

Risk Budgeting within an Asset Liability Modelling (ALM)
Framework, using Mean-Variance Optimisation

Prasheem Singh

November 1, 2006

Acknowledgements

I would like to express my sincere appreciation and gratitude to the following:

Andrew van Biljon for his unfailing encouragement, support, guidance and invaluable input throughout the process of completing my dissertation.

Jarred Glansbeck for his time, support, input, guidance and all round encouragement.

My parents, brother, close family members and friends for your boundless belief in me and your never-ending interest, support, encouragement and egging on throughout.

RisCura for their financial and data assistance, and also to all my colleagues at RisCura for their assistance, patience and encouragement throughout the period of my study.

The lecturers and other staff at UCT involved with the Mathematics of Finance programme for all your help and patience over the years and my fellow classmates for their encouragement.

Contents

1	Introduction	8
2	Overview of the Pension Fund Investment Process	11
2.1	The Asset Liability Modelling (ALM) Process	12
2.1.1	Effect of Diversification	13
2.1.2	Asset Allocation	14
2.1.3	Investment Strategies	17
2.1.4	Manager Selection	21
2.2	The Emergence of Risk Budgeting	22
3	Risk Budgeting in an ALM Framework	24
3.1	Traditional Mean Variance Optimisation	24
3.2	The Pension Fund Optimisation Problem	26
3.3	Solving the Optimisation Problem Numerically	27
3.4	Risk Budgeting and Manager Allocation	30
3.5	Inputs used in the Optimisation Process	32
3.5.1	Covariance Matrix Estimation	33
3.5.2	Expected Return Estimation	43
4	Application of Risk Budgeting in a Pension Fund Environment	50
4.1	Description of the Process and the Data used	50
4.2	Analysis of Risk Budgeting Results	54
4.2.1	The Asset Allocation Results	55
4.2.2	Balanced Mandate Results	57
4.2.3	Specialist Mandate Results	69
4.3	Discussion of Results	85

5	Conclusion	90
6	Glossary	92
A	Analysis of inputs used in the implementation of the optimisation	95
A.1	The Optimisation Technique	95
A.2	The Covariance Matrix	96
A.2.1	The Sample Covariance Matrix and Average Return case output	96
A.2.2	The Covariance Matrix output comparisons	97
A.3	Expected Returns	102
A.3.1	The Average Return output comparisons	104

List of Figures

2.1	Example of an Efficient Frontier	15
2.2	Example of an Relative Efficient Frontier	16
2.3	Schematic of the Investment Decision Process	17
4.1	Differences in Cashflows between True Liabilities and Investment Strategy	53
4.2	Efficient Frontier and Optimal Asset Allocation Solutions for Asset Allocation	56
4.3	Efficient Frontier Comparison of the Balanced Mandate Situations	59
4.4	Situation A's Optimal Solutions and their Active Risk and Return Component Breakdown	61
4.5	Situation B's Optimal Solutions and their Active Risk and Return Component Breakdown	63
4.6	Situation C's Optimal Solutions and their Active Risk and Return Component Breakdown	65
4.7	Situation D's Optimal Solutions and their Active Risk and Return Component Breakdown	68
4.8	Situation E's Optimal Solutions and their Active Risk and Return Component Breakdown	70
4.9	Efficient Frontier Comparison of the Specialist Mandate Situations	72
4.10	Situation 1's Optimal Solutions and their Active Risk and Return Component Breakdown	75
4.11	Situation 2's Optimal Solutions and their Active Risk and Return Component Breakdown	77
4.12	Situation 3's Optimal Solutions and their Active Risk and Return Component Breakdown	79
4.13	Situation 2's Optimal Manager Allocations relative to Situation 3's	81
4.14	Situation 4's Optimal Solutions and their Active Risk and Return Component Breakdown	83
4.15	Situation 5's Optimal Solutions and their Active Risk and Return Component Breakdown	86
A.1	Efficient Frontier and Optimal Asset Allocation Solutions for Scenario 1	97
A.2	Comparison of the Efficient Frontiers and Optimal Asset Allocation Solutions of the Scenario 2 cases	99
A.3	Weightings for Scenario 1 and the Scenario 2 cases.	101
A.4	Comparison of the Efficient Frontiers and Optimal Asset Allocations solutions for Scenario 4	103
A.5	Scenarios monthly weights relative to the Equity benchmark returns (ALSI50R)	105
A.6	Scenarios monthly weights relative to the Bond benchmark returns (ALBI)	106

A.7	Scenarios monthly weights relative to the Cash benchmark returns (AFMM)	107
A.8	Comparison of the Efficient Frontiers and Optimal Asset Allocations Solutions for Scenario 2e108	
A.9	Comparison of the Efficient Frontiers and Optimal Asset Allocations Solutions for Scenario 3	111

List of Tables

4.1	Expected Return estimates for Asset Allocation	55
4.2	Black-Litterman Views applied for Asset Allocation	55
4.3	Covariance Matrix estimates for Asset Allocation	55
4.4	The Fund's chosen Optimal Asset Class Allocation & Its Expected Risk and Return Characteristics	57
4.5	Litterman Views taken for all the Balanced Mandate Situations	58
4.6	Manager and Manager Benchmark Asset Allocations for Situation A	58
4.7	Expected Return estimates for Situation A	60
4.8	Covariance Matrix estimates for Situation A	60
4.9	Expected Return estimates for Situation B	62
4.10	Covariance Matrix estimates for Situation B	62
4.11	Manager and Manager Benchmark Asset Allocations for Situation C	64
4.12	Expected Return estimates for Situation C	64
4.13	Covariance Matrix estimates for Situation C	64
4.14	Manager and Manager Benchmark Asset Allocations for Situation D	66
4.15	Expected Return estimates for Situation D	66
4.16	Covariance Matrix estimates for Situation D	67
4.17	Manager and Manager Benchmark Asset Allocations for Situation E	67
4.18	Expected Return estimates for Situation E	69
4.19	Covariance Matrix estimates for Situation D	69
4.20	Black-Litterman Views for Situation 1, Situation2 and Situation 3	73
4.21	Expected Return estimates for Situation 1	74
4.22	Covariance Matrix estimates for Situation 1 ($\times 10^{-6}$)	74
4.23	Expected Return estimates for Situation 2	76
4.24	Covariance Matrix estimates for Situation 2 ($\times 10^{-6}$)	76
4.25	Minimum Asset Class Allocations imposed on Situation 3's Optimisation	78

4.26	Expected Return estimates for Situation 3	80
4.27	Covariance Matrix estimates for Situation 3 ($\times 10^{-6}$)	80
4.28	Situation 4's Manager Selection choices per Asset Class	82
4.29	Black-Litterman Views for Situation 4	82
4.30	Expected Return estimates for Situation 4	82
4.31	Covariance Matrix estimates for Situation 4 ($\times 10^{-6}$)	84
4.32	Situation 5's Manager Selection choice per Asset Class	84
4.33	Black-Litterman Views for Situation 5	84
4.34	Expected Return estimates for Situation 5	85
4.35	Covariance Matrix estimates for Situation 5 ($\times 10^{-6}$)	85
A.1	Breakdown of Inputs used in Optimisation Scenarios	95
A.2	Table of estimated Expected Return Values for Scenario 1	96
A.3	The estimated Covariance Matrix for Scenario 1	96
A.4	Comparison of Optimal Allocations and Expected Returns at 3% Risk level	98
A.5	Annualised Risk (Volatility) of the asset class benchmarks	100
A.6	Comparison of the Expected Return estimates for the Scenario 2 cases	104
A.7	Scenario 2 Weights at yearly end points of the analysis period	104
A.8	Black-Litterman Views for Scenario 3a	109
A.9	Black-Litterman Views for Scenario 3b	110
A.10	Expected Returns estimates for Scenario 3, adjusted by B-L Views	110

Chapter 1

Introduction

For a pension fund, risk is the possibility that the fund's accumulated assets will be insufficient to cover its future liabilities. This means that the fund will be unable to provide its members with adequate pensions on their retirement. To prevent this, trustees need to properly quantify their fund's liabilities in order to understand its funding status in terms of its current assets. If the trustees want to have any chance to improve a fund's funding status (particularly if it is under-funded), they will need to quantify an acceptable level of risk they are willing to take to try and boost their assets, even though there is simultaneously a possibility that they could weaken this funding position.

A risk budget is the quantification of an acceptable amount of risk that a pension fund is willing to take in its investment decisions, without significantly sacrificing its ability to meet its current and future liabilities. In other words, it indicates the degree to which the fund returns could deviate from the growth in the liabilities (or some benchmark representative of these liabilities), without sacrificing the fund objectives.

The choice of the risk budget determines how much risk the trustees are willing to take, which is influenced by their level of risk aversion and indeed the circumstances of the fund. The risk budget therefore determines the fund's investment strategy type as well as its implementation. For instance, if the fund chooses to have a zero risk budget then the strategy that needs to be implemented is a full immunization strategy, which can only be implemented if the fund is fully funded, or indeed if the members expectations are such that no further exceptional growth is required. This strategy allows for the exact replication of the cash flows required to match the fund's future liabilities. If a low to medium risk budget of about 2% is chosen, the fund would typically invest in lower risk assets like bonds and inflation-linked bonds. However if the fund chooses a large risk budget, say 10%, this means that the fund is willing to invest in assets that are quite risky compared to the profile of their liabilities, for example equity products.

At each stage of implementing the investment strategy, risk is introduced. Some of the implementation stages are:

- Choice of strategic benchmark to model liabilities:
- Choice of strategic benchmark to manage the assets:
- Choice of investment style, e.g. passive vs. active management:
- Choice of asset managers, and the effects of blending their investment styles.

The combined risk introduced by all these implementation decisions contributes to the overall risk profile of the fund and should lie within the fund's defined risk budget. The awareness and choice of a fund's risk budget is therefore an important process for trustees as it determines the potential risks to which the fund is exposed.

Once the trustees have established their risk budget it is an equally important process for them to allocate this budget across their various investments so as to optimally utilise the risk that they are willing to take on, and in so doing to achieve the highest possible return for the fund.

The aim of this thesis is to give an overview of the various steps and decisions in the risk budgeting process that trustees are confronted by, to ultimately allocate pension fund assets in the most optimal manner to investment managers, while functioning within the required risk budget. This process requires that the allocation of mandates to investment managers be carried out such that the mismatch and stock selection risk components introduced into the Fund's strategy is optimised so that the risk budget may be fully utilised. This will require a correct allocation to various mandates and investment styles.

Chapter 2 provides a brief overview of the asset-liability modelling (ALM) process so as to set out the various steps that lead to the ultimate allocation of pension fund assets to investment managers. It touches on the underlying concepts of investments by briefly describing the concept of diversification and how optimisation is useful in determining the pension funds optimal mix of asset classes. The different investment strategy options available to pension funds are also discussed along with a brief explanation of the manager selection procedure. Since there is risk introduced at every step in the process which contributes to the risk budget, risk budgeting is introduced as being integral to optimally allocating this risk budget.

Chapter 3 looks at the theory and technicalities involved in optimisation procedures. It also discusses the main inputs required and the various methods by which these inputs may be estimated. Traditional mean-variance optimisation is explained and it is shown that optimisation problems within a pension fund framework cannot be solved analytically given the required constraints, but rather that numerical solutions need be sought. Some of these solutions are also discussed and with the optimal asset allocation taken care of, risk budgeting and its role in manager allocation is discussed. The objective function for the optimisation is amended to include the mismatch and stock selection components of risk and return. Looking at the optimisation in this manner, allows the fund to identify mismatch and to decide on its manager allocation so as to optimally allocate its assets to achieve the best return. The rest of the chapter discusses the two main inputs into the optimisation procedure, i.e. the covariance matrix and expected return estimators, and discusses the pros and cons of some of the estimation methods used to estimate these inputs.

Chapter 4 describes the application of the risk budgeting process. It provides a description of the data used and analyses performed at the various steps in the risk budgeting process. The application takes us through the optimisation procedure to establish a pension funds optimal asset allocation and subsequently how that asset allocation is used to optimally allocate the funds assets to the relevant asset managers. Various situations within both the balanced and specialist mandate scenarios are analysed and discussed. The last section of the chapter discusses the results obtained.

Chapter 5 summarises the scope of the project and discusses some ways of potentially enhancing the analyses so that they may provide more realistic results that could be implemented in real life scenarios.

Preliminary analyses were done to establish the impact of different covariance matrix and expected return estimators on mean variance optimisation. This initial analysis was important to be able to establish the most appropriate estimators that should be used in the application of the risk budgeting process. The

various analyses performed and a discussion of their results have been included in the appendix to provide a better understanding of the rationale behind the chosen input estimators for the asset allocation and manager allocation optimisations.

Chapter 2

Overview of the Pension Fund Investment Process

Since no person can work continuously until the day that he or she dies, he or she eventually reaches an age where it is either impossible or unfeasible for him or her to work any longer. People in this situation, however, still most probably have many years left to live, and with this their many needs, wants and responsibilities remain a factor of their lives. They still need, for example, to eat and have a place to live, etc and therefore require some sort of income or savings to be able to afford and sustain these prerequisites of life.

In an ideal world, it is expected that everyone would work and save enough money to be able to provide for that time in their lives when they are unable to earn their living. In reality, however, many of us either neglect to or do not have the discipline to save enough or the skill required to ensure that we invest our savings in such a way so as to achieve the best possible return with the lowest possible risk. It is from this need that the Pension Fund industry was born.

A Pension Fund receives contributions (typically some percentage of one's salary while employed) from both, the member of the Pension Fund as well as from the Company at which one is employed. The percentage contributions made are usually governed by any relevant legislation and also the tax implications for both the employer and employee play a major role in its determination. Both the employee and employer contributions together form the total contribution for the member of the Fund. In addition, the advantage of investing in a Pension Fund is that an individual's relatively small contribution every month is combined with that of many others to form a much larger pool of funds for investment. This larger pool will be able to be spread across more investment opportunities and thus each Pension Fund member benefits from the size of the Fund's investments through cost savings and diversification.

