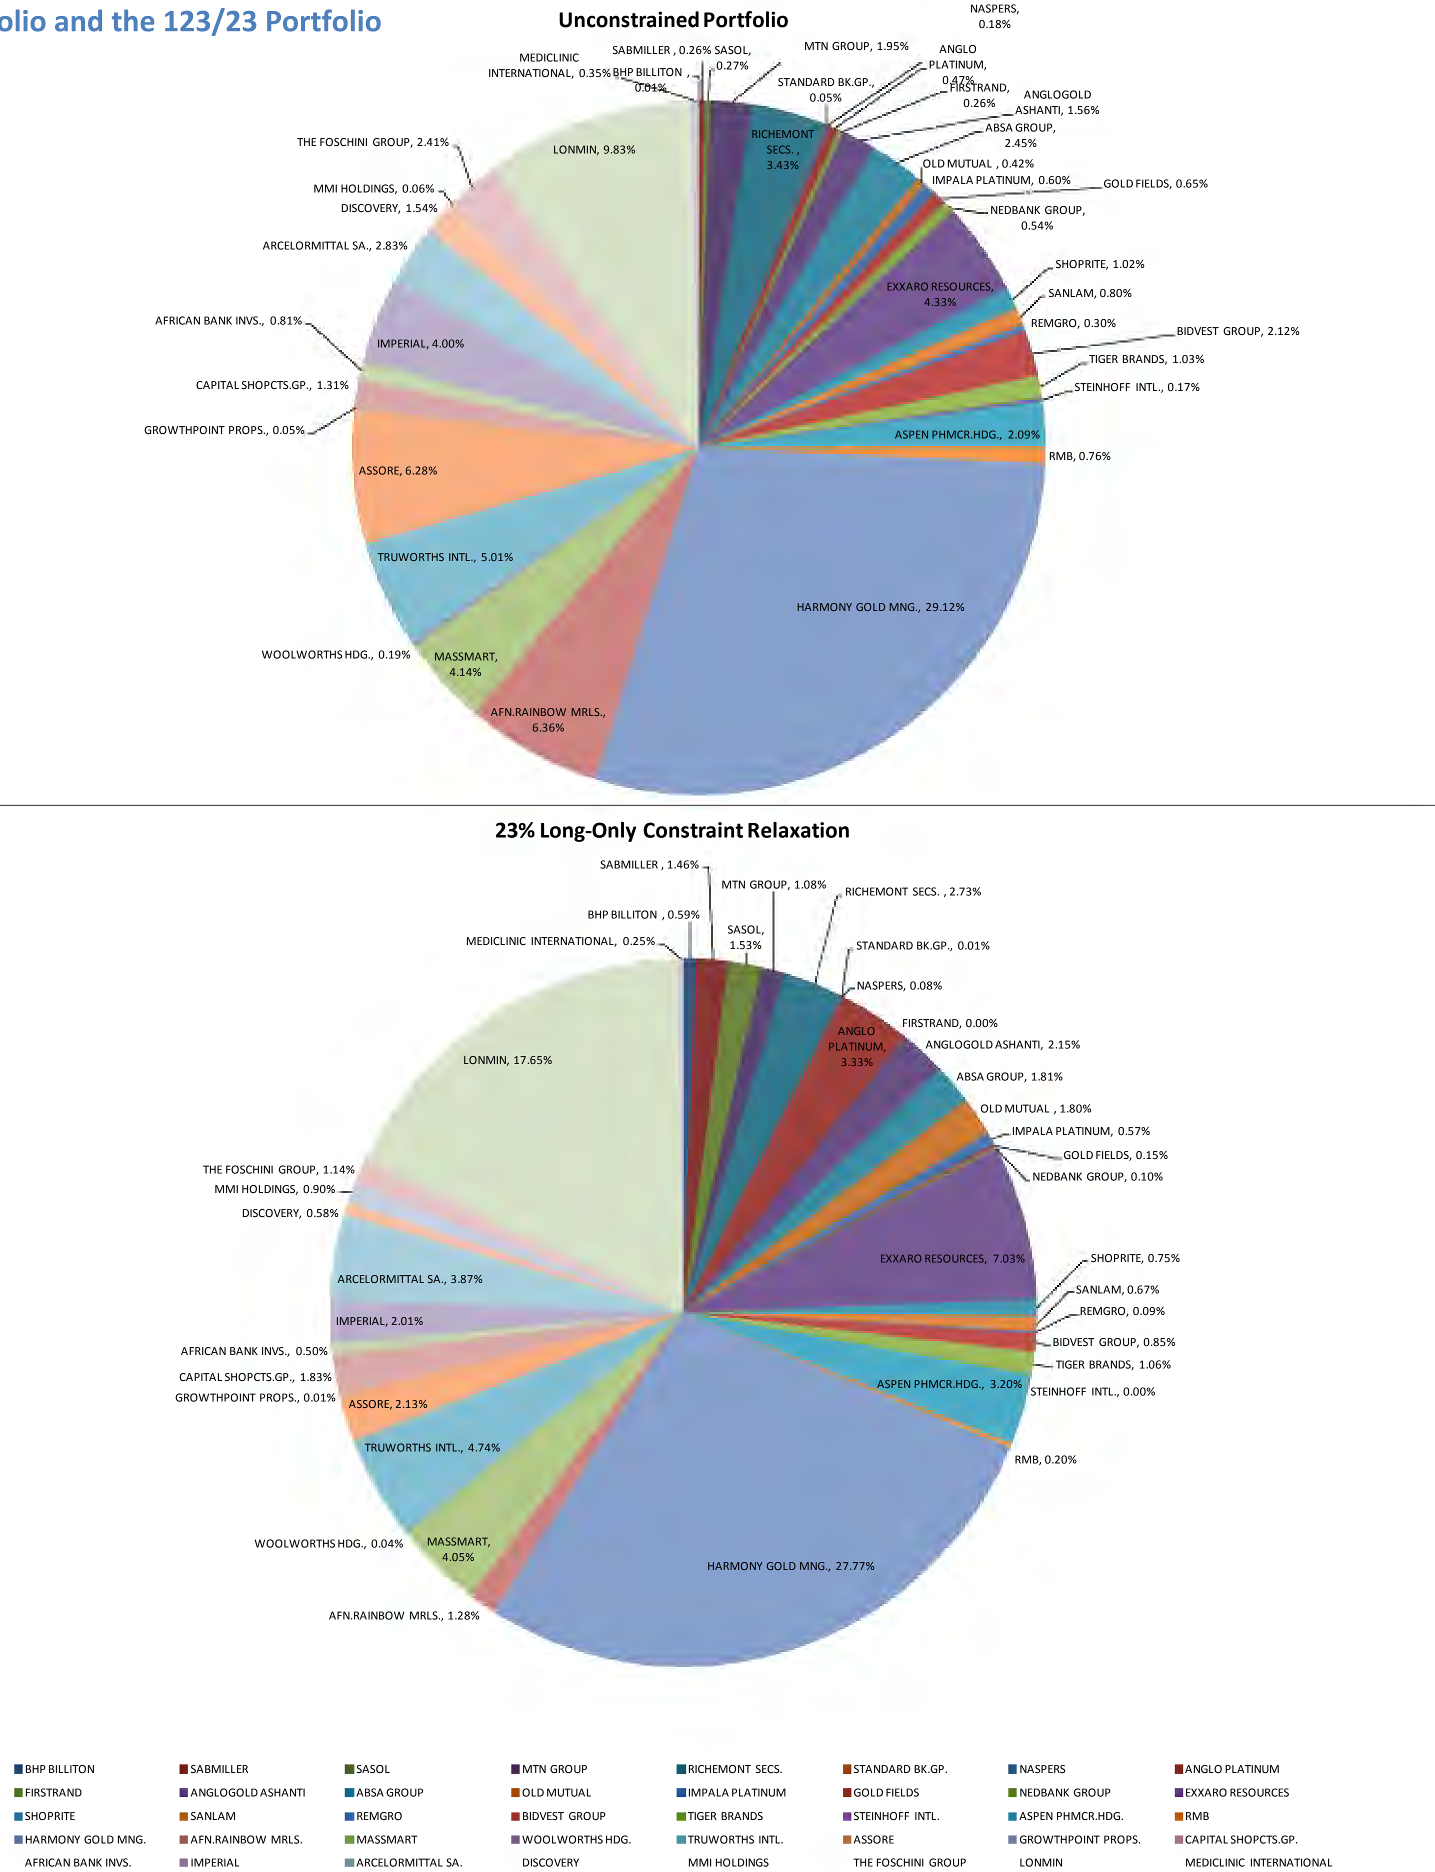

6.3.4 View Vector and Number of View Changes

In order to obtain an overall picture of how the level of view distortion decreases for relaxing the long-only constraint, figure 23 below plots the Mahalanobis distance between the constraint-implied view and originally expressed view vectors. As the Mahalanobis distance decreases, the constraint-implied view vector becomes increasingly similar to the view vector originally expressed by the investor. At a 19% relaxation of the long-only constraint, the rate of decrease between the two vectors decreases significantly and further relaxations of the long-only constraints provide small benefits, whereby a 23% relaxation achieves the smallest distance between the two vectors and no further improvement is obtained. The end result is a significant improvement in the efficiency for which the investor views are expressed.