The contributions to the Fund, together with any returns achieved thereon make up the assets of the Fund. These assets then need to be invested by the Pension Fund Trustees in order to achieve an adequate return to pay out pensions to the members on their retirement. The pension payouts that a Fund makes are known as the liabilities of the Fund. These liabilities are however not a fixed value and are expressed in "real terms", for example, either they are a function of an employee's salary near retirement or are linked to a Consumer Price Index (CPI), or are limited to the accumulated value of an investment account in the case of defined contribution funds. Also these liabilities do not all occur at the same time, but are spread

out over a time period of several years. The exact timing of these liabilities is dependent upon the scheme, the Fund members' mortality and morbidity dynamics, i.e. when they retire, when they die, if they have dependants, if they become disabled, etc.

The entity bearing the risk of the Fund not achieving its liabilities differs within the different types of Pension Funds, i.e. Defined Benefit (DB) Funds and Defined Contribution (DC) Funds. A Defined Benefit Fund is, as the name suggests, a fund where the benefits that the members receive are specified or defined up front, and is typically some specified percentage of the members' final salary. This percentage is not dependent on the Fund's investment performance and the company or plan sponsor thus bears the risk of the Fund not generating sufficient income to meet these liabilities (i.e. the Fund deficit). The company will have to ensure that the members receive their specified benefits irrespective of any potential loss to the Fund. On the other hand, surpluses are generated when the Fund's income exceeds its liabilities. However the ownership of these surpluses is still surrounded by much debate and regulatory changes with respect to acceptable distribution thereof.

The Defined Contribution Fund, as the name suggests, receives contributions, from its members, that are some specified percentage of their salary before retirement. This proportion is seldom changed except, where that decision is voluntary, and the payout on retirement is mainly dependent on the Fund's investment performance on these assets. The member receives a lump sum payout on retirement of their contributions (employer and employee) plus any investment performance achieved thereon. With this lump sum the member can then purchase some sort of annuity to provide for the duration of their retirement. In this manner the member bears the risk of the potential investment losses and also enjoys the benefit of potential gains. However the major risk is whether or not this lump sum will be sufficient to provide an income, in retirement, which will keep the member at the same standard of living to which the member has become accustomed.

Even though the DB fund has a well defined and fixed liability structure, and the DC fund a more floating liability structure, which is determined by investment performance, the goals of both fund types are the same. The Funds' ultimate aim is to provide adequate pensions for the members, so that their standards of living are not detrimentally altered after retirement.

2.1 The Asset Liability Modelling (ALM) Process

Asset Liability Modelling (ALM) tries to ensure that an investment structure today will meet the liabilities of the future. However, given the "real term" nature of pension fund liabilities the appropriate assets required to match these liabilities may not exist.

For example, one of the factors that determine the liabilities is salary, and salary increases. Although salaries generally rise with inflation, they can increase at faster rate than inflation due to promotions etc. Although these additional amounts are subjective and very specific to the individual members, they can be stochastically modelled in order to determine the probable liabilities for the fund. Unfortunately, there are no assets available that match the cashflow characteristics of salary increases. The assets that have the closest cashflow characteristics to these liabilities are inflation linked bonds, as they provide at least inflation protection. Also this increase has an impact on contribution levels and so significantly affects future income receipts into the Fund. In combination the impact of future salary, especially for the young, can be marked.

The coupon value of an Inflation linked bond is calculated by applying its coupon rate to its principal amount that continually grows by an inflation proxy, in most cases, CPI. Also the final principal amount paid on maturity of the bond, is the initial principal value grown by inflation over the term to maturity of the bond. Since these bonds remove inflation risk as a factor, the yields of these bonds are lower than those obtained for conventional bonds. This is acceptable to investors because they are willing to “pay the premium” for the guarantee of inflation protection, in keeping with the risk-reward tradeoff concept.

However, within South Africa particularly, there are currently very few Inflation linked bonds available in the market. This selection, particularly for larger Funds, thus provides a limited set of term to maturity options for investments and may not cover the whole spectrum of maturity dates required. In addition the inclusion of these assets in investment strategies may prove to be too costly relative to the benefit derived from reduced portfolio risk, thus reducing the Fund’s overall return potential. It may then be more feasible to incorporate the higher risk asset classes, i.e. equity, conventional bonds, etc within the portfolio strategy in order to optimize the Fund’s return potential.

With Pension Fund ALM it is important that the investments of the Fund achieve the correct mix of asset classes given the asset classes’ individual risk and return characteristics, to meet the unique characteristics of the liabilities of the Fund. In addition, ALM aims to minimise the risk involved in not meeting these liabilities.

2.1.1 Effect of Diversification

The minimisation (controlling) of risk is important in lowering (managing) the chances of potentially not meeting the Fund’s liabilities, and thus having a Fund shortfall (i.e. when the liabilities of the Fund exceed the assets), and is achieved somewhat through investment diversification. In general, diversification is the concept of “not putting all your eggs in one basket”.

The total risk of a number of combined assets is not calculated by merely adding together the individual asset risks, because there can be correlations between the assets which influences the total risk. Total portfolio risk rather depends on how the returns of the individual assets move together or are correlated with each other, and relative to the liability. Portfolio risk is generally obtained by calculating the standard deviation of the portfolio returns. To combine standard deviations one cannot simply add them together, but one needs to first calculate the variance of each asset’s returns (square of the standard deviations) and the covariances of returns between pairs of assets. As an example, the formula for combining two assets, A and B, is given by:

$$\sigma_p^2 = w_a^2 \sigma_a^2 + w_b^2 \sigma_b^2 + 2.w_a.w_b.\sigma_{ab} \quad (2.1)$$

Where w_a, w_b are the weights used to combine the assets, A and B, σ_p is the total portfolio risk, σ_a, σ_b are the risks of the assets, A and B respectively, σ_{ab} is the covariance between asset A and asset B.

From equation 2.1 it can be seen that one can achieve a potentially lower total portfolio variance by combining lesser-correlated or uncorrelated investments with higher individual variances. For example, assume asset A has a variance of 20% and asset B has a variance of 10% with a covariance of 4%, the

variance of the portfolio combining 60% of asset A and 40% of asset B is given by:

$$\begin{aligned}\sigma_p^2 &= 0.6^2 * 0.2 + 0.4^2 * 0.1 + 2 * 0.6 * 0.4 * (0.04) \\ &= 0.08992 \\ &= 8.99\%\end{aligned}$$

We see that the total portfolio variance at about 9% is lower than either individual assets' variances of 20% and 10% for asset A and asset B respectively. By combining the assets we were able to lower the total portfolio risk and thus achieve some level of diversification.

The effect of diversification is also crucial to pension fund investments in order to minimise their total portfolio risk, but probably more important than just lowering their total portfolio risk is lowering their relative risk, i.e. the risk of the portfolio relative to their liabilities. The risk of these liabilities is often encapsulated or estimated by the risk implicit in some well diversified benchmark. Equation 2.2 demonstrates the effect of diversification on a relative basis.

$$\begin{aligned}\sigma_{r,t}^2 &= \frac{1}{n-1} \sum_{i=1}^n [(R_{P_i} - R_{B_i}) - (\overline{R_P} - \overline{R_B})]^2 \\ &= \frac{1}{n-1} \sum_{i=1}^n [(R_{P_i} - R_{B_i})^2 - 2(R_{P_i} - R_{B_i})(\overline{R_P} - \overline{R_B}) + (\overline{R_P} - \overline{R_B})^2] \\ &= \frac{1}{n-1} \sum_{i=1}^n [(R_{P_i} - \overline{R_P})^2 + (R_{B_i} - \overline{R_B})^2 - 2R_{P_i}R_{B_i} + 2R_{P_i}\overline{R_B} + 2R_{B_i}\overline{R_P} - 2\overline{R_P}\overline{R_B}] \\ &= \sigma_P^2 + \sigma_B^2 - \frac{1}{n-1} \sum_{i=1}^n [2R_{P_i}R_{B_i} + 2R_{P_i}\overline{R_B} + 2R_{B_i}\overline{R_P} - 2\overline{R_P}\overline{R_B}] \\ &= \sigma_P^2 + \sigma_B^2 - \frac{1}{n-1} \sum_{i=1}^n [2(R_{P_i} - \overline{R_P})(R_{B_i} - \overline{R_B})] \\ \sigma_{r,t}^2 &= \sigma_P^2 + \sigma_B^2 - 2\sigma_{R_P R_B}\end{aligned}\tag{2.2}$$

Where $\sigma_{r,t}$ is the fund's relative risk, R_{P_i} and R_{B_i} are returns of each asset i in the fund's portfolio (P) and the fund's benchmark (B), σ_P and σ_B is the risk of the fund's portfolio and the fund's benchmark respectively, $\sigma_{R_P R_B}$ is the covariance between portfolio P and benchmark B .

From equation 2.2 it is evident that should the Fund's investment portfolio be perfectly correlated to its benchmark, the Fund's relative risk would be zero. This implies that the Fund could expect to match its liabilities exactly. Since in practice investment portfolios are very seldom, if at all, perfectly correlated to the benchmark representing their liabilities, it is important for the Fund to try to ensure that its investment portfolio is highly correlated to its benchmark so as to reduce its relative risk to a level at which the Fund is comfortable. The Fund can also further reduce its relative risk by ensuring that its investment portfolio is well diversified, thereby lowering the portfolio risk component of the equation and therefore lowering the total relative risk.

2.1.2 Asset Allocation

For each total portfolio risk level there are various combinations of asset classes that can be invested in. However only one of these combinations, per total risk level, is optimal, i.e. for a particular total risk level there is one combination of asset classes that achieves the highest expected return. The optimal

asset allocation is established by optimizing the expected (i.e. future forecasted) returns of the asset class benchmarks, given the risk tolerance level and using a relevant covariance matrix (to estimate the risk of potential returns around the forecast). These optimal portfolios are thus the most efficient of the many combinations that exist for a particular risk level and lie on an Efficient Frontier in risk-return space. The efficient frontier is concave since the optimal portfolios' risk and return characteristics change in a non-linear manner as its asset class weightings are changed [2]. An example of an Efficient Frontier can be seen in Figure 2.1.

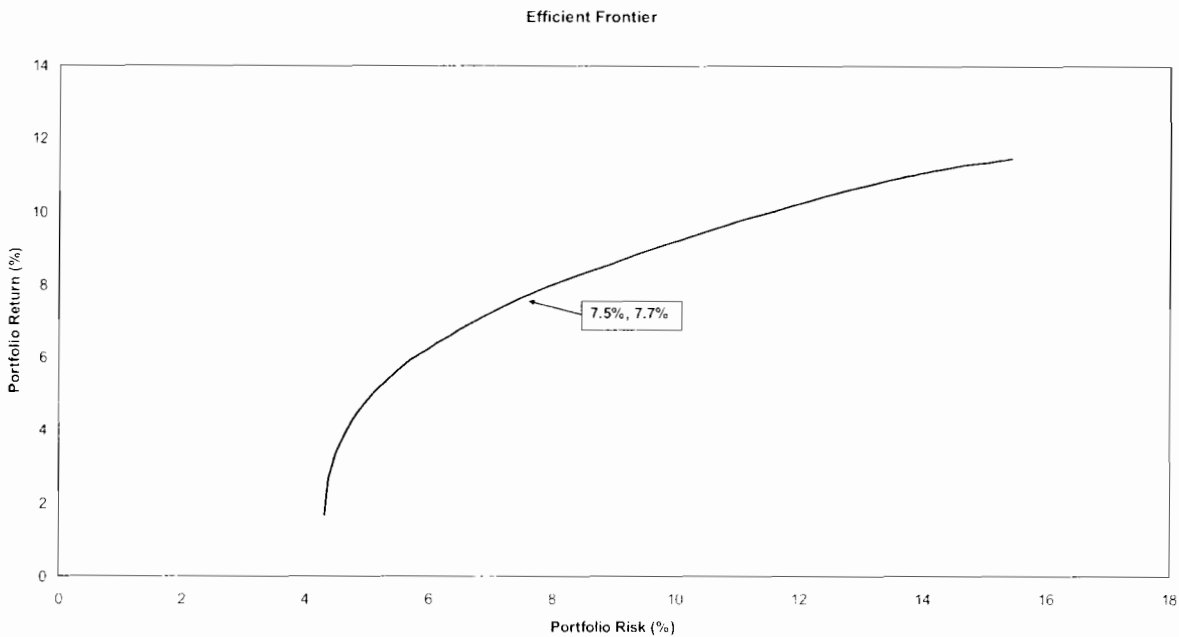


Figure 2.1: Example of an Efficient Frontier

As can be seen from the diagram, for a particular level of risk, for example 7.5%, the most optimal combination of asset classes lies on the curve with an expected return of about 7.7%. This whole concept underlies Modern Portfolio Theory introduced by Harry Markowitz in 1952 [17].

This concept is a part of Asset modelling; in particular the above Markowitz Theory is based on the Mean-Variance approach. This name came about because it utilizes returns, that are usually assumed to be normally distributed, which can thus be modelled and described by, their 1st and 2nd moments, i.e. the mean and variance respectively. Some other asset modelling techniques are the Stochastic or Monte Carlo methods.

Liability Modelling is also an important (and integrated) part of the ALM process. Its purpose is to model how future liabilities may change based on assumptions made with respect to inflation, mortality, morbidity, salary increases, growth in membership, etc. This can be done deterministically or stochastically.

Deterministic modelling assumes that for a given group of members with their associated mortality and morbidity statistics, a certain proportion must succumb to these events in a year on year time periods. Stochastic modelling acknowledges the fact that these mortality and morbidity statistics are merely an indication of the probability of such events occurring and is subsequently treated as such.

The components of both the asset modelling and the liability modelling are utilised during the ALM process.

The optimal assets selected during the asset modelling process need to meet the liability structure obtained from the Liability modelling. Also the minimum amount of risk required to meet the liabilities needs to be incorporated in the optimal asset allocation decision. This forms an integral input into the ALM process of establishing the optimal asset allocation strategy to meet the expected liabilities of the Fund.

The choice of optimal asset allocation infers the Funds risk tolerance level, as the allocation will provide information about the amount of risk that the Fund is comfortable with. This acceptable amount of risk could be referred to as the Fund's risk budget. The size of the risk budget represents a trade-off between the amount of risk that the Fund needs to achieve its objectives, i.e. cover its liabilities, and the amount of risk it can tolerate.

Like the efficient frontier shown earlier, a similar efficient frontier can be generated in order to establish the optimal portfolio asset allocations relative to that of the benchmark, representing the liabilities. This relative efficient frontier is shown in Figure 2.2, with the axes reflecting relative risk and relative returns instead.

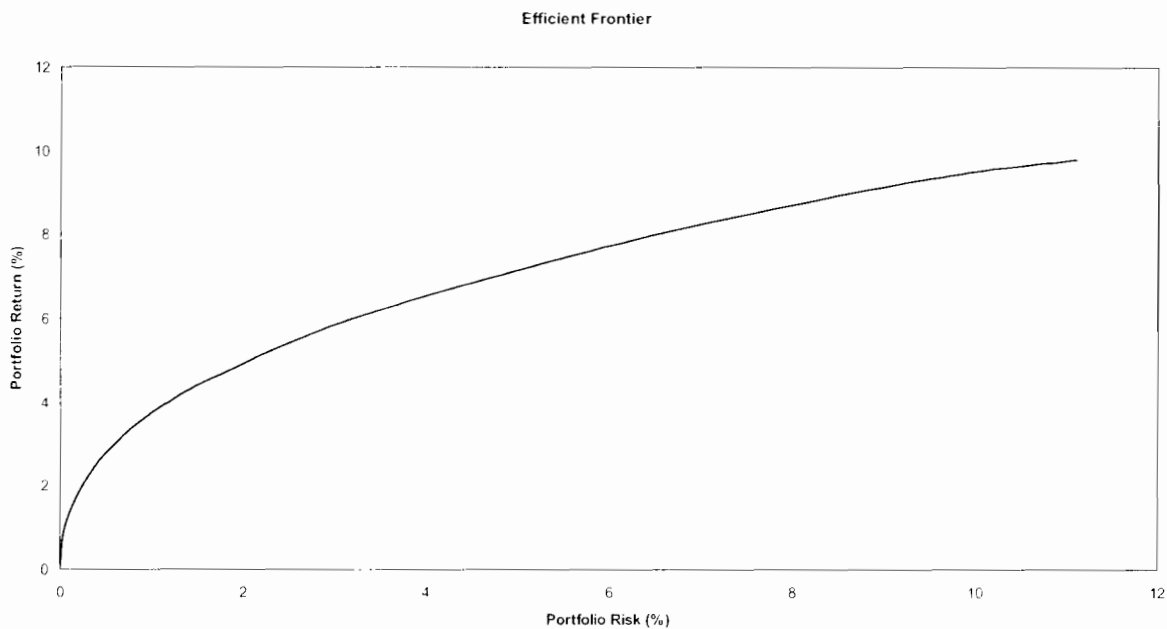


Figure 2.2: Example of an Relative Efficient Frontier

Once the optimal asset allocation to the various asset classes is established, the Fund needs to decide how best to achieve that allocation. This is addressed by deciding what the best investment strategy is, and coupled with that, the relevant choice of asset manager within that strategy.

Asset managers are, as their title suggests, institutions that undertake to manage the assets of investors (e.g. Pension Funds, Insurance Companies, Unit Trusts and some Private individuals) according to pre-specified guidelines (i.e. a mandate). Since all Pension Funds do not necessarily have the required investment skill and knowledge to carry out the investments in-house, this task is "out-sourced" to the asset managers for an agreed upon fee.

2.1.3 Investment Strategies

Figure 2.3 below illustrates the breakdown of the investment strategies that are available to investors and the management styles that asset managers follow. A Pension Fund would need to decide which strategy to follow, i.e. the strategy that best fulfilled the Fund's investment needs within its specified constraints.

Based on the Fund's objectives, the Fund would need to decide whether it would want to implement an active management or a passive management strategy, and then whether to implement that strategy using balanced or specialist mandates.

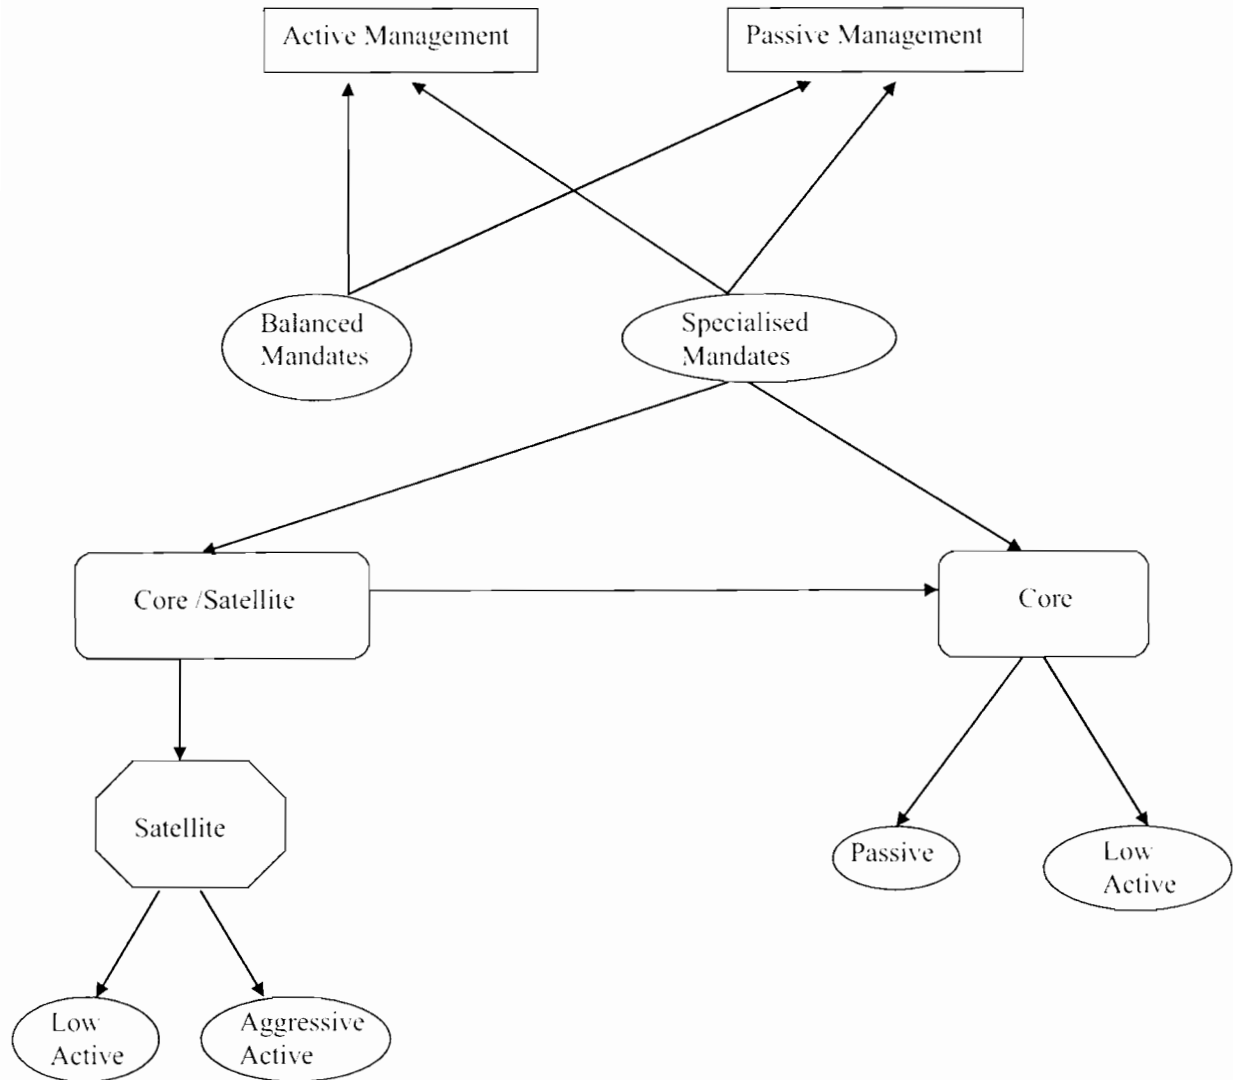


Figure 2.3: Schematic of the Investment Decision Process

Passive Management / Indexation

The Passive investment style is one where the assets are selected to replicate the asset class index that is used for performance measurement. A common name associated with this investment style is "Tracker fund", due to the nature of the investment process. These portfolios should be highly diversified so that no extra effort or resources are used in attempting to improve the portfolios performance through asset

analysis. As passive management generates very little to no risk over and above the benchmark risk, this strategy forms the foundation for the Fund's investments.

There are at least three ways in which to passively manage a portfolio. The first involves the full replication of the benchmark or index being managed against. This means that the manager's portfolio would include each and every asset, in exactly the same proportion, as that of the benchmark. This will ensure that the manager's portfolio achieves exactly the same return, pre cost, as that of the benchmark, and would track the benchmark exactly (i.e. zero active risk). This style of managing a portfolio can be very costly, as some of the assets in the benchmark may be very illiquid, and thus the manager would need to pay a higher liquidity premium in order to buy these assets.

The second way aims to generate the same return as the benchmark, but the portfolios need not necessarily hold every asset included in the benchmark. This method may utilize statistical methods, including optimization, to establish the combination of assets that exhibits the same characteristics as the benchmark. The combination of assets would make up the assets of the manager's portfolio and would typically comprise fewer assets than the total number of assets in the benchmark. These portfolios would be expected to generate returns similar to those of the benchmarks. Since fewer shares are purchased, this method reduces the costs associated with replicating the entire benchmark.

The third method is known as stratification. This method aims to have the same industry allocation as the benchmark, i.e. it maintains the same weights as the benchmark in the various industries. However within each industry, similar techniques to the second method may be utilised to establish the combination of assets within each industry which best matches the benchmark's industry returns. By replicating each industry's return and then weighting them by similar, or the same, weights as the benchmark, the portfolio thereby replicates the entire benchmark's return.

The exposures of passively managed portfolios are kept in line with those of the benchmarks by rebalancing the portfolios. Typically for the full replication portfolios, rebalancing is usually only required where there are corporate actions which affect the issued shares of a company in the benchmark or if the benchmark constituents change, while for optimised portfolios the rebalancing is usually due to asset price changes. There are also a variety of rebalancing frequency options which a Fund may choose to best suit its requirements, for example a Fund may decide to only rebalance monthly or even only quarterly etc. With the passive investment style there is generally very asset trading in the portfolio except where trading is required to bring the portfolio in line with its benchmark, as a result of a rebalancing criteria. The returns are thus based on the buy and hold method, i.e. the returns are calculated using the portfolio's asset exposures at the beginning of the period. These exposures are essentially kept constant or unchanged (except where corporate actions are involved) for the duration of the period over which the return is calculated. Benchmarks are thus passive portfolios and their returns are calculated as described above. Passively managed portfolios thus aim to provide returns in line with the specified benchmark, and promise no lower or higher return above that. The passive manager thus does not expose the investment to substantial active risk (i.e. risk over and above that of the benchmark)

Active Management

The Active investment style, on the other hand, attempts to improve performance over and above that of a benchmark by either identifying mis-priced assets or by timing the performance of asset classes. This style involves taking bets on particular assets, i.e. having higher exposures than the benchmark in assets

in which one has conviction in their superior characteristics and return potential and lower exposures in those assets that are felt not to provide as good value. This management style thus promises active return, however it also introduces active risk into the investment.

Within active management, there are a number of different types of active managers. The pure active managers invest as described above and could take bets that vary from neutral to the extreme positive or negative positions. Among these pure active managers there is a further distinction made, depending on the extent of the managers extreme negative bets. The most extreme negative bet would be one where the manager was allowed to go short on the share exposures, i.e. the manager was allowed to hold a negative exposure to a share, even if the share was not in the benchmark universe. Some hedge fund managers fall into this category, that could possibly be considered the purest form of active managers. The other pure active managers have limited extreme negative bets, in that their negative bet is limited to the size of the benchmarks exposure. This type of pure active manager is typically referred to as a "long-only" manager. Given the extreme nature of pure active managers' bets, they tend to have high active risks or tracking errors. The Core active managers tend to be more conservative than their counterparts, the pure active managers. While they adopt the similar investment method, the bets they take would be somewhat less than the extreme bets and would keep the active risk within a moderate risk range. Another type of active manager would be the Enhanced Index manager. These managers aim to slightly out-perform the benchmarks they manage against without taking too much risk. They take relatively small bets and have correspondingly very low tracking errors.

All these active managers can be further categorized according to how investment decisions are arrived at within the asset management institutions. These investment decisions can be taken based on quantitative analysis or on fundamental analysis. Quantitative managers usually make their investment decisions based on the outcome of some quantitative model or analysis. This could be, for example, that the shares that obtain a high ranking on a quantitative model would be bought while those with a low ranking would be sold. These rankings could be the exposures of assets to each factor of some factor model, where the specific factor in question is thought to be the one driving performance. In general most enhanced index managers tend to make use of quantitative analysis to make investment decisions, as they need to exploit the perceived opportunities while keeping their active risk low. In some cases quantitative decisions are overridden or checked against the outcome of fundamental analysis. This is due to the fact that quantitative models can not always account for all market dynamics, as models are not all encapsulating and usually contain assumptions that may not be valid for all realistic situations.

Those managers that base their decisions on fundamental analysis do in-depth analyses on the companies in which they intend to invest. The investment decisions are based purely on the merits researched and observed within the financial statements of the companies and economic and market trends. Some of the more fundamental managers ignore the benchmarks as investment guidelines and base their decisions solely on their fundamental analysis, while others use the benchmark as an initial starting point, and then take bets driven by the outcome of their fundamental analysis process.