Figure 23: The Mahalanobis Distance Between the Constraint-Implied and Original View Vectors for a Relaxation of the Long-Only Constraint

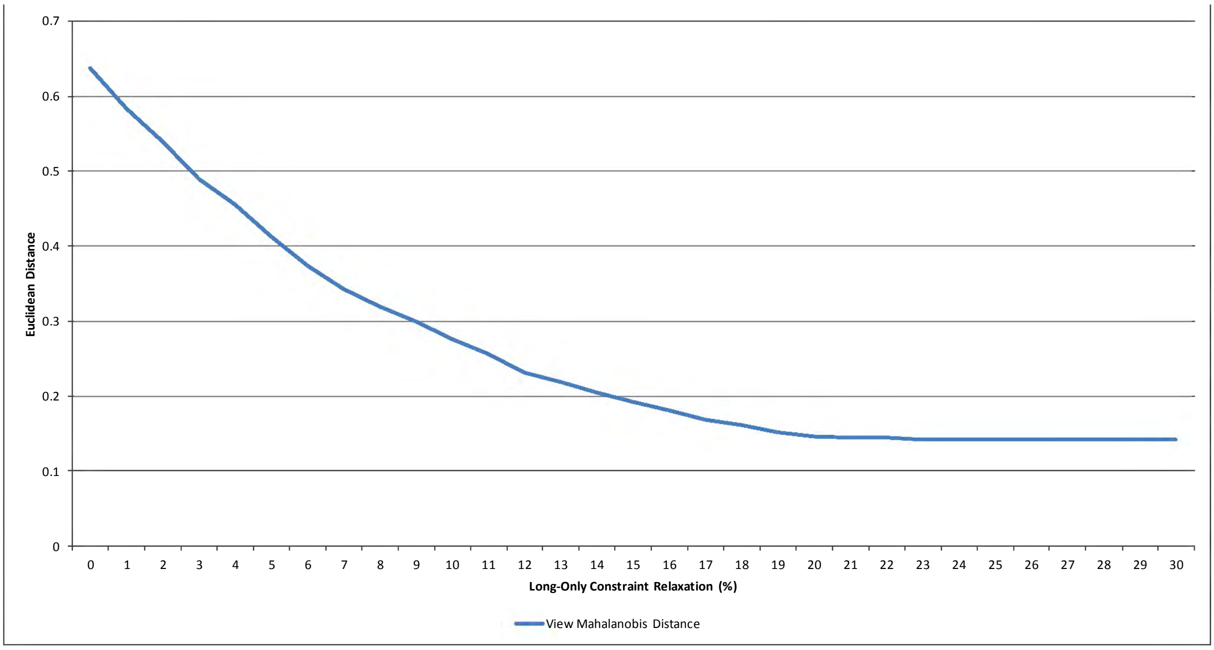

In order to further illustrate how the view expression improves as the long-only constraint is relaxed, the number of view score changes for relaxing the long-only constraint was recorded and represented in figure 24 below. In particular, a view score change was defined as an absolute difference of 2.5 between the constraint-implied view score and the original view score[83].

Observing figure 24 below, as the long-only constraint is relaxed, the number of view score changes significantly decreases from 28 changes (0% relaxation) to a minimum of 3 view changes (20% relaxation), whereby the respective view changes recorded were for: BHP Billiton (-2.88 difference), SAB Miller (-3.341 difference) and Tiger Brands (3.04 difference).

[83] For example, if the original view score for asset A was $\eta^{original} = 5.5$ and the corresponding constraint-implied view score is $\eta^{con} = 2$, then a view score change is recorded since $|\eta^{original} - \eta^{con}| \geq 2.5$

Closely observing figure 24, it becomes apparent that, for certain marginal relaxations of the long-only constraint, the number of view changes increases, rather than decreasing as expected. However, it must be noted that figure 24 cannot be interpreted in isolation and that figure 23 must be considered when interpreting the number of view changes taking place. In particular, it can be noted that, despite the occurrence of a number of view changes *increase* from 3 view changes to 5 view changes for a marginal relaxation of the long-only constraint from a 20% to 21% relaxation, figure 15 shows the overall view tracking error contributions and portfolio active weights becoming increasingly similar to the unconstrained portfolio. In other words, the number of recorded view changes is cannot be considered as a fully concise measure of the improvement of the constrained portfolio and that, despite the number of view changes increasing, the overall constrained portfolio becomes increasingly similar to the unconstrained portfolio case.

Figure 24: Observing the Number of View Score Changes as the Long-Only Constraint is Relaxed

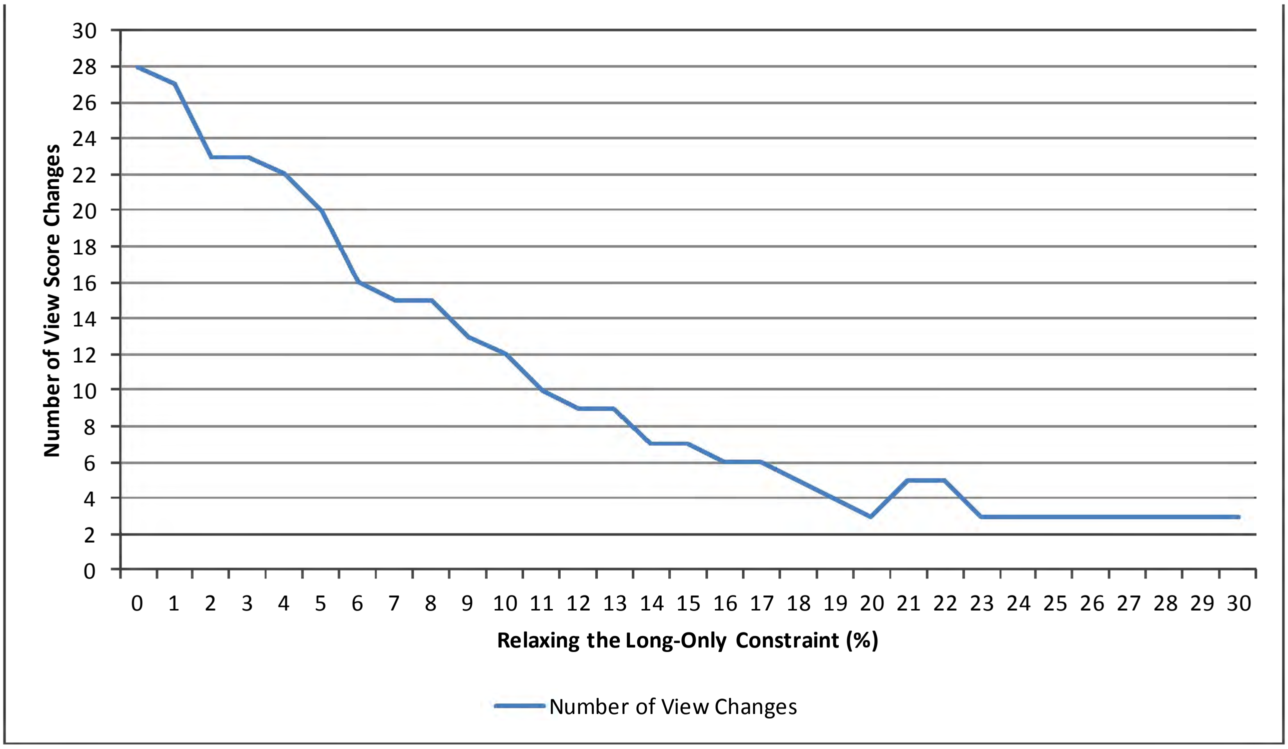

6.3.5 Impact on Utility

In addition to examining how marginal relaxations of the long-only constraint improve the consistency of the constraint-implied views relative to the originally expressed constraints, the change in the expected investor utility is of significant importance to the investor.

As discussed in section 5.1.2, when constraints are imposed on a portfolio, a significant reduction in utility is typically observed. This is due to the inconsistencies in the portfolio risk and return estimates that are created in order for the portfolio to meet the imposed constraints. In order to measure the relative impact on utility as a result of the constraints, the percentage of utility shortfall of the constrained portfolio relative to the unconstrained portfolio[84] is computed. The greater the level of utility shortfall of the constrained portfolio relative to the unconstrained portfolio, the greater the *negative* impact of the constraints on the final portfolio is experienced. It therefore follows that the observed utility shortfall percentage provides and overall summary of the level of impact of the constraints.