Balanced Mandates

Balanced mandates are such that the chosen asset manager manages the Fund's assets across all the specified asset classes. As can be seen from Figure 2.3, balanced mandates can adopt either the active or the passive management style. In addition to the underlying processes governing these different styles, there are further responsibilities placed on the manager.

With an actively managed balanced mandate the asset class allocation within the portfolio can vary, relative to the prescribed benchmark's asset class allocation, in keeping with the manager's unique economic outlook and views for the various asset classes. With passively managed balanced mandates the manager is compelled to "track" the prescribed benchmark's asset class allocations, i.e. the manager cannot take bets around the benchmark's asset class allocation.

Specialised Mandates

This type of mandate, allocates the Fund's assets to asset managers with the proviso that they may only manage the money with assets within one particular asset class or sector within an asset class. For example, the Fund may allocate an equity mandate to manager A and a bond mandate to manager B. Both these managers then need to be fully invested in assets within their prescribed asset classes with maybe some cash for working capital. As seen from the diagram above, specialized mandates can also adopt either the active or passive management style or even in some cases the strategy could combine both management styles.

Core These mandates tend to be more conservative and some may even be passively managed. However, in some cases they can also be actively managed for Funds that are willing to take slightly more risk than a purely passive strategy. In this case the manager would deviate from the asset exposures of the benchmark, but very little.

Core/Satellite This strategy combines the Core strategy, as described above, with the Satellite approach. The Core portion of this strategy forms the foundation of the Funds investments. Here, as described above, the Fund has a very low risk tolerance level over and above the benchmark's risk, and would typically require a conservatively managed or enhanced index actively managed investment designed to track or outperform the fund benchmark. This component is also low cost. The Satellite portion of this strategy, the higher cost component, is then the focussed portion of assets on which the Fund has a higher risk tolerance level. The fund would be willing to take more risk on this portion of the assets provided that the expected returns adequately compensated for the additional risk that would be taken on. This thinking is in keeping with the Risk/Return tradeoff. The assets would thus be managed aggressively, in that large bets would be taken around the benchmark's exposures, with the expectation of achieving higher returns than the benchmark. The satellite portfolios also allow the Fund to neutralise any style biases that may have arisen during the manager selection process. Since some managers have clear biases toward a particular style (e.g. Growth), the fund would want to neutralize that implicit bet by appointing another manager with an opposing style bias (e.g. Value), so that both ends of the style spectrum are covered, and neutralised relative to the benchmark (unless of course there is clear favouring of either style by the Trustees).

Once the strategy has been decided on, the Fund then needs to decide how to allocate the funds amongst a variety of asset managers.

The process of allocating assets to asset managers requires some further research and analysis. It is not just a matter of allocating the assets to the manager that negotiated the lowest fee, but rather a comprehensive Manager Selection process needs to be carried out.

2.1.4 Manager Selection

As discussed in Modern Investment Management [15], and touched on briefly here, the first step in this process would be to

Determine the universe of managers This would involve compiling a complete list of potential asset managers that the Fund would probably approach to manage its assets.

This universe for some mandates is vast, making it impractical and time consuming to thoroughly analyse each manager on the list, in order to ultimately arrive at the appropriate managers for the Fund. The Fund may thus want to narrow down the universe by doing some quick, easily accessible research into what the various managers have to offer. This quick analysis may involve looking at performance and risk statistics that would be easily available from the various asset manager surveys or from the managers themselves. This type of screening process would be largely quantitative and is described in more detail below.

Quantitative screens

These screens look at a variety of factors including style consistency, risk-adjusted performance over rolling periods of time, performance consistency and risk analysis, to name a few. Two key factors that identify successful managers would be

- Superior risk-adjusted performance in various market environments and over the market cycle.
- Consistent results relative to an appropriate benchmark.

Comparing and contrasting all the assessed managers' information, both historical and the current, will allow the Fund to make informed choices. Even though past performance is not an indicator of future performance, the analysis will provide some information into the managers' reliability to meet targets and provide an out-performance. It may thus also serve as a "historical benchmark" or "report-card" for future manager performance and positioning.

Once this initial process manages to narrow down the list, further in-depth analysis would be required, since looking at only these quantitative screens fails to shed light on other equally important factors within the asset management organization. For example, even though a manager may show an exemplary risk-adjusted performance record historically, the pattern may have been changing significantly recently for reasons not apparent from the quantitative screens. Also, given that specialist skills are sometimes hidden by broad mandates, looking at the total returns of these mandates and deciding to eliminate managers purely as a result of the quantitative analyses, may lead to the incorrect elimination of potentially star managers. An attribution analysis and performance breakdown pre-screen may help to avoid this problem.

Qualitative screens

The qualitative screening process is generally conducted on-site and would thus require that the manager-selection team pay a visit to the asset managers' offices. Qualitative screens can be used to gain information about for example the experience of the team; the length of time the team was at the organization, the depth of the organizations resources, investment philosophy and process, ownership structure, ability to retain staff, staff movements and recent changes in processes which could affect the consistency of performance. Having discussions or conducting interviews with the investment team and management will be helpful in

establishing this kind of information. In addition, being on the premises will allow the manager-selection team to observe the organization's physical setting, corporate culture, and the quality of interpersonal relations and communication. This insight into the asset manager organizations would be very useful in establishing the stability of the investment team and the organisation as a whole and would greatly assist in spotting potential pitfalls.

All asset management organisations have some positives and some negatives that emerge from the screening processes and the various analyses conducted. As a result, the screening processes and the analyses need to be combined and looked at in totality, to provide an overall picture of each investigated organisation. Then the information gathered for each organization needs to be filtered and compared with that of other asset management organizations within its appropriate peer group. The manager-selection process would then allow the Fund to identify future pitfalls associated with the individual asset management organizations and thus establish whether it is worth investing with them.

2.2 The Emergence of Risk Budgeting

ALM and management in a nutshell involves:

- Assessing the liabilities of a Fund and utilizing the assets in an optimal manner so as to meet these liabilities by.
- Establishing an optimal asset allocation for a Fund based on expected risk and return numbers for the passive asset class benchmarks.
- Deciding on the investment strategy that best suits the Funds Risk/Return profile.
- Based on the asset allocation and the investment strategy, choosing Asset managers, via an in depth asset manager selection process, that best suits the investment requirements of the fund.
- Continuous monitoring of the managers' performance, risk statistics and the portfolios positioning relative to the specified guidelines and benchmarks.
- Continuous monitoring and assessment of the overall Fund's Investment strategy and whether it is still viable for the Fund.
- Implementing changes where necessary.

These optimal asset allocations obtained from the ALM process, as discussed earlier, are based on the expected risk and return numbers of asset class benchmarks relative to liabilities. Therefore, the asset allocation to a specific asset class has implicitly committed those asset managers mandated to that asset class to maintain reasonably similar or identical characteristics to that benchmark on average over time [22]. Investments that mimic or that aim to be identical to a specified benchmark is obtained through passive management or pure index tracking.

However, most investors, and thus Pension Funds, would like to obtain higher returns, rather than settling for just the average or benchmark returns, as this allows contributions to be reduced or pensions increased. Higher than benchmark returns are usually expected to be obtained through active investment management. As a result many Funds, after the asset allocation policy has been established, allocate mandates to active managers in an effort to obtain active return.

With the inclusion of active managers within the investment strategy, the active asset managers will deviate from the specified benchmark in terms of the assets held in their portfolios, as per the active management style discussed earlier. They will thus potentially provide a positive active return (i.e. out-perform the benchmark), however there is also a possibility that they may provide a negative active return (i.e. under-perform the benchmark). This possibility of out/under performance is due to the additional risk over and above that of the benchmark (active risk) that an active asset manager exposes the Fund to [22].

The Funds allocation of mandates to both active and passive managers alters the Risk/Return profile of the asset allocation policy that was established through the initial Asset Liability Modelling exercise. The active return and active risk introduced by the active managers upsets the optimal allocation already established rendering it potentially sub-optimal.

This inconsistency with ALM processes was the driving force behind the emergence of Risk Budgeting. A risk budget is simply a particular allocation of portfolio risk. An optimal risk budget is an allocation of risk such that the optimal asset allocation is still the most efficient for that particular level of risk, i.e. it lies on the efficient frontier. The risk budgeting process is the process of finding an optimal risk budget [15].

Chapter 3

Risk Budgeting in an ALM Framework

3.1 Traditional Mean Variance Optimisation

Harry Markowitz's work in 1952 was highly influential to the development of portfolio theory and, at the time, provided a breakthrough in developing a model that fully accounted for the covariances between asset returns.

Markowitz's paper entitled "Portfolio Selection" [17] dealt with the process of establishing an optimal portfolio based on some assumptions about relevant future performances and a need by the investor to control the standard deviation of returns or volatility of the portfolios performance. The optimization made use of the assumption that returns are normally distributed and could thus be characterized by its first two moments, i.e. the mean and variance of returns. This concept gave rise to the name "Mean-Variance" optimization that is now widely recognized.

A key to Markowitz's optimisation process is that expected returns are desirable to investors while expected variance is undesirable. The investor would thus invest in a portfolio that yielded the best return for a given level of risk, or similarly, a portfolio which yielded the lowest risk for a given level of return.

For an N asset portfolio, x_i represents the weight of asset i such that

$$x^T = (x_1, \dots, x_N) \tag{3.1}$$

such that $\sum_{i=1}^N x_i = 1$

The portfolio's return, at a given time t , is then given by

$$R_x(t) = \sum_{i=1}^N x_i R_i(t) \tag{3.2}$$

The mean and variance of R_x are given by the following equations

$$\mu_x = E[R_x] = \sum_{i=1}^N x_i E[R_i] = \sum_{i=1}^N x_i \mu_i \quad (3.3)$$

And

$$\sigma_x^2 = Var(R_x) = \sum_{i=1}^N \sum_{j=1}^N x_i x_j Cov(R_i R_j) = \sum_{i=1}^N \sum_{j=1}^N x_i \sigma_{ij} x_j \quad (3.4)$$

Where σ_{ij} is the covariance of asset i and asset j

Since risks are not additive but rather combine in a way that depends on how returns move together, it can be seen that the spreading of investments across less correlated assets tends to reduce overall portfolio risk, while never completely eliminating it. An example of this was demonstrated earlier for a two asset case.

Markowitz's method of optimization has proved popular and for the optimization investors have two objectives, while taking cognizance of the portfolio weight constraints: [19]

- Maximisation of the expected value of the portfolio return
- Minimisation of portfolio risk

However a portfolio with maximum expected return is not necessarily the one with minimum variance. An investor can gain expected return by taking on variance or reduce variance by giving up expected return. According to the investor's individual preferences a weighting can thus be applied to the investor's conflicting objectives. This weighting is called the investor's risk aversion represented by parameter in the equation below. Thus, according to Markowitz the optimization problem to be solved is:

$$\begin{aligned} \max_{x \in \mathbb{R}^N} \{ \mu_x - \lambda \sigma_x^2 \} \\ \text{such that } \sum_{i=1}^N x_i = 1 \end{aligned} \quad (3.5)$$

The objective function above is quadratic in the weights, thus this problem is referred to as a quadratic optimization problem.

As noted in Panjer [19], the Mean-Variance optimization method has several important advantages. The following comments list some these.

1. Only the first and second moments of the asset returns are required to proceed with the optimization.
2. As discussed and as can be seen from the objective function, the preferences of the investor are taken into account simply by determining the investor's risk aversion factor.
3. If no additional constraints are imposed in the optimization problem, then the optimal weights can be calculated analytically. From this fact the most popular results in modern portfolio theory were derived, such as the structure of the efficient frontier, mutual fund theorems, and the capital-asset-pricing model.
4. Even if additional linear constraints are imposed on the weights, the optimization problem remains quadratic convex, and numerical algorithms can be used to obtain optimal solutions.

However, the Markowitz Mean-Variance approach's assumption that returns are normally distributed and that the portfolio returns can be characterised purely by the first two moments of the distribution is a severe limitation to its practicality in some situations. This method cannot account for the effect that options and other derivatives have on the distribution of returns. The inclusion of derivatives in a portfolio changes the characteristics of the return's distribution, by possibly introducing fatter tailed distributions that cannot adequately be characterised by just the mean and variance of the distribution. The inability of Markowitz's method to cope with this has been one of its major criticisms.

The optimization method is discussed in more detail in the following sections.

3.2 The Pension Fund Optimisation Problem

A pension fund's main objective is to always ensure that its assets are sufficient to cover its liabilities. As such, it is imperative for a pension fund to always approach optimisation on a relative basis and to optimise their investment portfolios relative to their liabilities, rather than just optimising their portfolios in isolation. This allows the Fund trustees to invest in portfolios which are expected to yield the highest return for a particular level of relative risk, at which the Fund is comfortable.

For illustration purposes it has been assumed here that the Fund trustees are concerned about volatility of the portfolios returns and thus the risk is defined by the standard deviation, which has a high degree of similarity to the goal protecting against the loss of capital or more accurately the protection of return stability. To establish the portfolios that are optimal or efficient, i.e. have the maximum return for a given level of risk, the following equation is used.

$$\begin{aligned} \text{Max} U &= \mu^T x - \lambda x^T \Sigma x & (3.6) \\ \text{subject to } e^T x &= 1 \end{aligned}$$

Where μ is a vector of expected returns of assets or asset classes, etc. Σ is the covariance matrix of the returns with the elements given by σ_{ij} , λ represents the risk aversion factor and, e is a vector of 1's.

Varying the weight vector, x , maximizes the above equation. The weights represent the proportion of the total or the exposure to the asset or asset class, etc. Note that equation 3.6 is just equation 3.5 written in matrix notation.

The optimisation constraint above is commonly referred to as the Budget Constraint, as it ensures that all the invested funds or assets are allocated. The risk-aversion factor is a parameter that characterizes how much risk one is willing to have in the investment portfolio. By varying this parameter, you vary the total portfolio risk. A lower risk-aversion factor, will give rise to an optimal portfolio with lower portfolio risk, while a higher risk-aversion factor, would give rise to potentially higher return optimal portfolios due to the higher portfolio risk.

To calculate the set of efficient portfolios we assume:

1. The covariance matrix is positive definite, so that the inverse of the matrix can be calculated.
2. The vectors e, μ are linearly independent. This is required since the optimization will not work if all the assets have the same return and risk.

On an efficient frontier, the portfolio with the lowest variance on the efficient part of the curve (i.e. the

turning point of the graph or when the risk aversion factor (λ) is at its maximum) is referred to as the minimum variance portfolio. The weights of which can be calculated analytically, as shown below [19].

$$\min = \frac{1}{c^T \Sigma^{-1} c} \Sigma^{-1} c \quad (3.7)$$

For all other λ values, the solution to the optimization problem is:

$$x^* = \frac{1}{c^T \Sigma^{-1} c} \Sigma^{-1} c + \frac{1}{\sqrt{\lambda}} \left(\Sigma^{-1} \mu - \frac{c^T \Sigma^{-1} \mu}{c^T \Sigma^{-1} c} \Sigma^{-1} c \right) \quad (3.8)$$

or simplified

$$x^* = x^{\min} + \frac{1}{\sqrt{\lambda}} z^* \quad (3.9)$$

$$\text{where } z^* = \left(\Sigma^{-1} \mu - \frac{c^T \Sigma^{-1} \mu}{c^T \Sigma^{-1} c} \Sigma^{-1} c \right)$$

Efficient portfolios are thus expressed in the form shown in equation 3.9. The covariance matrix is required to be positive definite because the solution to the optimization problem uses the inverse of the covariance matrix (Σ^{-1}). Above, x^{\min} is the minimum variance portfolio, which depends on the covariance matrix but not on μ , and z^* depends on Σ and μ and has the property $\sum_{i=1}^N z_i^* = 0$. Thus z^* is a self-financing portfolio in that the long positions, i.e. where the weights are positive, are financed by corresponding short positions, i.e. where the weights are negative.

The above optimization problem thus provides a solution, which allows investments to have long and short positions in its assets. Generally investment mandates, particularly within the Pension Fund environment do not allow such freedom with its assets. Pension funds are subject to legislated rules and regulations in an attempt to safeguard the pensions of its members. As such, the optimization problems for Pension funds include further constraints. An example of such constraints is: the exposure of any asset invested in has to be greater than or equal to zero (i.e. no short selling is permitted). The "no short selling" constraint is standard across most Funds, while other additional constraints may vary from Fund to Fund, in keeping with their unique requirements. These may include maximum or minimum limits on the weights of individual assets, etc.

The solution in equation 3.9 only allows for optimization with respect to the budget constraint. To try and obtain an analytic solution where the budget and the additional constraints are taken in account is usually difficult. Thus the Funds have had to rely on answers obtained numerically, by utilising numerical optimization methods.

3.3 Solving the Optimisation Problem Numerically

Using mathematical packages like Matlab or more commonly Excel Solver, numerical solutions are sought for the constrained optimization problem that Pensions Funds are often faced with when trying to establish a suitable Investment strategy.

The optimization equation 3.6, with the additional constraints, as described above, can be written as:

$$\begin{aligned} \text{Max } U &= \mu^T x - \lambda x^T \Sigma x & (3.10) \\ \text{subject to } & c^T x = 1 \text{ and} \\ & a_i \leq x_i \leq b_i \text{ for all } i=1 \text{ to } N & (3.11) \end{aligned}$$

Where μ is a vector of expected returns, x is the vector of weights, λ represents the risk aversion factor, c is a vector of 1's, Σ is the covariance matrix of returns, and a_i represents the minimum and b_i the maximum allowed exposure or weight to asset 'i' for the relevant constrained assets. a_i must always be positive.

The x_i 's can be varied numerically to obtain the maximum utility for a particular risk-aversion level. The numerically obtained weights established for each asset or asset class, per risk aversion factor, make up the optimal asset allocation choices for the Fund (the portfolios that lie on the efficient frontier). The Fund would then be required to establish its asset allocation related to the risk aversion factor that it is comfortable with.

Because it is difficult to quantify ones risk aversion, often risk-adjusted performance ratios are utilised to assist in establishing this. These ratios in general provide information about the risk-adjusted return that a Fund can expect (based on the expected returns and risk numbers). These measures are thus often more tangible, in that one has better feel for the range of values that one would be comfortable with. Some common examples of these ratios are the Information ratio and Sharpe Ratio [9].

The Information (IR) and Sharpe ratio (SR) are defined as the expected active return relative to active risk. The only difference is that the Information ratio looks at the active return and risk relative to the Fund's benchmark, while the Sharpe ratio uses a risk free rate as the benchmark. The return achieved over and above the return of the risk-free rate, specifically, is commonly referred to as the excess return [6]. Thus the Sharpe ratio is more commonly defined as the expected excess return relative to excess risk. The measures can be represented as described below:

$$\text{IR} = \frac{(\mu - \mu_b)^T x'}{x'^T \Sigma' x'} \quad (3.12)$$

Where μ is a vector of Fund expected returns, μ_b is a vector of Benchmark expected returns, x' is the vector of active weights (i.e. the bet), Σ' is the covariance matrix of active returns.

Similarly,

$$\text{SR} = \frac{(\mu - \mu_f)^T x}{x^T \Sigma^* x} \quad (3.13)$$

Where μ is a vector of Fund expected returns, μ_f a vector of the risk-free rate expected returns, x is the vector of weights, Σ^* is the covariance matrix of excess returns.

It can be deduced from the definition of the ratios, that the higher the final ratio values the better. This can be achieved by either increasing the numerator, which represents the expected returns, or by decreasing the denominator, which represents the expected risk. Thus, for example, a Fund that achieved a high return in a moderate risk environment would be preferred to a Fund that has achieved a similarly high return in a high-risk environment, and this would be evident by its higher risk-adjusted performance.

The risk ratios can be utilized in two ways in the optimization. First, is to find the optimal weights that give the maximum Information or Sharpe Ratio. For example if the objective were to maximise the Sharpe

ratio, the objective function would be, as per the definition in equation 3.13:

$$\begin{aligned} \text{Max}U &= \frac{(\mu - \mu_f)^T x}{x^T \Sigma^* x} & (3.14) \\ \text{subject to } e^T x &= 1, a_i \leq x_i \leq b_i \text{ for all } i = 1 \dots N \end{aligned}$$

Where μ is a vector of Fund expected returns, μ_f a vector of the risk-free rate expected returns, x is the vector of weights, Σ^* is the covariance matrix of excess returns, and a_i represents the minimum and b_i the maximum allowed exposure or weight to asset 'i' for the relevant constrained assets. a_i must always be positive.

This optimization problem is equivalent to solving for the optimal weights for the different risk-aversion levels, as per equation 3.10, and then choosing those weights and thus the corresponding risk-aversion level that gives the maximum Sharpe ratio. In doing this, the Fund had made an implicit risk-aversion level choice without needing to actually quantify the risk-aversion level with which it is comfortable. The optimization would yield a single solution providing the weights that maximize the Sharpe Ratio.

Second is if the Fund is targeting a specific value for the risk-adjusted performance ratios. There will however be several optimal portfolios that yield the targeted ratio value as obtained if using equation 3.10. The objective function thus changes to either maximise return or to minimise risk, with the further constraint that the ratio value must be equal to some chosen constant, c . Using the targeted Sharpe ratio as the chosen ratio, these optimization problems are shown below:

$$\begin{aligned} \text{Max}U &= \mu^T x & (3.15) \\ \text{subject to } \frac{(\mu - \mu_f)^T x}{x^T \Sigma^* x} &= c \\ \text{subject to } e^T x &= 1, a_i \leq x_i \leq b_i \text{ for all } i = 1 \dots N \end{aligned}$$

Or

$$\begin{aligned} \text{Min}U &= x^T \Sigma^* x & (3.16) \\ \text{subject to } \frac{(\mu - \mu_f)^T x}{x^T \Sigma^* x} &= c \\ \text{subject to } e^T x &= 1, a_i \leq x_i \leq b_i \text{ for all } i = 1 \dots N \end{aligned}$$

Where μ is a vector of Fund expected returns, μ_f a vector of the risk-free rate expected returns, x is the vector of weights, Σ^* is the covariance matrix of excess returns, and a_i represents the minimum and b_i the maximum allowed exposure or weight to asset 'i' for the relevant constrained assets. a_i must always be positive.

The optimization of equation 3.15 or equation 3.16 will also provide a single solution yielding the optimal weights that will provide the targeted ratio value. The weights will vary depending on the choice of objective function and any other additional constraints. Again, this optimization is equivalent to solving for the optimal weights for the different risk-aversion levels and then from the choice of optimal portfolios where the Ratio is at the targeted value, choosing either the portfolio with the maximum return or the minimum risk.

3.4 Risk Budgeting and Manager Allocation

Having established the Fund's optimal asset allocation, as described in the preceding section, it needs to be decided how best to invest the assets. The first step of this decision process is to develop an investment strategy.

The Investment strategy forms the core of the investment decision, as it takes into account all possible ways of investing the assets and arrives at an answer that would best suit the Fund. All the relevant bases need to be covered, in terms of the choices as outlined in the Investment Strategy section earlier. The Fund needs to decide whether it requires an active or passive management style to achieve its objectives. Also within that decision, whether to allocate balanced or specialized mandates or a combination thereof. In the case of specialized mandates the fund needs to further contemplate whether there will be a further mix of active and passive management styles. The final stage of the process would be deciding how many and to which managers to allocate assets for management.

In this process the Fund would need to guard against the common mistake of diversifying the number of managers and not the number of styles and strategies [20]. Increasing the number of asset managers used to manage a Fund's assets may very well spread the Fund's manager specific risk and provide diversification benefits. However the Fund also needs to diversify itself with respect to the styles and strategies inherent within the chosen managers. It may not be worthwhile having assets managed by say five different managers, if for example they all managed assets according to a Value process. In this case, even though the Fund would have diversified its manager specific risk, it would have a significant Value style bias. This implicit bet against Growth would be fine if the Fund intentionally preferred Value to Growth but should be neutral if the Fund has no particular preference or view with regard to Value and Growth.

The process of getting all these details correct is an intricate one of finding the right balance and something that the Fund Trustees are comfortable with. Also the process of selecting managers could be time consuming and requires in-depth research into the people and processes of the asset management organizations, as described in the section on Manager Selection.

With Risk Budgeting however, there is an extra step required before deciding how much of the available assets to allocate to the managers. As mentioned briefly earlier, the decision to allocate mandates to active managers, once the asset allocations have been established as per the asset allocation process described in the previous chapter, upsets the implicit risk allocation that has been determined and agreed upon by the Fund initially. Since the initial risk allocations were made based on passive benchmarks, the addition of active risk (tracking error) to the Fund structure renders the optimal asset allocation potentially sub-optimal.

Taking on active risk may generate positive active return, however there is also a chance that it could generate negative active return, thereby detracting from the investor's expectations. Thus tracking error is a scarce resource, in that the Fund would only be willing to take on a limited amount of additional risk over and above that of the benchmark. The Funds appetite for tracking error must thus be weighed against several factors, including its ability to sustain losses in excess of its strategic benchmark. It is thus important that the Fund does its best to allocate it efficiently, so as to obtain the best possible return. The Fund also needs to budget for a realistic level of active risk commensurate with its ability to tolerate persistent negative active return [15].

As explained by Waring [22], whenever there is a trade-off of expected return against expected risk, there exists an opportunity for optimization. In this case the objective would be to maximise the active return for

a given level of active risk across all the asset management organizations being considered for inclusion in the Fund's strategy. The objective function is similar to the one used for asset allocation in equation 3.10, except that the variables are in active space. For this optimization too, as discussed previously, optimal allocations are obtained for varying levels of risk-aversion.

$$MaxU_A = \mu_A - \lambda_A \sigma_A^2 \quad (3.17)$$

Where μ_A is expected active returns of managers, λ_A represents the risk aversion factor with respect to active risk, σ_A^2 is the variance of active return.

This optimization solves the problem of optimally allocating active risk, however it ignores another problem that exists. Different asset management organizations may have different internal asset class benchmarks to those that the Fund has decided on. This difference introduces mismatch risk into the Fund's strategy, and the aggregation of all the possible mismatches at Fund level could move the total portfolio significantly away from its intended benchmark, thus once again rendering the investment strategy sub-optimal or ineffective. It is thus useful to separate the active return and risk into two components, i.e. selection active return (and its risk) and the mismatch active return (and its risk) according to an agreed upon benchmark for each manager. The selection return is the difference between the managers fund return relative to the return of the benchmark agreed upon with the Fund. The mismatch return is the difference between the agreed upon manager benchmark returns and the Fund's initial asset allocation benchmark returns [20].

The combination of the component active risks and returns combine to the total active risk and total active return respectively, as defined in equation 3.17.

The decomposition into these component risks and returns allows the Fund to optimize its total active risk, taking into account the optimal levels for both the selection and mismatch risks. The objective function for this optimization is of the same form as in equation 3.17, just with the component parts included, as derived by Waring, et al [22], and is shown below.

$$MaxU_A = [(h^T X - h_b^T)r_k + h^T \alpha_s] - \lambda_A [(h^T X - h_b^T)\Sigma_k(h^T X - h_b^T)^T + h^T \Sigma_s h] \quad (3.18)$$

Where α_s is the vector of managers' expected selection alphas, i.e. the managers excess return relative to the benchmark, Σ_s represents the covariance matrix of selection alphas, λ_A is the risk-aversion factor to active risk, r_k is the vector of expected returns for the asset class benchmarks, Σ_k represents the covariance matrix of asset class benchmark returns, X is the matrix of managers' asset class benchmark weights, h_b is the vector of Fund's benchmark weights and h is the vector of Optimal Managers' weights, arrived at and solved for during the optimization, process

$(h^T X - h_b^T)$ represents the mismatch weight, i.e. the difference between the managers' asset class benchmark weights relative to the Fund's asset class benchmark weights, $(h^T X - h_b^T)\Sigma_k(h^T X - h_b^T)^T$ is the corresponding risk associated with this mismatch of manager and Fund asset class benchmark allocations.