Figure 25 below illustrates the decrease of the utility shortfall percentage as the long-only constraint is marginally relaxed. It can be noted that for the fully constrained portfolio (0% relaxation), the portfolio utility for the constrained portfolio utility decreases by 30.15% relative to the unconstrained portfolio. This therefore suggests that the nearly a third of the portfolio utility is lost as a result of the imposed constraints. For a 1% relaxation of the long-only constraint, the constrained portfolio utility increases by 4.08%, which corresponds to a 26.07% utility shortfall relative to the unconstrained portfolio. Further incremental relaxations of the long-only constraint result in further additional improvements to the constrained portfolio's expected utility. Relaxing the long-only constraint by 9% results in an additional 20.40% increase in the expected utility of the portfolio, corresponding to a 9.75% utility shortfall. As a result of a 23% long-only constraint relaxation, the utility shortfall percentage drops to its minimum value of 4.34%, which is a significant improvement over the 30.15% utility shortfall percentage experienced by the fully constrained portfolio.

Given the observed improvements of the constrained portfolio expected utility for relaxing the long-only constraint and assuming that the investor has control over the level of long-only constraint relaxation, the investor can decide on the optimal level of long-only constraint relaxation relative to the gain in the expected utility.

Figure 25: Decrease in the Utility Shortfall Percentage for Relaxing the Long-Only Constraint

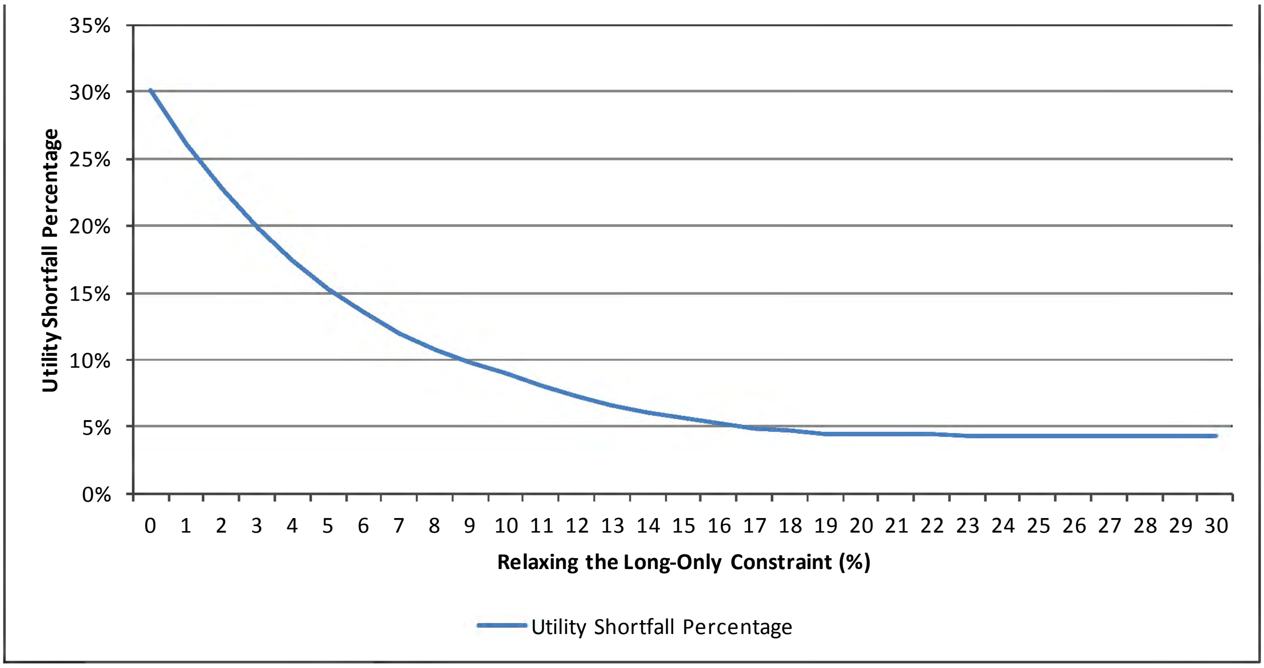

[84] Utility shortfall describes the percentage loss in utility of the constrained portfolio relative to the unconstrained portfolio.

As illustrated above, the benefits of marginally relaxing the long-only constraint have resulted in significant improvements to the efficiency of the constrained portfolio. In particular, not only are there significantly improvements in the consistency in which the views are expressed (in terms of view score values and risk contributions) but corresponding improvements in the expected utility are overwhelming.

6.3.6 Summary and Conclusion

The previous sections illustrated how a mild relaxation of the long-only constraint results in significant improvements to the optimal constrained portfolio. In particular, as the long-only constraint was marginally relaxed, there were observed decreases in the constraint proportions to the portfolio active weights, tracking error and other statistics of the portfolio. In addition, in order to measure the degree of similarity between the constrained portfolio active weights vector and the tracking error contribution vector with the corresponding unconstrained portfolio vectors, decreases in the Mahalanobis distance measure were observed. As final measures of how relaxing the long-only constraint has resulted in an improvement in which the investor views are expressed as well as the optimality of the portfolio, the number of view score changes and the decrease in the portfolio utility shortfall was recorded.

Table 11: Summary of the Results Obtained at Various Intervals of the Long-Only Constraint Relaxation

Long-Only Constraint Relaxation	Active Weights		Tracking Error		Views		Utility Shortfall Percentage
	Constraint Proportion	Distance Decrease[85]	Constraint Proportion	Distance Decrease	View Score Distance Decrease	Number of View Score Changes	
0%	35.46%	0%	33.83%	0%	0%	28	30.15%
5%	26.20%	39.48%	26.48%	47.70%	39.71%	20	15.30%
10%	19.14%	56.65%	18.93%	78.60%	56.81%	12	8.93%
15%	14.07%	61.63%	14.24%	90.06%	61.82%	7	5.64%
20%	10.96%	67.56%	6.67%	95.53%	67.70%	3	4.47%
23%	10.75%	69.03%	4.06%	96.34%	69.17%	3	4.34%
Unconstrained Portfolio	0%	100%	0%	100%	100%	0	0%

In order to summarise the above mentioned results, table 11 above summarises the observed results for relaxing the long-only constraint in increments of 5%, including a 23% relaxation. For a simple comparison, the fully constrained (0%) and unconstrained portfolio statistics were included.

The first category summarises the statistics of the active weights vector for a marginal relaxation of the long-only constraint. In particular, a significant decrease in the constraint proportion (from 35.46% to 10.75%) is observed for a 20% relaxation, after which, marginal improvements are observed for a 23%

[85] The *Distance Decrease* refers to the percentage decrease of the Mahalanobis distance between the marginally relaxed portfolio and the unconstrained portfolio active weight vectors, *relative* to the Mahalanobis distance between the *fully* constrained portfolio and the unconstrained portfolio vectors.

relaxation. A 69.03% decrease in the Mahalanobis distance between the marginally constrained active weight vector and the unconstrained active weight vector is observed for a 23% relaxation. This therefore indicates a significant increase in the similarity between the active weight vectors relative to the fully constrained portfolio.

It can be noted that the portfolio tracking error experiences the greatest improvement as the long-only constraint is relaxed. Similar to the active weight constraint proportion, 33.83% of the tracking error proportion is explained by the constraints. In contrast, for a 23% relaxation, the constraint proportion decreases significantly to 4.06% and the Mahalanobis distance measure decreases by 96.34%. This therefore results in an overwhelming improvement in the control and consistency of the view risk budget allocations.

Observing the final impact of relaxing the long-only constraint on the view scores, the Mahalanobis distance decrease between the marginally relaxed portfolio and the original view scores was calculated. For a 23% relaxation, a 69.17% decrease was observed, therefore implying that a more consistent constraint-implied view score vector was obtained. In addition, it can be noted that this particular reduction translated into a significant decrease in the number of view changes. In particular, the number of recorded view changes decreased from 28 (out of 40 views) to 3 view changes – which is an overwhelming improvement in the consistency of which the views are expressed.

The utility shortfall percentage, an overall measure of how relaxing the long-only constraint has benefitted the constrained portfolio, is the percentage for which the constrained portfolio utility is lower than the unconstrained portfolio. For a 5% relaxation, the 30.15% shortfall dramatically decreases to 15.30%. Relaxing the long-only constraint by 23% results in further significant decreases, where a 4.34% utility shortfall is observed. This is a significant improvement over the fully constrained portfolio and is sure to result in superior risk and return characteristics.

As illustrated above, marginally relaxing the long-only constraint has shown to result in significant improvements to the constrained portfolio efficiency. In particular, a relaxation of 23% resulted in the greatest improvement in the consistency for which the investor views are expressed in terms of minimising the view distortion as well as from a risk budgeting perspective. It must be noted that for different portfolios with different constraints and objectives, the optimal level of relaxation is likely to differ. It can also be noted that the optimal 23% level of short-selling observed for this particular example is somewhat consistent with the 30% optimal level of short-selling found by Johnson *et al.,* (2007), who conducted an extensive empirical analysis on various asset portfolios listed on the Russell 1000 and EAFE indices. In conclusion, it can be noted that the modest use of short positions as a result of marginally relaxing the long-only constraint, not only decreases the distortion present in the constrained portfolio but may also be more acceptable to investment committees who would otherwise not allow a full relaxation of the long-only constraint.