$h^T \alpha_s$ is the weighted manager selection return, i.e. the return the manager's portfolio generates relative to the return achieved by the agreed upon benchmark. This return could be a positive or a negative value. Trying to forecast a managers excess return is extremely difficult, given that managers have different sources of value add and trying to forecast the extent and consistency of these would be difficult indeed. $h^T \Sigma_s h$ is the risk associated with this selection alpha.

However, one's aversion for the mismatch risk may be completely different to ones aversion for the selection risk. Even though in its current form in equation 3.18 it is assumed that these risk aversion levels are equal.

the equation could easily be modified to take into account varying risk aversion levels. This could be done by assigning each risk its own unique risk aversion level.

With the optimization in equation 3.18 we are again faced with the difficult task of quantifying ones risk aversion level. Again, as with the asset allocation optimization, we can use the different risk-adjusted performance measures, like the Sharpe ratio, to circumvent this problem.

This total active risk optimization calculates allocations to particular asset managers based on their individual capabilities with respect to the benchmarks that are at the centre of their investment processes. This enables the Fund to then assess and evaluate these managers' performances with respect to the benchmarks that they are comfortable with and have explicitly agreed upon, as opposed to the potentially different Fund benchmark that the managers may have had to try to adapt their processes to. This allocation process should then produce the optimal result from each manager and thus greatly assist the Fund to best meet and achieve its objectives.

For these reasons equation 3.18 will be the central optimization equation that we will use in our optimization process. Once we have established our optimal asset allocations by optimizing equation 3.9 using the relevant chosen constraints, we will establish the optimal asset manager allocations by optimizing equation 3.18.

Once the manager choice and the allocations to these managers has been decided upon, the Fund should implement some restrictions and "rules" relevant to the various manager portfolios being managed, so that the combination thereof still allows the Fund to be within its allowed limits, as per legislation. Some typical limits placed on managers are tracking error limits, maximum allowed cash weightings, share and industry exposure limits, etc. These limits aim to ensure that the managers manage the assets to the best of their abilities but also within the framework of the Funds goal of achieving its objectives. Should managers stray away from these limits, it has implication on and may prejudice the Fund at a total asset level. For example, if a manager has been allocated active risk of 10% but chooses to manage rather at a level of about 5%, the manager is not fully utilizing the portion of the risk budget that was allocated to it to generate alpha, and thus the excess could rather have been allocated elsewhere to obtain and fulfill the Fund's objectives better. In the extreme, these diversions from the optimal risk allocation could affect the Fund in that it would be unable to or fall short of meeting its objectives. A higher tracking error is as much of a concern as low tracking error as it too could potentially cause the manager to fail to achieve its return targets.

For this reason and also just as good practice, it is imperative that the Fund monitors and evaluates its current strategy and objectives and the achievement thereof on a regular basis. This would include monitoring the chosen managers with respect to the risk and return profiles and whether they are consistently achieving the objectives set for them.

3.5 Inputs used in the Optimisation Process

The Pension Fund optimization processes and how we arrive at results that are optimal have been discussed in earlier sections. Also the requirements and inputs for the optimization have been covered, however the intricacies of arriving at those inputs have thus far been omitted. These issues will be explored and discussed in this section.

The two main inputs into any optimization problem, as can be seen from the objective functions, are the Covariance matrix and the Expected Returns. Depending on the optimization required, these are calculated for assets and portfolios of assets. These portfolios assets could be, for example, managers' portfolios or

benchmarks and asset class portfolios or benchmarks, etc.

3.5.1 Covariance Matrix Estimation

Since the true covariance matrix of asset returns is not observed, it must be estimated using statistical techniques. The estimation of the covariance matrix is however difficult and there is no one clear method of estimating it. Each estimator has its own inherent limitations, which could affect the outcomes of the optimization procedures. It is thus best to be cognizant of the limitations of the estimator being used and its impact on the optimization results, as investment decisions and performance may be significantly affected by the choice of estimator

The Sample Covariance Matrix

The sample covariance matrix is the estimation method that is most commonly used and makes use of historical returns.

$$Cov[r_i(m), r_j(m)] = \frac{\sum_{t=1}^T [r_{i,t}(m) - \overline{r_i(m)}][r_{j,t}(m) - \overline{r_j(m)}]}{T} \quad (3.19)$$

Where $r_{i,t}(m)$ is the return on an asset i between month $t - 1$ and month t , $\overline{r_i(m)}$ is its sample mean and T is the number of observations.

This method is thus easy to compute; however it has many limitations [15]. Some of these are:

- This estimator requires many parameters (i.e. variances and covariances) to be estimated, because it imposes too little structure and overfits the sample data [5]. The sample covariance matrix requires $n*(n+1)/2$ parameters to be estimated, where n is the number of assets, etc. For example, if we were calculating the covariance matrix for 100 assets we would need to estimate 5050 parameters (100 variances, 4950 covariances). The problem of estimating these large numbers of parameters is known as the "curse of dimensionality" [16]. In addition, the minimum number of historical returns required to generate a positive semi-definite covariance matrix for this case would be 100 historical returns, which is equivalent to about 8 years of monthly data for each asset [15]. This subsequent requirement of a long history to estimate these parameters is a problem in that such a long history may not be available for each asset, and even if it were, the problem could be further compounded if the older data were no longer relevant. If we consider equity assets (shares), it is entirely possible that the returns that a particular company's shares achieved years back is no longer relevant given possible changes, be they advantageous or detrimental, to either the company's management teams, processes, core operations and services etc. The problem of irrelevant data would thereby bias the risk estimates.
- The sample covariance matrix assigns the same weight to all the observations used in the sample. With the limitation, mentioned above, of requiring many data points for the estimation, the fact that the same weight is applied to all the data points could be problematic. It is possible that the older return data may no longer be relevant to the current conditions but would have been allocated an equal weight to the current, most relevant data, which could lead to a bias in the risk estimates.
- The covariance matrix needs to be positive definite, thus the covariance matrix has to be non-singular. This property will ensure that the covariance matrix is invertible, which it is required to be for the

optimization. For the matrix to be invertible, the columns of the matrix are required to be linearly independent. The larger the dimensions of the covariance matrix, the larger the amount of historical returns required to satisfy the invertibility condition [16]. Again, with the long history of returns required, even if it were available, the relevance of the historical data is questionable. If the covariance matrix were singular, the following condition would apply, and the following situations occur [19].

$$\text{Var}(a_1R_1 + \dots + a_NR_N) = 0, \text{ for some } (a_1, \dots, a_N) \neq (0, \dots, 0)$$

- There exists a riskless investment strategy.
- One asset is a combination of other assets, i.e. the columns of the covariance matrix are linearly dependent.
- The market allows for arbitrage.

Equation 3.19 shows how to estimate the monthly covariance matrix, but could be generalized to accommodate any time horizon or data frequency, i.e. daily, quarter, year, etc. The limitations discussed above thus apply to all the sample data frequencies used. Due to the relatively short history of financial data that is available, particularly within South Africa, when using longer time horizons, i.e. quarterly, yearly, etc. one would require more observations than are available. This may exacerbate the dimensionality problem associated with the covariance estimation.

One way to overcome the problem of using longer time horizon data is to use higher frequency data, i.e. daily, weekly, etc. Sampling data at a higher frequency within a specified given period, rather than just extending the sampling period while keeping the sample frequency constant, improves the accuracy of the variance and covariance estimators [1]. If one then wanted to calculate the covariance matrix over a longer horizon, one would have to aggregate the covariances of the higher frequency data. However, in order to do this one needs to take into account the serial correlation between returns to construct the covariance matrix for a longer horizon (e.g. monthly). As a result of the serial correlation one cannot estimate the monthly covariance by just merely multiplying the daily covariance by the number of business days within that month¹.

The estimation of longer horizon covariance matrices using higher frequency data has been derived in Modern Investment Management [15] and can be expressed in the following equation in matrix notation.

$$\Sigma(m) = p \cdot \Sigma_0(d) + \sum_{k=1}^q (p - k) [\Sigma_k(d) + \Sigma_k(d)'] \quad (3.20)$$

Where $\Sigma_0(d)$ is the daily covariances between the daily assets returns (i.e. lag = 0), p is the number of daily returns within the month, $\Sigma_k(d)$ is the daily covariances between assets on different days, i.e. lagged. For e.g. if $k = 1$, then $\Sigma_1(d) = \text{cov}[r_{i,t+1}(d), r_{j,t}(d)]$ and $\Sigma_1(d)' = \text{cov}[r_{i,t}(d), r_{j,t+1}(d)]$, k is the lag between the assets returns (i.e. returns are days apart), q is the level of serial correlation assumed.

The equation expresses the calculation of the monthly covariance from daily covariances, but could be generalized for longer horizons. In the derivation of the above equation it was assumed that returns are

¹Within the Finance industry the standard practice of annualizing monthly volatilities by multiplying it by the square root of 12, makes the assumption that there is no serial correlation between returns.

continuously compounded and thus in this example the monthly return for each asset is calculated by adding together the asset's daily returns for that month. An advantage of this method is that it accommodates for the existence of serial correlation in high frequency data and the monthly covariance matrix generated will be positive semi-definite, however the other limitations that pertain to Sample covariance matrices, as discussed earlier still apply.

The Observation Weighted Sample Covariance Matrix

The problem of the equal weighting applied to all the observations used in the estimation of the sample covariance matrix, can be somewhat alleviated by modifying the above estimation method. This is done by applying weights to the observations used. Typically a weight of 1 is applied to the most recent observation and the weights decrease for the preceding observations by some rate, called the decay rate (δ). If w_t is the weight applied to the observation at time t , then the sequence of weights can be calculated recursively by $w_{t-1} = (1 - \delta)w_t$. The observation weighted sample covariance matrix equation is given below [15], by modifying equation 3.19.

$$Cov[r_i(m), r_j(m)] = \frac{\sum_{t=1}^T w_t^{1/2} [r_{i,t}(m) - \overline{r_i(m)}] w_t^{1/2} [r_{j,t}(m) - \overline{r_j(m)}]}{\sum_{t=1}^T w_t} \quad (3.21)$$

Again, this equation represents the estimation for monthly data, but may be generalized for other higher frequency data or for longer horizons.

Also, equation 3.20 can be similarly modified to incorporate a weighting system for the observations.

$$\Sigma(m) = p \cdot \Sigma_0(d) + \sum_{k=1}^q (p - k) [\Sigma_k(d) + \Sigma_k(d)'] \quad (3.22)$$

Where $\Sigma_0(d) = \frac{R(d)'R(d)}{\sum_{t=1}^T w_t}$ and $\Sigma_k(d) = \frac{R(d)'R_{-k}(d)}{\sum_{t=1}^T w_t}$

Where $R(d)$ is the weighted daily returns.

Here, both the order of serial correlation (q) and the decay factor (δ) will have to be estimated in order to estimate the covariance matrix.

Although the observation weighted covariance matrix takes care of the problem of applying the same weight to all the observations, it is still prone to the other limitations of the sample covariance matrix, as well as the additional limitation discussed below.

- The observation weighted covariance method applies the same weights across the entire covariance matrix. This could be problematic when evaluating a covariance matrix with different asset classes (i.e. equity, bonds, etc) in that it may be desirable to allow for different weighting schemes or decay factors for the various asset classes. Also even if looking within just one specific asset class, the same weighting is applied to both the variances and correlations. Often however it is argued that although volatilities tend to change quickly to market shocks and occur in clusters, correlations are more likely to move slowly over time, and it would thus be more relevant to have different decay rates for each one [15].

This can be addressed, by decomposing the covariance matrix into its components, i.e. correlations and volatilities. One can then estimate the volatilities and correlations using different assumptions for their decay rates, while still preserving the positive semi-definite nature of the covariance matrix, as shown below [15].

$$Cov[r_i(d), r_j(d)] = corr[r_i(d), r_j(d)] * \sqrt{var(r_i(d))} * \sqrt{var(r_j(d))} \quad (3.23)$$

$$\text{Where } var[r_i(d)] = \frac{\sum_{t=1}^T w_t r_{i,t}^2(d)}{\sum_{t=1}^T w_t} \text{ and } corr[r_i(d), r_j(d)] = \frac{[\sum_{t=1}^T v_t^{1/2} r_{i,t}(d) v_t^{1/2} r_{j,t}(d)] / \sum_{t=1}^T v_t}{\sqrt{var[r_i(d)]} * \sqrt{var[r_j(d)]}}$$

In this case the weights w and v indicate that different decay rates were used in the calculations.

GARCH Processes

As discussed in detail in Risk Management and Analysis [4], and briefly touched on here, GARCH is the acronym for Generalised Autoregressive Conditional Heteroscedasticity. The term "Heteroscedasticity" may be quite a tongue twister but simply means, "changing variance". The conditional heteroscedasticity of a data series is observed when periods of high (low) volatility tend to be followed by more periods of high (low) volatility, i.e. volatility occur in clusters.

The Autoregressive Conditional Heteroscedasticity model (ARCH) is based on a regression process that captures the conditional heteroscedasticity of financial returns by using a moving average of the past squared returns. The ARCH (p) conditional variance specification uses p time periods and is given by the following equation.

$$\sigma_T^2 = \alpha_0 + \alpha_1 r_{T-1}^2 + \dots + \alpha_p r_{T-p}^2 \quad (3.24)$$

$$\text{Where } \alpha_0 > 0, \alpha_1, \dots, \alpha_p \geq 0$$

The constraints ensure that the conditional variance is always positive.