Chapter 7: Conclusion

In light of Harry Markowitz's (1959) statement, "The rational investor is Bayesian", throughout this thesis, the Bayesian approach to asset allocation has been described as the most appropriate methodology for incorporating human decision making in an asset allocation framework. One of the most attractive features of the Bayesian approach is the ability to combine sample information with information not available within the sample. In this manner, the Bayesian approach has proven to result in robust portfolios that intuitively reflect the individual preferences of the investor.

In order to further improve the practical applications of the Bayesian methodology, various extensions and additions to the Black-Litterman framework were explored. In particular, robust higher moment estimates for the implied equilibrium returns were developed. Including higher moments in the estimation of the implied equilibrium returns overcomes a significant practical shortcoming of the original Black-Litterman methodology, whereby normally distributed returns were assumed.

As a result of considering the higher moments present in the data, the higher moment implied equilibrium returns are better equipped at capturing the phenomenon of fat left tails inherent in many asset returns. In particular, many asset returns exhibit significant losses with an increased probability than what is expected by the normal distribution (Harvey & Siddique, 2000). Therefore, incorporating higher moments within the Black-Litterman framework will provide the investor with the necessary information in order to better capture the information available in the data as well as to avoid the large catastrophic losses that may occur. In addition, the investor may also express individual risk aversions towards negative skewness or lower kurtosis, which would inevitably result in an increased realised expected utility for the investor should preferences for higher moments exist.

In order to test the effects of including higher moments in the optimisation, a generic mean-variance-skewness utility function optimisation as well as a generic mean-variance-kurtosis optimisation was applied to a 6 asset domestic balanced fund portfolio. In particular, the effects of simultaneously varying the risk aversions to skewness and variance as well as for the simultaneous risk aversions to kurtosis and variance were examined. The results indicated that when an aversion to variance was included, the effects of increasing the risk aversion to skewness resulted in insignificant changes to the optimal asset allocation. In addition, similar results were obtained when a risk aversion to variance was included for varying the risk aversion to kurtosis; no significant changes to the optimal asset allocation were observed unless unrealistically high risk aversions to skewness and kurtosis were expressed. In summary, the inclusion of a risk aversion to variance showed to have dominated the optimal asset allocations when higher moments were considered. However, despite the disappointing results in favour of higher moments, it can be noted that a different dataset with an increased number of assets as well as the inclusion of hedge funds may have resulted in significantly different results.

While it may be attractive for investors to obtain a more complete picture of the underlying asset distribution as a result of including higher moments in the optimisation, it can be noted that a major practical shortcoming is the computing time involved when estimating the higher order moments and co-moments. In practise, investment portfolios typically consist of hundreds (if not thousands) of assets (Scherer, 2007) and the computing time involved for estimating the co-moments between every pair of assets would most likely be tremendously too long given the time constraints. In addition, investors would also need to weigh up the costs of obtaining higher moments versus the usefulness of the solution, where there may be a strong possibility that the variance-covariance matrix alone may sufficiently estimate the optimal portfolio.

Further research may involve the improved estimation of higher moments in order to significantly decrease the computing time involved. In addition, additional methodologies that consider the non-normality of asset returns, such as the mean-CVaR optimisation model can be compared to the higher moment utility functions described in this thesis and applied to the Black-Litterman model. More specifically, the higher-moment implied equilibrium returns derived in this thesis can be compared to the methodologies proposed by Giacommetti et al., (2005).

In addition to the development of robust higher moment implied equilibrium returns, the second major topic involved the development of a suite of diagnostic tools. These diagnostic tools were developed with the aim of improving the practical use of the Black-Litterman model and, more specifically, the transparency of which the optimal portfolios are formed. In particular, the diagnostic tools enable the investor to identify the individual active weight, return and tracking error percentage contributions of each particular view to the final portfolio. In addition, the contribution of the portfolio constraints and the associated distorting effects on the optimal portfolio were also quantified using these diagnostic tools.

It can be noted that very little literature exists that enables the transparent observation of the effects of the investor views and constraints on the optimal portfolio using the Black-Litterman methodology. Herold (2003) briefly describes a methodology whereby investors can observe the percentage contribution of the views to the portfolio tracking error in the unconstrained portfolio setting. Building on Herold's (2003) approach, the diagnostic tools derived in this thesis provided a novel approach whereby investors are able to transparently observe the effects of each view to the portfolio statistics within the *constrained portfolio setting*. In particular, the active weights arising as a direct result of the expressed views are identified. In this manner, the contribution of each view to the portfolio tracking error, beta, variance, skewness and kurtosis can further be derived and quantified. In addition, using the same methodology for the investor views and applying it to the Lagrange multipliers of the constraints, investors are able to directly observe exactly *how* the imposed constraints distort the expressed views as well as the individual statistics of the optimal portfolio.

Applying the diagnostic tools to an example of a portfolio consisting of the FTSE/JSE Top 40 assets and a random set of investor views expressed on the asset returns using the Black-Litterman framework, the results exposed the respective contributions of each view and the imposed constraints to the portfolio statistics. In particular, the emphasis of the example was to expose the distorting effects of imposed constraints[86] on the expressed views. Disturbing results were obtained whereby expressed investor views, implied by the constrained portfolio allocations, had been severely distorted to the extent that many of the views on assets had effectively been *reversed* (changed from positive (negative) to negative (positive) on the asset returns). More specifically, the results indicated that the imposed constraints accounted for over a third of the portfolio asset allocations, thereby distorting the views to the extent that 32 out of the 40 views had changed to a different view category. In summary, the results indicated that the imposed constraints were a major driving force behind the asset allocations, therefore dramatically reducing the investor's control whereby the views expressed within the Black-Litterman framework had a significantly diminished influence on the optimal asset allocations.

In order to mitigate the extreme distorting effects of the constraints, the view distortion percentages were examined for marginal relaxations of the long-only constraint. The results for the particular example illustrated that a 23% relaxation of the long-only constraint resulted in a significant reduction in the distortion percentages. More specifically, the 123/23 portfolio resulted in a decrease in the utility shortfall percentage from 30.15% to 4.34% relative to the unconstrained portfolio utility, therefore indicating a significant improvement in the risk and return characteristics of the portfolio. In addition, the number of view score changes (as a result of the view distortions) dramatically decreased to 3 out of the 40 expressed views. It can also be noted that the optimal 23% relaxation of the long-only constraint in order to form a 123/23 portfolio for the particular example conducted is somewhat consistent with the results obtained from Clarke *et al.,* (2004) who found that, in general, a 30% relaxation of the long-only constraint is required in order to optimally take advantage of the negative alpha opportunities present within the portfolio.

The practical implications of the results obtained in this thesis are that the expression of the investor's views and convictions are severely limited when constraints are imposed on a portfolio. In particular, the results indicated that constraints significantly limit the usefulness of the Black-Litterman model in practise as investor views cannot be clearly expressed as intended, therefore leading to portfolios that do not represent the views of the investor. It can therefore be noted that the severe distortion of the investor

[86] The imposed constraints were the 3% tracking error, zero net active position, long-only and the beta neutrality (portfolio beta equal to one) constraints.

views effectively renders the costly view generating process as a futile exercise given that the final portfolio is essentially driven by the imposed constraints.

In addition, the results illustrated that a full relaxation of the long-only constraint is not necessary as marginal relaxations of the long-only constraint resulted in substantial improvements in the consistency of the expression of the investor views as well as significant improvements in the constrained portfolio utility relative to the unconstrained portfolio case. The implications of such a result indicate that investment committees may rather be willing to accept a small relaxation of the long-only constraint as opposed to completely removing the long-only constraint in order to improve the efficiency of the investment portfolio.

Despite the notable practical advantages of the diagnostic tools, there are practical difficulties involved when using the tools to calibrate the views and constraints of the portfolio. In particular, the Lagrange multipliers arising as a result of the imposed constraints can be considered as non-linear variables that may change considerably depending on the specification of the constraints, the degree to which the constraints are binding as well as the associated interactions between the constraints. It therefore follows that the calibration of the respective views within a constrained portfolio setting, in order to achieve specified risk and return targets, may be rather challenging. More specifically, when constraints are particularly binding, small changes in the expressed views may result in large changes to the respective contributions of the individual views to the portfolio characteristics. It therefore follows that the tracking error and active weight contributions of the expressed views may be unstable and will therefore make the achievements of particular risk and return objectives rather challenging. It can be further noted that, while the phenomenon of unstable view contributions proves to be a practical *usability* weakness of the diagnostic tools, it also further illustrates the destructive effects of the imposed constraints.