The GARCH model, on the other hand, includes a conditional mean equation and an equation that models the conditional variance. Because the focus of GARCH is on the conditional variance, the conditional mean equation is usually simple. Thus, the generalization of the ARCH (p) model, into a vanilla GARCH model, adds q autoregressive terms to the existing terms in equation 3.24. This gives rise to the GARCH (p, q) model with conditional variance equation given below.

$$\sigma_T^2 = \omega + \alpha_1 r_{T-1}^2 + \dots + \alpha_p r_{T-p}^2 + \beta_1 \sigma_{T-1}^2 + \dots + \beta_q \sigma_{T-q}^2 \quad (3.25)$$

$$\text{Where } \omega > 0, \alpha_1, \dots, \alpha_p, \beta_1, \dots, \beta_q \geq 0$$

GARCH processes have become popular methods to estimate financial market volatility because they allow for the variances of a series of data to be heteroscedastic by relating the variance, to the previous values of the variance of the series, and to past "shocks" in the series [19]. This conditional variance equation provides an easy analytic form for the stochastic volatility process in financial returns. Taking the square root of the GARCH conditional variance series and expressing it as an annualized percentage gives rise to a time-varying volatility estimate.

The GARCH (1,1) process for the daily volatility on an asset can be expressed as:

$$\sigma_T^2 = \omega + \alpha r_{T-1}^2 + \beta \sigma_{T-1}^2 \quad (3.26)$$

Where β the lag coefficient captures the persistence in volatility. The closer it is to 1 the larger the persistence, i.e. shocks to conditional variance take a long time to die out. α the return coefficient reflects the tendency for volatility to adjust in reaction to market surprises. ω is a constant that determines the long term average level of volatility to which GARCH forecasts converge. It is sensitive to the length of the data horizon used to estimate the model. The value will be high if a period of many years, during which there were extreme market movements, were used.

With this process, the volatility for the asset at time T depends on the volatility of the asset at time $T - 1$ (i.e. the previous period) and on the squared return on the asset at time $T - 1$.

The 2 main considerations in choosing data for GARCH models are:

- the data frequency, and
- the data time horizon.

Usually higher frequency data, like daily or even intra-day data is preferred to lower frequency data, like weekly or monthly, etc. This is because convergence problems could be encountered on low frequency data due to insufficient ARCH effects. The choice of time horizon or length of historical data used would directly influence whether it would be required that the major market events of many years ago should influence your current forecasts. Usually one or two years of daily data are necessary to ensure proper convergence of the model [4].

However the current estimate is not taken to be the forecast of volatility over all future time horizons. By first estimating the GARCH model parameters, mean-reverting forecasts of volatility can be constructed. Maximum Likelihood Estimation is used to estimate the GARCH model parameters. Based on an assumption about the shape of the distribution of the data generating process, the parameter estimate that maximises the likelihood of the data is chosen.

One advantage of GARCH is that the volatility and correlation forecasts for any horizon can be constructed from the one estimated model. The estimated GARCH model gives the forecasts of instantaneous forward volatilities, i.e. the volatility of r_{T+j} made at time T and for every step ahead, j . These forecasts are calculated analytically, shown for example in the GARCH (1,1) model below [4].

$$\sigma_T^2 = \omega + \alpha r_{T-1}^2 + \beta \sigma_{T-1}^2$$

and the j -step ahead forecasts are computed iteratively as

$$\sigma_{T+j}^2 = \omega + (\alpha + \beta) \sigma_{T-j-1}^2 \quad (3.27)$$

The volatility term structure is a plot of the volatility of returns for $n = 1, 2, 3, \dots$, where the returns (logarithmic) at time T over the next n periods is $r_{T,n} = \sum_{j=1}^n r_{T+j}$. Thus,

$$Var_T(r_{T,n}) = \sum_{i=1}^n Var_T(r_{T+i}) + \sum_i \sum_j Cov_T(r_{T+i}, r_{T+j}) \quad (3.28)$$

The GARCH forecast of n -period variance is, as can be seen from the equation above, the sum of the instantaneous GARCH forecast variances and the double sum of the forecast autocovariances between the returns [4]. However, in most cases the simple conditional mean equation in a GARCH model is merely a constant, so the returns are independent and the double sum of autocovariances is zero. The n period volatility forecasts are calculated by summing the j step-ahead GARCH variance forecasts.

The relationship between GARCH volatility and Implied volatility is used to predict future price movements. Implied volatilities are those volatilities that are implicit in the prices of options, and can be backed out from the explicit option pricing formulae, like the Black Scholes formula. Since these volatilities are based on prevailing market prices rather than on the past history of returns, they are forward looking measures. Implied volatilities are however different from statistical volatilities in that they are volatility forecasts, with a horizon given by the maturity of the option, rather than an estimate for current volatility.

Similarly, Implied correlations are those correlations that are implicit in the prices of options and can be calculated from three implied volatilities by rearranging the formula for the variance of a difference. The correlation between x and y is calculated below, using their implied volatilities.

$$\rho = \frac{\sigma_x^2 + \sigma_y^2 - \sigma_{x-y}^2}{2\sigma_x\sigma_y} \quad (3.29)$$

There are only a limited number of ways to back out implied correlations from option prices however, and they may be unstable.

Factor Models

A factor, as defined in Modern Investment Management [15], is a random variable that, at a particular point in time, can explain or account for the variation among a set of asset returns. It can also be looked at as, a variable that is common to a set of asset returns, influencing each return through its factor loading.

Assets' returns, however, do not necessarily depend on only one factor to explain how they vary, but may rely on a combination of factors. Thus a factor model provides the structure that allows the variation among asset returns to be explained in terms of this combination of factors.

Factor models, in general, are represented by the following mathematical expression:

$$r_n(t) = \sum_{k=1}^K \beta_{nk} f_k(t) + \varepsilon_n(t) \quad (3.30)$$

Where $r_n(t)$ is the return of asset n at time t , β_{nk} is the loading (exposure) of asset n to factor k , $f_k(t)$ is the return of factor k at time t , $\varepsilon_n(t)$ the disturbance term of asset n at time t .

Factor models have been widely used within the investment management field to quantify a portfolio's return and risk characteristics. For example, they may be used for performance evaluation and attribution or for portfolio risk optimization. Factor risk models, in particular, are of interest in this chapter as they allow the risk of a portfolio to be decomposed into various risk sources, represented by risk factors. The variance matrix of the asset returns at a particular time, t , is given by:

$$V_t = B\Sigma_t B^T + \Delta_t \quad (3.31)$$

Where B is the $N \times K$ exposure matrix, Σ_t is the $K \times K$ covariance matrix of risk factor returns, Δ_t is the $N \times N$ diagonal matrix of the disturbance term variances (assuming the returns are uncorrelated)

A major advantage of factor risk models is that they offer a better solution to the "curse of dimensionality" problem, in that their use significantly reduces the number of parameter estimates required to construct the variance matrix. Given the formula in equation 3.31 for a multi-factor model, one would require $k(k+1)/2 + (n*k) + n$ parameter estimates, where k is the number of factors and n is the number of assets. For example, in the case of a 100-asset portfolio and a five-factor model, one would require estimates for

615 parameters, as opposed to the 5050 estimates required with the sample covariance matrix, as discussed earlier.

The risk factors can be estimated using cross-sectional or time series estimation, and most often, by some combination of the two. Within cross sectional estimation, it is assumed that each asset has a measurable, known exposure (loading) to the factors. These exposures may be, for example, industry classifications or style exposures etc. The returns of each asset are then cross-sectionally regressed on the factor exposures and repeated over time to obtain a time series of factor returns [15]. With time series estimation, the factor returns over time are pre-calculated and assumed known, and the factor exposures are estimated on a stock-by-stock basis [16].

There are four main types of factor models, i.e. The Market model, Macroeconomic Factor models, Fundamental Factor models and Statistical Factor models. The Market model and the Macroeconomic Factor Models are examples of models that utilize observed factor returns, while the Fundamental and Statistical Factor Models utilize factor returns that are unobserved, and therefore the returns need to be estimated.

The Market Model The market model is the most simplistic model in that it has only one factor, i.e. the market portfolio return. In keeping with equation 3.30, the following one-factor model describes asset n 's excess return (over the risk free rate) in terms of the market portfolio's return, at time t .

$$r_n(t) - r^f(t) = \beta_n(t)[r^m(t) - r^f(t)] + \varepsilon_n(t) \quad (3.32)$$

Where $r_n(t)$ is the return of asset n at time t , $r^f(t)$ is the return on a risk-free asset at time t , $\beta_n(t)$ is the loading (exposure) of asset n to the market. It is referred to as the Market Beta, and measures the covariation between the market and asset return, $r^m(t)$ is the return on the market portfolio at time t , $\varepsilon_n(t)$ is the disturbance term of asset n at time t . In this model it is assumed that it has a zero mean, as it is uncorrelated to the market returns and is asset specific.

To estimate the risk of the asset the market beta needs to be estimated using time series regression. An R-squared statistic measures how much the variation in the asset's excess return is explained by the variation in market returns. However, as mentioned earlier, the market return is not the only factor that may explain movements in asset's excess returns, and it has been shown that the market portfolio is in fact poor at explaining movements in individual asset returns [15].

The Macroeconomic Factor Model Given that the country's economy and some other global economic forces influence companies and subsequently their assets' returns, macroeconomic factor models utilize these economic variables, as factors, to try to explain the variation in asset returns. Examples of these macroeconomic factors are, interest rates, exchange rates, inflation, growth rates in industrial production, etc.

One way of interpreting macroeconomic factors is that they influence companies in addition to and independently of the effect of the market. If viewed in this light, the macroeconomic factor model would, as in the case of the market factor model, comprise the market portfolio return as its first factor. The macroeconomic factors would then, because of their independent impact on individual assets, be a part of the disturbance term of each asset. This is represented by the equations below.

$$r_n(t) - r^f(t) = \beta_n(t)[r^m(t) - r^f(t)] + \varepsilon_n(t) \quad (3.33)$$

$$\text{and } \varepsilon_n(t) = \gamma_{n,1}(t)f_1(t) + \gamma_{n,2}(t)f_2(t) + \dots + \gamma_{n,K-1}(t)f_{K-1}(t) + u_n(t) \quad (3.34)$$

Where $r_n(t)$ is the return of asset n at time t , $r^f(t)$ is the return on a risk-free asset at time t , $\beta_n(t)$ is the loading (exposure) of asset n to the market, $r^m(t)$ is the return on the market portfolio at time t , $\varepsilon_n(t)$ is the disturbance term of asset n at time t . In this model it is no longer assumed that it has a zero mean. $f_k(t)$ is the return on the k th macroeconomic factor at time t , $\gamma_{n,k}(t)$ is the loading of the k th factor on the n th asset, $u_n(t)$ is the specific return of asset n at time t .

Alternatively the macroeconomic model can be viewed as a multifactor model, with each factor trying to explain a portion of the variation in asset returns. This would be represented mathematically by equation 3.30, where $f_k(t)$ would be the returns of the relevant k macroeconomic variables, and $\beta_n(t)$, asset n 's respective loadings for each factor at time t .

Fundamental Factor Models Some examples of fundamental factors are size, growth, dividend yield, price to book, industries, etc. While the exposures to these factors are measurable, their returns are unobserved. With unobserved factor returns, the values need to be estimated using information on their exposures and asset returns. Again, this can be done using either the cross section or the time series of returns. A linear cross sectional model is a popular factor model where the unobserved returns are based on fundamental variables, and is represented in matrix notation below citeLitterman.

$$R(t) = B(t-1)F(t) + u(t) \quad (3.35)$$

Where $R(t)$ is the N vector of one period asset return, $B(t-1)$ is the $N \times K$ matrix of asset exposures to factors as of time $t-1$, $F(t)$ is the K vector of one-period factor returns, $u(t)$ is the N vector of one period specific returns.

Statistical Factor Models These models use various maximum-likelihood and principal-components based factor analysis procedures on cross-sectional/time series samples of asset returns to identify the factors in returns [8]. As a result, one of the major criticisms of statistical factor models is that it is difficult to relate the statistically derived factors to intuitive economic variables. However, it is possible to model other 'observed' or defined factors, like value, growth and size etc, within a principal component model. This is useful for handling measured exposures and controlling the risk of those exposures.

Principal Component Analysis (PCA), is a multivariate procedure that reduces matrices of data to orthogonal factors. These factors are thus independent or uncorrelated with each other. These uncorrelated factors are essentially linear combinations of the original variables and are ordered by reducing variability. Thus the last of these factors can be removed with a minimum loss of data, thereby reducing the dimensionality of the data while still retaining as much information as possible [3]. For this, however, it still needs to be determined exactly how many of the principal components are significant. The first principal component explains the greatest amount of the variation and the second principal component explains the next largest amount of variation, but it is independent to the first principal component, and so on.

As discussed by Litterman in Modern Investment Management [15], the standard principal component method to estimate factors is shown below.

Assuming that specific returns are uncorrelated, the variance of $R(t)$ is

$$V(t) = B(t-1)B(t-1)^T + \Delta(t) \quad (3.36)$$

Where $V(t)$ the variance of the asset returns, $B(t-1)$ is the $N \times K$ matrix of asset exposures to factors as of time $t-1$, $\Delta(t)$ is the diagonal matrix of residual returns.

Assuming $\Delta(t)$ is small enough to ignore,

$$V(t) \cong B(t-1)B(t-1)^T \quad (3.37)$$

To estimate the exposures matrix, a sample estimator of $V(t)$ is

$$\hat{V}(t) = \frac{1}{T} \sum_{j=0}^{T-1} R(t-j)R(t-j)^T \quad (3.38)$$

By decomposing $V(t)$ in terms of its eigensystem, the exposures matrix, B , can be found

$$V(t) = P(t)\Theta(t)P(t)^T \quad (3.39)$$

Where $P(t)$ is $N \times N$ matrix of eigenvectors and $p_n(t)$ represents the n th column of $P(t)$, $\Theta(t)$ is $N \times N$ diagonal matrix of the eigenvalues $\theta_n(t)$ ($n = 1, \dots, N$) as its elements.