Further research could involve analysing the distorting effects of the constraints specified by Regulation 28[87]. More specifically, the diagnostic tools developed in this thesis can be applied to a pension fund portfolio that is subjected to the constraints specified by Regulation 28. In addition, the impact of marginally relaxing the constraint on foreign assets from 20% to 25% can be analysed and, more specifically, it can be analysed whether the additional 5% relaxation results in any significant changes to the optimal portfolio. In addition, the diagnostic tools can be used to determine *how* the optimal Regulation 28 portfolio is influenced.

In this thesis, the Bayesian approach to asset allocation was described as a powerful approach that enables investors to construct robust portfolios that incorporate human decision making in an intuitive manner. The Black-Litterman model was used as an example of a Bayesian asset allocation model. In order to improve the practical applicability of the Black-Litterman model, various extensions and improvements

[87] Regulation 28 that gives effect to section 36(1)(BB) of the South African Pension Funds 1956, 2010

were discussed. In particular, robust higher moments were integrated into the implied equilibrium returns in order to better represent investor preferences for skewness and kurtosis of the asset return distribution. In addition, a novel approach was followed whereby the Black-Litterman model was used to measure the distorting effects of constraints on the views. The results obtained not only apply to users of the Black-Litterman model but also to any investment portfolio that is subject to constraints. The results indicated that investors need to be cautious when naively enforcing constraints in an attempt to control portfolio risk, as the constraints may effectively *increase* portfolio risk by forcing active allocations into positions whereby the investor does not exhibit significant skill. As a result of improving the transparency for which the active allocations are obtained, investors are provided with the ability to review and calibrate the respective views and constraints. In this manner, as a result of more effectively being able to observe the direct effects of human decision making on the optimal asset allocations, the relationship between the art and science of asset allocation is strengthened, therefore leading to optimal portfolios that intuitively represent the preferences and objectives of the investor.

References

Allen, D., & Powell, R. (2009). *Thoughts on Var and CVaR.* Edith Cowan University, School of Accounting, Finance and Economics, Perth.

Avramov, D., & Zhou, G. (2010). Bayesian Portfolio Analysis. *The Annual Review of Financial Economics* .

Bevan, A., & Winkelmann, K. (1998). *Using the Black-Litterman Global Asset Allocation Model: Three Years of Practical Experience.* Goldman, Sachs & Company.

Black, F., & Litterman, R. (1992). *Global Asset Allocation with Equities, Bonds and Currencies.* Fixed Income Research, Goldman, Sachs & Company.

Ceria, S., & Stubbs, R. A. (2006). *Incorporating Estimation Errors into Portfolio Selection: Robust Portfolio Construction.* Axioma.

Clarke, R. G., de Silva, H., & Sapra, S. (2004). Toward More Information-Efficient Portfolios: Relaxing the Long-Only Constraint. *The Journal of Portfolio Management* .

Da Silva, A., Lee, W., & Pornrojnangkool, B. (2009). The Black-Litterman Model for Active Portfolio Management. *The Journal of Portfolio Management* , 61-70.

Elton, E. J., Gruber, M. J., Das, S., & Hlavka, M. (1993). Efficiency with Costly Information: A Reinterpretation of Evidence from Managed Portfolios. *Review of Financial Studies* , 1-22.

Fabozzi, F. J., Focardi, S. M., & Kolm, P. N. (2006). Incorporating Trading Strategies in the Black-Litterman Framework. *The Journal of Trading* .

Fabozzi, F. J., Kolm, P. N., Pachamanova, D. A., & Focardi, S. M. (2007). *Robust Portfolio Optimization and Management.* New Jersey: Johan Wiley & Sons.

Giacommetti, R., Bertocchi, M., Rachev, S., & Fabozzi, F. (2005). Stable Distributions in the Black-Litterman Approach to Asset Allocation. *Journal of Portfolio Management* .

Harvey, C. R., Liechty, J. C., & Liechty, M. W. (2008). Bayes vs. Resampling: A Rematch. *Journal of Investment Management* .

Harvey, C., & Siddique, A. (2000). Conditional Skewness in Asset Pricing Tests. *Journal of Finance* (55), 1263 - 1295.

Herold. (2003). Portfolio Construction with Qualitative Forecasts. *The Journal of Portfolio Management* , 63.

Howson, C., & Urbach, P. (1989). *Scientific Reasoning: The Bayesian Approach.* Illinois: Open Court Publishing.

Hwang, S., & Satchell, S. E. (1999). *Modelling Emerging Market Risk Premia Using Higher Moments.* University of Cambridge, Department of Applied Economics.

Idzorek, T. (2005). *A Step-by-Step Guide to the Black-Litterman Model: Incorporating User-Specified Confidence Levels.* Zephyr Ass.

Idzorek, T. M. (2006). *Developing Robust Asset Allocations.* Chicago: Ibbotson Associates.

Johnson, G., Ericson, S., & Srimurthy, V. (2007). An Empirical Analysis of 130/30 Strategies: Domestic and International 130/30 Strategies Add Value Over Long-Only Strategies. *The Journal of Alternative Investments* , 31-42.

Jorion, P. (1986). Bayes-Stein Estimation for Portfolio Analysis. *Journal of Financial and Quantitative Analysis* , 279-292.

Kandel, S., McCulloch, R., & Stambaugh, R. F. (1995). Bayesian Inference and Portfolio Efficiency. *Review of Financial Studies , 51*, 1-53.

Kritzman, M., Myrgren, S., & Page, S. (2007). Optimal Execution for Portfolio Transitions. *Journal of Portfolio Management* .

Lamoureux, C., & Zhou, G. (1996). Temporary Components of Stock Returns: what do the data tell us? *The Review of Financial Studies , 9*, 1033-1059.

Ledoit, O., & Wolf, M. (2001). *Improved Estimation of the Covariance Matrix of Stock Returns.*

Lee, J. L. (2007). *Bayesian Reasoning Method for Intelligence Using Natural Frequencies.* Mercyhurst College, Department of Intelligence Studies, Pennsylvania.

Levy, H., & Sarnat, M. (1984). *Portfolio and Investment Selection: Theory and Practise.* Prentice Hall International.

Liew, J., & French, C. (2005). Quantitative Topics in Hedge Fund Investing . *The Journal of Portfolio Management* , 21-32.

Mankert, C. (2006). *The Black-Litterman Model: Mathematical and Behavioural Finance Approaches Towards its use in Practice.* Department of Industrial Economics and Management: The Royal Institute of Technology.

Markowitz, H. (1959). *Portfolio Selection: Efficient Diversification of Investments* (2nd Edition ed.). New York: Wiley.

Markowitz, H., & Usmen, N. (2003). Resampled Frontiers Versus Diffuse Bayes: An Experiment. *Journal of Investment Management , 1* (4), 9-25.

Martellini, L., & Ziemann, V. (2010). Improved Estimates of Higher-Order Comoments and Implications for Portfolio Selection. *The Review of Financial Studies* , 1467-1502.

Meucci, A. (2009). Enhancing the Black-Litterman and Related Approaches: Views and Stress-Test on Risk Factors. *Journal of Asset Management , 10*, 89-96.

Michaud, R. (2008). Estimation Error and Portfolio Optimization: A Resampling Solution. *Journal of Investment Management , 6* (1), 8-28.

Michaud, R. O., & Michaud, R. O. (2008). *Efficient Asset Management: A Practical Guide to Stock Portfolio Optimization* (2nd Edition ed.). Oxford University Press.

Michaud, R. O., & Michaud, R. O. (2003, June). Optimal and Investable Portfolios.

Michaud, R. (1989). The Markowitz Optimization Enigma: Is "Optimized" Optimal? *Financial Analysts Journal* , 31-42.

Michaud, R., & Michaud, R. (2004). *Resampled Efficiency vs. Bayes: Implications for Asset Management.* Boston: New Frontier Advisors.

Montier, J. (2002). *Behavioural Finance: Insights into Irrational Minds and Market.* Chichester: John Wiley & Sons.

Padberg, H. (2007). *Portfolio Management Using Black-Litterman* . The Munster School of Business Administration and Economics.

Popova, I., Morton, D. P., Popova, E., & Yau, J. (2007). *Optimizing Benchmark-Based Portfolios with Hedge Funds.* Retrieved from SSRN: http://ssrn.com/abstract=988176

Rachev, S. T., Hsu, J. S., Bagasheva, B. S., & Fabozzi, F. J. (2008). *Bayesian Methods in Finance.*

Satchell, S., & Scowcroft, A. (2000). A Demystification of the Black-Litterman Model: Managing Quantitative and Traditional Construction. *Journal of Asset Management* , 138-150.

Scherer, B. (2007). *Portfolio Construction and Risk Budgeting* (3rd Edition ed.). Risk Books.

Schneeweis, T., Crowder, G. B., & Kazemi, H. (2010). *The New Science of Asset Allocation: Risk Management in a Multi-Asset World.* New Jersey: John Wiley & Sons.

Stokes, G. (2011, April 13). *The nuts and bolts of Regulation 28*. Retrieved April 2, 2012, from FA News: http://www.fanews.co.za/article.asp?Front_Page_Features~25,Featured_Story~1147,The_nuts_and_bolts_of_Regulation_28~9360

Tversky, A., & Kahneman, D. (1984). Rational Choice and Framing of Decisions. *The Journal of Business* .

Welch, S. (2008). The Hitchhicker's Guide to Core/Satellite Investing. *The Journal of Wealth Management* , 97-111.

Witten, G., & Wilson, B. (2011). Sensitivity of Sector Allocations to the Black-Litterman Model Parameters. *Unpublished manuscript* .

Wolf, M. (2006). *Resamplng vs. Shrinkage for Benchmarked Managers.* University of Zurich, Institute for Empirical Research in Economics, Zurich.

Xiong, J. X., & Idzorek, T. (2010). *Mean-Variance Versus Mean-Conditional Value-at-Risk Optimization: The Impact of Incorporating Fat Tails and Skewness into the Asset Allocation Decision.* ibbotson.

Zlotnick, J. (1972). Bayes' Thereom for Intelligence Analysis. *Studies in Intelligence , 16* (2), 43-52.

Appendix A

A.1 The Framework of the Black-Litterman Model

The Black-Litterman model was designed to aid portfolio managers in constructing optimal portfolios that produce sensible portfolio weights which accurately reflect the manager's unique views. The Black-Litterman methodology can be describes as follows:

- The first step of the Black-Litterman model-building approach is to obtain the market equilibrium returns vector. This first input to the model is used as a neutral starting point and is extracted by a process of reverse optimization from known information.
- The investor then specifies a number of views on the market in the form of expected returns and provides a level of confidence in each view.
- These expected returns are then combined with the equilibrium returns in order to produce a set of Black-Litterman expected returns.
- The set of posterior Black-Litterman returns are then optimised in a mean-variance manner in order to obtain the new posterior recommended weights.

In this way, a recommended portfolio is created whereby asset weight deviations (known as bets) from the market weights are only taken in positions where the investor has expressed specific views. The size of each deviation in relation to the equilibrium portfolio weights is heavily dependent on the confidence level specified by the investor on the particular view.

The final Black-Litterman formula for the vector of posterior returns can be expressed as follows:

$$E[R] = [(\tau\Sigma)^{-1} + P^T\Omega^{-1}P]^{-1}[(\tau\Sigma)^{-1}\Pi + P^T\Omega^{-1}Q]$$

In formulating the Black-Litterman model, the following notation was used:

$\mathbf{E[R]}$ Is the new (posterior) vector of Black-Litterman returns (Nx1 Column Vector)

$\boldsymbol{\tau}$ Is the covariance shrinkage factor (a scalar)

$\boldsymbol{\Sigma}$ Is the covariance matrix of excess returns (NxN matrix)

$\mathbf{P}$ Is the view participation matrix that identifies the assets involved in the views (KxN matrix) for K views on N assets.

$\boldsymbol{\Omega}$ Is the diagonal covariance matrix of error terms from the expressed views. It represents the uncertainty in each view (KxK matrix)

$\boldsymbol{\Pi}$ The implied equilibrium return vector (N x 1 column vector)

Q Is the View Vector (K x 1 column Vector)

For obtaining the implied equilibrium excess returns, the following notation was used:

$\mathbf{\Pi}$ Is the Implied Excess Equilibrium Return Vector (Nx1 Column Vector)

$\boldsymbol{\lambda}$ Is the Risk Aversion Coefficient

$\mathbf{\Sigma}$ Is the Covariance Matrix of Excess Returns (NxN matrix)

$\mathbf{w_{mkt}}$ The Market Capitalization Weight vector[88] (Nx1 Column Vector)

$\boldsymbol{w_{BL}}$ Final Black-Litterman recommended posterior weight vector

In summary of the above, in the context of the Black-Litterman approach, the distribution of the market equilibrium returns (Π) is the *prior* distribution. The investor views (Q), are therefore known to be the *frequency/sampling* distribution which when combined with the prior distribution, leads to the formation of the updated *posterior* distribution of Black-Litterman returns $\mathrm{E[R]}$.

It is the above Black-Litterman formula from which the optimal portfolio weights are derived. The set of posterior Black-Litterman returns are then substituted into the *unconstrained* optimiser in order to obtain the following expression for the final set of Black-Litterman recommended portfolio weights:

$$w_{BL} = (\lambda\Sigma)^{-1}E[R]$$

The following proceeding sections describe the various inputs of the Black-Litterman model as well as how the various components of the model fit together.

A.1.2 Implied Equilibrium Returns

As discussed above, the implied equilibrium returns are used as a neutral starting point for the Black-Litterman model. Investors are assumed to have the following utility function (Padberg, 2007):

$$U = \mathrm{w^T\pi} - \frac{1}{2}\mathrm{\lambda w^T\Sigma w}$$

Maximising the utility function *U* with respect to the weight invested in each asset leads to the following vector solution of optimal portfolio weights (w^*):

$$w^* = (\lambda\Sigma)^{-1}\Pi$$

Given that we are using the market portfolio as a benchmark, it can be assumed that $w^* = \mathrm{w_{mkt}}$ and therefore, in order to obtain an expression for the implied equilibrium excess returns, the above formula can be expressed as follows:

$$\Pi = \lambda\Sigma w_{mkt}$$

88 An alternative to market capitalisation weights include presumed efficient benchmark weights.

This vector of implied equilibrium returns represents the returns estimated from the market (or benchmark) portfolio. It is composed of three sources of information: the risk aversion parameter(λ), the covariance matrix of excess returns (Σ) and the market capitalisation weights (w_{mkt}). The risk aversion parameter can be interpreted as the rate at which additional return is expected per additional unit of risk taken on. The risk aversion parameter is often represented by the Sharpe ratio (reward-to-volatility ratio) of the portfolio (Idzorek, 2005).

The vector of implied equilibrium excess returns represents the first source of information for the Black-Litterman optimisation procedure. The vector of implied equilibrium returns are assumed to be normally distributed, with a mean of "$\prod$" and a variance of "$\tau\sum$" [89].

A.1.3 Investor Views and Confidence Levels

The Black-Litterman model allows investment managers to express specific views (which differ from the implied market equilibrium returns) on particular assets in the portfolio. These views are generated through extensive research conducted by dedicated research analysts. It should be noted that views allow the investor to take active positions in terms of a tactical asset allocation procedure. The views result in deviations of the portfolio weights from the market equilibrium. If the investor does not specify a view on a particular asset, then the investor will simply hold the market portfolio.

Investor Views

Within the Black-Litterman framework, there are two methods in which the investor views can be expressed. Views can be expressed in *absolute* or *relative* terms. Each view is expressed with a corresponding level of confidence. The confidence can be expressed as the standard deviation of the expected value of the return of the view or alternatively on a 0% to 100% scale. A view expressed with a low level of confidence will prove to have a small impact on the final recommended portfolio weights. In contrast, a view that is expressed with a high level of confidence will have a considerable impact on the final portfolio weights of the model. An example of an absolute view is: "*Asset A will have an absolute excess return of 5.25% (Confidence = 20%)*" and an example of a relative view is: "*Asset A will outperform asset B by 0.75% (Confidence = 45%)*". A point to note is that in traditional portfolio mean-variance optimisers, relative views cannot be expressed (Mankert, 2006).

It should be noted that the model does not require the investor to specify views on all assets. This proves to be favourable to the investor as it is most often not possible to possess views on all the assets in the portfolio.

In most cases, the investor does not express views on an individual asset or asset class but instead expresses views on a set of linear combinations of assets (Herold, 2003). The respective weights of these

[89] The prior distribution, $\prod$ is assumed to be distributed; $\Pi \sim \mathrm{N}(\Pi; \tau\Sigma)$.

portfolios are then expressed in the matrix **P**; the matrix that identifies the assets involved in the views. The matrix **P** is composed of the respective view portfolios in which each element represents the weight of a particular asset in a certain view portfolio. Each view portfolio is represented by a row in the matrix P. The matrix P can be expressed as:

$$P = \begin{bmatrix} p_{11} & \cdots & p_{1N} \\ \vdots & \ddots & \vdots \\ p_{K1} & \cdots & p_{KN} \end{bmatrix}$$

where p_{ij} is the weight of asset *i* in view portfolio *j* for a total of *K* views on *N* assets.

In addition to the view participation matrix P, the investor views are also composed of the view vector Q[90]. Due to the uncertainty involved in the investor's views, the result is the inclusion of a random, unknown, independent and normally distributed error term vector ε [91]. It therefore follows that a view has the form $Q + \varepsilon$. The general form is given below:

$$Q + \varepsilon = \begin{bmatrix} Q_1 \\ \vdots \\ Q_K \end{bmatrix} + \begin{bmatrix} \varepsilon_1 \\ \vdots \\ \varepsilon_K \end{bmatrix}$$

The views expressed in the view vector Q are then matched to the specific assets by the view participation matrix P. In order to better understand the construction of investor views and specifically the matrix P, consider the following adapted example from Idzorek (2005):

The following three views are given on a portfolio of eight assets:

View 1: South African Property will have an absolute excess return of 5.25% (25% Confidence)

View 2: Foreign Equity will outperform South African Equity by 3.5% (50% Confidence)

View 3: SA Large Growth and SA Small Growth will outperform SA Large Value and SA Small Value by 2% (65% Confidence)

View 1 is an example of an absolute view where as views 2 and 3 are examples of relative views. It is necessary to consider the implied returns of the assets involved in the views in order to determine whether the effect of the view have a positive or negative effect[92] on the asset with regards to its respective weighting in the portfolio.

Considering View 2; if the implied equilibrium returns of Foreign Equity and South African Equity are 6.7% and 0.8% respectively, then the difference in returns is 5.9%. In contrast to the 5.9% difference in equilibrium returns, View 2 expects the difference to be smaller at 3.5%. Therefore, since $3.5\% \leq 5.9\%$,

[90] The distribution of the views are assumed to be: $Q \sim N(Q; \Omega)$.
[91] The view error term vector ε is distributed with a mean of zero and a covariance matrix Ω.
[92] A view that results in an increase in the weighting of a particular asset in a portfolio is regarded as a positive effect.

View 2 is therefore negative on Foreign Equity and is positive on South African Equity. One would therefore expect the final portfolio weights to tilt away from Foreign Equity and towards South African Equity. It can therefore be concluded that if the view is greater than the difference between the two implied equilibrium returns, the model will always tilt the portfolio towards the *outperforming* asset (Idzorek, 2005).

In this example, there are three views and eight assets. The matrix P is therefore a *3x8* matrix and Q is a *3x1* column vector. Given the above example, the matrix P can be expressed as follows:

$$P = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ -1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & .5 & -.5 & .5 & -.5 & 0 & 0 \end{bmatrix}$$

As shown above, View 1 (an absolute view) is expressed by the first row. This particular view involves only one asset and is identified as being the 7th asset in the eight asset example, as shown by the value of "1" in the 7th column. The last two rows represent views 2 and 3 respectively. It can be noted that the rows of relative views each sum to zero as the positive weightings assigned to outperforming assets cancel out the negative weightings assigned to underperforming assets.

Uncertainty of Views

As discussed above, the uncertainty of the views results in an error term vector ε, that is random, unknown, independent and normally distributed. Due to the fact that it is an error term, it therefore has a mean of zero and a covariance matrix Ω. In the hypothetical case whereby the error term takes on a value of zero, it implies that the investor is 100% confident in the expressed view.

The covariance matrix Ω is a diagonal matrix composed of zeros on all the off diagonal elements and the elements ω_j on the main diagonal. The ω_j's are the variances of the error terms (ε) and represent the uncertainty of the view. It therefore follows that the larger the variance of the error term, the greater the uncertainty associated with the view. The off diagonal elements are zero as the model assumes that the views are independent of each other. The variance of the error term is the equivalent to the magnitude of the error. In other words, it is the absolute difference of the error value and its expected value of zero[93] (Idzorek, 2005). It should be noted that the error term vector does not directly enter the Black-Litterman formula and that it is the variance of the error term that enters the formula.

The general structure of the covariance matrix is given as:

$$\Omega = \begin{bmatrix} \omega_1 & 0 & 0 \\ 0 & \ddots & 0 \\ 0 & 0 & \omega_K \end{bmatrix}$$

It should be noted that determining the diagonal elements of the parameter Ω is one of the most complicated aspects of the model. Idzorek (2005) discusses two methods of incorporating confidence levels

[93] The magnitude of the error term is derived as: $|\varepsilon_j - E(\varepsilon_j)| = |\varepsilon_j - 0| = |\varepsilon_j|$.

on the user specified views. The first of which discusses the notion of incorporating an *implied level* of confidence into each view. The second approach builds on the framework of the first and is deemed to be the most intuitive approach whereby the elements of Ω are derived by a method for controlling the tilts of the portfolio weights for which the user is able to specify a 0% to 100% confidence level.

Black and Litterman (1992) make an assumption about the value of the covariance shrinkage factor τ. Once the matrix P is defined, the variance of each individual view is calculated using $p_j\Sigma p_j^T$ where p_j is a single 1 x N row vector constituent of the matrix P corresponding to the j^{th} view. The confidence of the view is then calibrated so that $\frac{\omega}{\tau} = p_j\Sigma p_j^T$ (the variance of the j^{th} view portfolio). The covariance matrix can therefore be expressed as:

$$\Omega = \begin{bmatrix} (p_1\Sigma p_1^T) * \tau & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & (p_k\Sigma p_K^T) * \tau \end{bmatrix}$$

In this way, the value of τ becomes irrelevant and changing it leaves the final expected returns unaffected.

A.1.4 Covariance Shrinkage Factor

Together with the view covariance matrix (Ω), the value of the covariance shrinkage factor (τ) are the most abstract and difficult to specify parameters in the Black-Litterman model. As explained by Padberg (2007), the covariance shrinkage factor plays a major role in the Black-Litterman model in minimising estimation error and is a key advantage over the Markowitz model.

Practically, the value of the covariance shrinkage factor can be interpreted as the remaining uncertainty in the estimate of the implied equilibrium returns (Π) (Rachevet et. al, 2008). In other words, it can be described as a scaling parameter of the estimated implied equilibrium returns (Idzorek, 2005). The excess return covariance matrix (Σ) is multiplied by the covariance shrinkage factor in order to reduce the model's sensitivity and improve its stability. Since $\Pi \sim N(\Pi\ ;\ \tau\Sigma)$, it should be noted that increasing the value of τ increases the uncertainty (variance) of the equilibrium returns and places a greater weighting on the investor views.

Unfortunately, in literature, there is very little guidance as well as many ideas that contrast each other on choosing the value for τ. A few guidelines and explanations for setting the value of τ are presented below:

- Satchell and Scowcroft (2000) suggest that the value for τ be set to a value close to or less than one. To justify such a value, they provide an advanced mathematical discussion of a conditional value for τ.
- In contrast to the above, Black and Litterman (1992) suggest a value close to zero. This is due to the fact that the uncertainty in the mean (implied returns) is less than the uncertainty (variance) of the

return. To further justify such a small value for τ, one would expect the implied equilibrium returns to be less volatile than historical returns (Idzorek, 2005).

- Lee (2000), sets the value of τ to be between 0.01 and 0.05 and then calibrates the model according to a target level of tracking error. This proves to be very useful as many active portfolio managers are subject to a target level of tracking error for which they are expected to maintain.
- Blamont and Firoozye (2003) interpret the value of τ to be approximately inversely proportional to the number of oberservations[94]. Therefore, as the sample size (period of observation) increases, the estimation uncertainty decreases (τ decreases).

As discussed in the previous section, the easiest and most popular method of dealing with the parameter τ is to make an assumption and set the ratio $\frac{\omega}{\tau}$ equal to the implied variance of the view portfolio $(p_j \Sigma p_j^T)$. In this way, the value of τ becomes irrelevant.

A.2 Conclusion

The above sections have discussed how the individual components of the Black-Litterman model are determined as well as various illustrations of how the Black-Litterman model can be applied in practice. As discussed above, within the Black-Litterman framework, the investor views are combined with the equilibrium returns of the market. The size of the bets taken on by the investor depends on three factors: the level of confidence assigned to each particular view, the size of the bet taken on and size of the specified covariance shrinkage factor. It should be noted that views which differ dramatically from the market consensus contribute to larger bets taken. Views that are assigned a confidence level of zero result in the investor simply holding the market portfolio. It therefore follows that when no views are expressed, the Black-Litterman model suggests holding the market portfolio.

[94] If the number of observations is represented by T, then $\tau = \frac{1}{T}$.

Appendix B

B.1 Robust Optimisation Methodologies

In order to account for estimation error present in the returns, there are various optimisation methodologies aimed at achieving portfolios that account for the inherent risk of estimation error. The following sections describe the respective merits of two additional robust optimisation methodologies and briefly compare them to the Bayesian methodology.

B.1.1 Portfolio Resampling

Traditional portfolio optimisation is based on optimising a particular set of inputs based on point estimations. For example, means, variances and correlations are estimated based on one realisation of return history. However, as discussed in previous sections, estimates of parameter values based on a finite sample are bound to exhibit a significant degree of estimation error. Portfolio resampling is a methodology aimed at producing an optimiser that reduces estimation error within the estimated parameters and therefore results in portfolio allocations that are less sensitive to small changes in the inputs.

Portfolio resampling reduces estimation error by repeatedly resampling from the observed data. For each resampled data set, an optimal portfolio is constructed by feeding the resampled data into a mean-variance optimiser. The process is repeated many times until a sufficient number of optimal resampled portfolios have been obtained. Once a sufficient number and variety of optimal resampled portfolios are obtained, the resampled portfolios are then averaged in order to find the optimal portfolio. The central idea behind finding the average of all the resampled portfolios is that all the positive and negative estimation errors will cancel each other out. Michaud (2008) states that because a number of simulated portfolios are averaged, the resulting averaged resampled portfolio will result in superior performance on average.

Michaud (2008) finds that the averaged resampled portfolio is typically more diversified than the single portfolio obtained from the original sample estimates. The increase in diversity therefore eliminates the extreme allocations that are typical of the traditional optimised portfolios. It can be noted that the efficient frontier obtained from a mean-variance optimisation ranges from the minimum-variance portfolio to the maximum risk portfolio. Naturally, these two portfolios are heavily concentrated in assets that exhibit the lowest and highest risk respectively. In contrast, the average resampled efficient frontier plots below and to the right of the original Markowitz efficient frontier; this is due to the fact that the resampled portfolio translates estimation error into increased portfolio risk.

The benefits of using portfolio resampling can be illustrated by plotting the asset allocations for a portfolio over the minimum-maximum risk spectrum and compare the results to classical mean-variance portfolio allocations. Michaud (2008) illustrates that the resampled portfolio is significantly more diversified and results in smooth transition changes in the allocations for increases in portfolio risk. In contrast, the classical mean-variance portfolio is highly concentrated in a few assets and dramatic changes in the asset allocations can be observed for small changes in the levels of risk. The resampled portfolio can therefore be described as a robust portfolio optimisation model as it has shown to result in portfolios that are well diversified and are stable over the risk-return space, which is highly favourable amongst investment practitioners.

Despite the favourable characteristics of portfolio resampling, various authors (Scherer, 2007; Fabozzi *et al.,* 2007; Wolf, 2006) criticise the resampling algorithm as being a heuristic that has no economic justification as it is not clear why averaging over resampled portfolios would be optimal in any way. In addition, it is criticised for its relatively poor performance in the absence of investment constraints, whereby portfolio resampling has resulted in portfolios that do not offer any significant improvement over the classical mean-variance optimisation. However, in the presence of the long-only constraint, the portfolio resampling algorithm has shown to result in portfolios that are more diversified and exhibit superior out-of-sample performance. As described by Wolf (2006) and Fabozzi et al. (2007), the increased diversification in the presence of the long-only constraint can be attributed to the averaging process, whereby all assets in the portfolio will most likely receive a positive nonzero weight. A resampled asset can either be more or less favourable than its original value, but when the long-only constraint is applied, the asset weight can either be zero or positive. Averaging over a sequence of zero or positive resampled weights will undoubtedly lead to a final positive allocation.

An additional criticism of portfolio resampling is the existence of a non-concave segment of the efficient frontier. A non-concave segment violates the assumption of optimality as it would theoretically be possible to construct a portfolio consisting of a linear combination of assets with the same level of risk with a greater level of return. Michaud (2008) states that this phenomenon is associated with a possible gap or shortage in the estimates corresponding to the particular region of the efficient frontier. Non-concavity may also be an indication that the simulation has not converged and that more simulations may be required.

A final criticism proposed by Fabozzi et al. (2007), is that the final averaged portfolio may not adhere to the imposed investment constraints. This is as a result of many "constraint satisfied" portfolios being averaged into one portfolio that may not necessarily satisfy the investment constraints. In order to rectify the problem, Michaud (2008) suggests that a two-stage optimisation be performed, whereby the first stage optimisation aims to achieve a theoretically optimal portfolio and the second stage aims to optimise the portfolio in order to satisfy the imposed constraints. Michaud (2008) further stresses that an understanding of the non-ad hoc Bayesian role of constraints is essential to avoiding inadvertent biases as well as benefiting from the implementation of constraints in portfolio optimisation.

It must be noted that additional methods, such as Bayesian estimation and shrinkage methods for improving parameter estimates, are not mutually exclusive to portfolio resampling (Michaud & Michaud, 2004). In fact, a combination of Bayesian improved estimates and the portfolio resampling methodology, implemented in the correct manner, may result in an effective portfolio management framework that performs well in practice.

Bayesian Methods versus Resampling

As illustrated in previous sections, there are many competing methodologies for developing optimal portfolios that have shown to produce impressive out-of-sample performance, the most remarkable of which are the portfolio resampling algorithm and Bayesian portfolio optimisation methods. Portfolio resampling aims to reduce estimation error via repeatedly resampling from the return distribution and to average the optimal portfolio weights. Bayesian portfolio optimisation aims to reduce estimation error by imposing a prior distribution on the parameters, thereby treating the sample parameters as random variables that occur within a specific confidence interval as opposed to point estimates.

Harvey *et al.,* (2008) explicitly tests the out-of-sample performance of the two competing methods using a series of simulated returns. In their testing methodology, they specify 10 true parameter sets from a multivariate normal density. Each true parameter set summarises the distribution of eight asset returns. For each of the 10 true parameter sets, 100 histories, each with 216 simulated observations were generated. The two optimisation methodologies were then applied to each of the 100 histories and sets of optimal portfolio weights were applied. Three levels of risk aversion were tested for each of the 10 truth sets, namely $\lambda = 0.5, 1, 2$. The best methodology for a particular history is the one that recommends portfolio weights that result in the greatest expected utility using the original parameters from the respective truth set.

For the above testing methodology, Bayesian methods resulted in the greatest expected utility for 9 out of the 10 truth sets for $\lambda = 0.5$. However, for greater levels of risk aversion $(\lambda = 1\ and\ 2)$, portfolio resampling results in the greatest utility for 7 and 8 of the truth sets respectively. It was therefore shown that the two methodologies yielded similar results whereby portfolio resampling performed better for higher levels of risk aversion and Bayesian methods performs better for lower levels of risk aversion.

Harvey *et al.,* (2008) further test the two methodologies assuming a shorter investment horizon whereby future returns are assumed to be consistent with the past return history. The testing methodology is almost identical to the previous one with the exception that returns are drawn from the predictive density. The results show that Bayesian methods dominate portfolio resampling for all 10 true parameter sets. However, it should be noted that the portfolio resampling methodology produced almost identical mean returns but at a higher level of overall portfolio variance. The increasing portfolio variance for increasing risk aversion resulted in portfolio resampling producing a consistently falling expected utility as compared to the Bayesian methods.

The respective tests show that the portfolio resampling and Bayesian methods have their particular merits and both perform comparatively well in most investment situations. The optimal choice of portfolio optimisation methodology simply depends on the current situation. Harvey *et al.,* (2008) state that if the future distribution is known to be similar to the past, then the Bayesian methodology will always result in superior performance. However, if there is uncertainty regarding the current return distribution (predictive distribution is different from the current distribution), then portfolio resampling will result in superior performance, especially for high levels of risk aversion. It is therefore easy to see that the Bayesian methodology relies on a well specified prior distribution[95] in order to perform well in practice, while the portfolio resampling methodology approach exhibits significant robustness towards distributional uncertainty (Harvey, Liechty, & Liechty, 2008). Therefore, in the absence of knowledge of an appropriate prior distribution, portfolio resampling methods should be used.

B.1.2 Robust Portfolio Optimisation

It must be noted that the term “robust optimisation” can be misleading and ambiguous. In literature, it has been used to describe several different concepts involving the procedure of obtaining portfolios that are “robust”. As previously described, there are many different types of robust portfolio optimisation methodologies available, such as portfolio resampling, shrinkage methods and Bayesian model averaging. Therefore, it must be noted that the particular “Robust Optimisation” methodology in this section, refers to the objective of providing a good outcome even if the worst parameter specification becomes true. In other words; maximising the worst case return and obtaining a portfolio that performs better under many scenarios within a range of uncertain returns (Fabozzi, Kolm, Pachamanova, & Focardi, 2007). Robust portfolio optimisation is generally concerned of whether allocation decisions are adequate even if the input parameters are poorly estimated.

Despite the improved out-of-sample results, it can be noted that Robust portfolio optimisation simply interpolates between the minimum variance portfolio and the maximum Sharpe ratio portfolio, depending on the level of estimation error aversion specified. Robust portfolio optimisation can therefore be viewed as a shrinkage estimator that that combines the minimum variance portfolio with an uncertain investment portfolio. Scherer (2007) states that this is no different from a typical Bayesian shrinkage estimators, as presented by Jorion (1986), which shrinks the estimated expected return towards the minimum variance portfolio. He further states that the return adjustment process is largely in-transparent relative to Bayesian alternatives.

[95] The process of selecting an appropriate prior distribution may be difficult as there are many possible prior distributions available. The uncertainty regarding the most appropriate prior distribution is referred to as distributional uncertainty.