By substitution, from equation 3.37 and 3.39 it follows that

$$BB^T = P(t)\Theta(t)P(t)^T$$

and the factor exposures matrix, B , is determined by the K largest eigenvalues and their corresponding eigenvectors: i.e.

$$\hat{B} = P\Theta^{-1/2} = [\sqrt{\theta_1}p_1(t) | \sqrt{\theta_2}p_2(t) | \dots | \sqrt{\theta_K}p_K(t)] \quad (3.40)$$

Given the estimate of $B(t-1)$, the factor returns, $F(t)$, can be estimated by regressing $R(t)$ on B , which gives:

$$F(t) = \left(B^T B \right)^{-1} \hat{B}^T R(t) \quad (3.41)$$

or

$$\hat{F}(t) = \begin{bmatrix} \frac{1}{\sqrt{\theta_1(t)}} p_1(t)^T R(t) \\ \frac{1}{\sqrt{\theta_2(t)}} p_2(t)^T R(t) \\ \vdots \\ \frac{1}{\sqrt{\theta_K(t)}} p_K(t)^T R(t) \end{bmatrix} \quad (3.42)$$

Where $p_K(t)^T R(t)$ represents the K th principal component of returns. From equation 3.42 it can be seen that each estimated factor return is a weighted average of the asset returns, where the weights are given by its corresponding eigenvector.

Shrinkage Estimators

The shrinkage concept was first introduced by Professor Charles Stein in 1955, and is thus sometimes referred to as the Stein Estimator. Subsequently, Olivier Ledoit has also researched and written several papers on the mechanics and benefits of this type of estimation, as explained later in this section.

The main objective of shrinkage estimators is to reduce estimation error implicit in the use of covariance matrix estimators for portfolio optimisation. Given previous research in this field it is intuitively believed, within the Finance profession, that the best model lies somewhere between an unbiased estimator (which

would have a lot of estimation error) and a biased estimator (with little estimation error). One way of obtaining this is to simply take a weighted average of the biased and unbiased estimators, thereby pulling the unbiased estimator towards a fixed target (Shrinkage target) represented by the biased estimator [13].

As discussed earlier, the most common unbiased estimator is the sample covariance matrix and it is estimated using historical asset returns. The sample covariance matrix estimation requires a large number of variance and covariance parameters to be estimated, when a large number of assets are under consideration. Should the historical data available not be sufficient, the estimation of these parameters is riddled with error. In this case the estimation of the most extreme coefficients of the covariance matrix tend to take on extreme values, that are not necessarily a true reflection of the data, because the estimates are distorted by an extreme amount of error. These error-ridden estimates may thus completely distort the results of any mean variance optimization.

The shrinkage estimator transforms the sample covariance matrix by pulling the most extreme coefficients towards more central values, thereby reducing the estimation error where it is at its maximum. This "shrinkage" process would reduce the extremely high estimates, which contain positive error, and similarly pull the extremely low estimates, which contain negative error, upwards. However the key to this process is determining the shrinkage target and the optimal shrinkage intensity [14].

As discussed by Ledoit and Wolf [14], a shrinkage estimator has 3 main components, i.e. an estimator with no structure (i.e. unbiased (the sample covariance matrix)), an estimator with a great deal of structure (i.e. biased (the shrinkage target)) and the shrinkage constant. Shrinkage thus results in a covariance matrix estimate that is a compromise between the sample covariance matrix and the highly structured estimator, and which performs better than either extreme. This is represented by the equation below.

$$\Sigma_{shrink} = \delta F + (1 - \delta)S \quad (3.43)$$

Where Σ_{shrink} the shrinkage estimator of the covariance matrix, F represents the shrinkage target, S represents the sample covariance matrix and δ is the shrinkage constant and $0 < \delta < 1$

The choice of shrinkage target may vary depending on its appropriateness to the optimization problem at hand. Some examples are, the single factor model (market model), the multi-factor model or as discussed by Ledoit and Wolf [14] the constant correlation model. The single-factor model and the constant correlation model are discussed briefly below.

The formula for the single factor or market model is shown below using the same assumptions as described earlier in the previous section.

$$r_n(t) = \beta_n(t)[r^m(t)] + \varepsilon_n(t)$$

The covariance matrix can be estimated by running a regression of asset n 's returns on the returns of the market to yield the following estimator.

$$F = s_{oo}bb^T + D \quad (3.44)$$

Where s_{oo} is the sample variance of the market returns, b is the vector of estimated market factor loadings and D is the diagonal matrix of residual variance estimates.

The "constant correlation model" assumes that all pairwise correlations are identical. The common constant correlation is estimated by the average of the sample correlations, and together with the vector of sample variances, forms the shrinkage target. This is shown formally in equation 3.45 below.

Let $y_{it}, 1 \leq i \leq N, 1 \leq t \leq T$, denote the return of asset i during period t . Assuming that returns are independent and identically distributed over time and that they have finite fourth moments, the sample average of asset i 's return is $\bar{y}_i = T^{-1} \sum_{t=1}^T y_{it}$. The entries of the sample covariance matrix (S) are given by s_{ij} .

The sample correlation between asset i and j is calculated

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

and the average sample correlation

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

Thus the Sample constant correlation matrix, F , shown in terms of the sample variances and average sample correlation is:

$$f_{ii} = s_{ii} \text{ and } f_{ij} = \bar{r} \sqrt{s_{ii}s_{jj}} \quad (3.45)$$

The advantage of using the constant correlation model is that it is easily implemented; however a major disadvantage is that the model is not appropriate if assets are in different asset classes. This is a particularly significant disadvantage with respect to the Financial industry, in which Pensions Funds play a major role, in that various asset classes are investigated and utilised during optimization exercises for diversification purposes. Some of the other highly structured estimators, like the multi-factor model, can cater for assets in different asset classes and may thus be better shrinkage targets in these cases.

The other important input for shrinkage estimation is the shrinkage intensity, quantified by the shrinkage constant, λ . As seen in equation 3.43, the shrinkage constant is a measure of the weight given to the structured (biased) estimator, i.e. the shrinkage target. The shrinkage constant can be values between 0 and 1, and any value in this range would yield a compromise between the sample covariance matrix and the shrinkage target. However, there is an optimal shrinkage constant that minimises the difference between the estimated covariance matrix, using shrinkage estimation, and the actual covariance matrix. There are several statistical techniques that allow one to obtain this optimal shrinkage constant, and the technical detail of these can be found in Ledoit and Wolf [13].

The discussion in this section has outlined a variety of methods in which one can estimate a covariance matrix. Each method has its respective advantages and disadvantages, and in some cases, a method builds on another method in an attempt to address some of that method's disadvantages. Ultimately, it is important that one understands the ins and outs of the methods, so that one understands the strengths and failings of the outputs, and thus makes informed decisions as to which method best suits one's needs.

3.5.2 Expected Return Estimation

Another important input into the optimization process is the expected returns of the assets or asset classes utilised. The future is uncertain, and unfortunately we all do not have a perfectly accurate "crystal ball" with which to predict what it holds and how it will affect our investments. With the many influences on asset returns, like the state of the economy, exchange rates, inflation, etc. it is difficult to know the true future returns of assets. Similarly it is extremely difficult to forecast a manager's return or even more importantly for pension funds, a manager's relative return, given the different sources of value-add

that managers utilise. Therefore one has to utilise estimation techniques in order to obtain some expected return.

However with estimation, there is the inherent risk that the estimate of the return may be very different from what the true return is. This potential error when used in an optimization process could be further exacerbated by the problems implicit with optimization techniques and thus lead to distorted results. It is with these problems in mind that many methods to estimate expected returns have arisen in an attempt to minimize the error implicit in the return estimation.

Some of the various ways in which to estimate expected returns are discussed in this section.

Average/Mean Returns

The simplest estimator for expected returns is the average. This average may be calculated by using daily, weekly, monthly, etc periods. The choice of analysis period however will largely depend on the amount of data that is available.

Standard Historical Average With this estimator, the expected return of the next period is the long-term average of all the return data used. This is calculated using the formula shown below:

$$E[R_{t+1}] = \frac{\sum_{i=1}^t X_i}{t} \quad (3.46)$$

Where $E[R_{t+1}]$ is the expected return for the next period, X_i the return data points for each period i , such that $i = 1 \dots N$.

This estimation method is very dependent on the amount of data used in its calculation. One of the major disadvantages of this though is ensuring that one has sufficient history over which to draw conclusions. It can be argued that the more history that is used, the more likely the estimate is to being close to the true value.

From equation 3.46 it can be seen that the historical average calculation allocates an equal weight to each of the historical data points used. Another concern related to this is, even if one has sufficient historical data, is the oldest data still relevant to current times. It could be that there may have been structural changes, over the period being used, that have either changed the returns for the better or for the worse. Equally weighting all this data may, in this case, lead to an incorrect estimate. Also the average expected return is also prone to error if there are just one or two extreme events, either large or small [18].

Weighted Historical Average The weighted average method assigns a weight to each of the data points used in the calculation. By choosing appropriate weights, this method can alleviate the concern about the relevance of the older data and the impact this older data should have on the expected return estimates. For this method, the standard historical average method is modified as shown below.

$$E[R_{t+1}] = \frac{\sum_{i=1}^t (w_i X_i)}{\sum_{i=1}^t w_i} \quad (3.47)$$

Where w_i is the weight applied to the i th data point, $E[R_{t+1}]$ is the expected return for the next period, X_i the return data points for each period i , such that $i = 1 \dots N$.

Typically the most recent observation (t) receives the highest weighting and the weights then decrease based on some, decided upon, constant decay rate or by some fixed value. However if all the weights don't sum to 1, the calculation would not really be an average, and thus to correct this, each weighted data point is divided by the sum of the weights.

Moving Average Returns

The moving average return estimate for expected returns aims to address the issue of using too much of the older, possibly irrelevant, data that is experienced when using average returns.

Standard Moving Average As mentioned previously, the use of older, irrelevant data would distort the expected return estimate [21]. The moving average method decides the data period over which to calculate the expected return in advance. It thus omits a certain period of the history. The formula used is shown below:

$$E[R_{t+1}] = \frac{\sum_{i=t-n}^t X_i}{n} \quad (3.48)$$

Where $E[R_{t+1}]$ is the expected return for the next period, X_i the return data points for each period i , such that $i = t - n \dots t$ and n is the number of most recent data points included in the calculation.

Weighted Moving Average With the use of the standard moving average as an estimator, the problem of using older data is addressed. However, the standard moving average still equally weights all the data points it uses for the estimation.

Even though the chosen amount of data (n) may be relevant to the calculation of the expected return, it may still be better to give the more recent points in this history more consideration, and thus a higher weight, than the earlier data [21]. The formula for this weighting moving average is shown below.

$$E[R_{t+1}] = \frac{\sum_{i=t-n}^t w_i X_i}{\sum_{i=t-n}^t w_i} \quad (3.49)$$

Where w_i is the weight applied to the i th data point such that $i = t - n \dots t$, $E[R_{t+1}]$ is the expected return for the next period, X_i the return data points for each period i such that $i = t - n \dots t$ and n is the number of most recent data points included in the calculation.

Here, also as with the Weighted Average return method, the Sum of the weighted returns is divided by the sum of the weights to ensure that the weights sum to 1, so that the average can be calculated.

Exponential Smoothing Exponential smoothing, sometimes referred to as the Exponentially Weighted Moving Average, is another method used to estimate expected returns. This method is in some ways like the weighted moving average method, except that it only requires that one parameter, α , which is a value between 0 and 1, instead of the various w_i , be chosen in order for the weighting factors to be calculated. Also, with the weighted moving average methods the weights decrease over the chosen range of history (n to t) and the rest of the history is omitted, i.e. the older data has a 0% weight). Exponential smoothing

on the other hand allows for a more gradual decrease in the weighting factors [1]. The formula used to calculate the expected return for the next period is shown below.

$$E[R_{t+1}] = \alpha * X_t + (1 - \alpha) * ER_t \quad (3.50)$$

Where $E[R_{t-1}]$ the expected return for the next period, $E[R_t]$ is the expected return of the current period, X_t is the actual return of the current period and α is the optimal weighting factor.

As is evident by equation 3.50, the expected return for the next period is a proportion of the current return and the current expected return. This formula however is recursive in that the current expected return is also a function of the previous return and the previous expected return, etc. Equation 3.50 can thus be re-written as:

$$E[R_{t+1}] = \alpha * X_t + (1 - \alpha) * [\alpha * X_{t-1} + (1 - \alpha) * ER_{t-1}] \quad (3.51)$$

By continuing to substitute the previous expected returns formulas into equation 3.51 recursively, and rearranging the terms, we obtain the following equation:

$$E[R_{t+1}] = \alpha * X_t + \alpha(1 - \alpha)X_{t-1} + \alpha(1 - \alpha)^2X_{t-2} + \alpha(1 - \alpha)^3X_{t-3} + \dots \quad (3.52)$$

From equation 3.52 it can be seen that the expected return for the next period is just a weighted sum of the historical returns. The weights are represented by the α and $(1 - \alpha)$ terms. Since α is a value between 0 and 1, $(1 - \alpha)$ is thus also a value between 0 and 1, and by multiplying these numbers together, the outcome is a smaller number. This results in the weights getting smaller exponentially [21].

Any value within the allowed range can be chosen for α to calculate the expected return for the next period. However an 'optimal' α can be calculated such that the sum of the squared errors are minimized. The error for each period is calculated by subtracting the expected return for period i from the actual return for period i .

Choosing α in such a manner results in the expected return for the next period being, deliberately, fitted to the past or historical data. This may be risky, considering the tagline "past returns are not an indication of future returns".

Reverse Optimisation

As mentioned previously, the estimation of the expected returns used in optimization problems is critical to the results obtained from the optimization process. It has been noted that the optimal solutions are very sensitive to the expected return inputs used [19], and as a result it is crucial to obtain the best estimate of expected returns.

While the methods mentioned previously have been methods of estimating the expected return using historical data to make an assumption of future returns, the reverse optimization method estimates the expected returns in an indirect way.

In standard optimization problems, the objective is to solve for the optimal weights, x^* , given in equation 3.53. This optimal solution assumes there are no constraints imposed on the weights.

$$x^* = \frac{(E[R] - r)}{\lambda \Sigma} \quad (3.53)$$

Where x^* is the optimal portfolio weights, representative of the average investor, $E[R]$ is the vector of expected return estimates for the assets, r is the risk free rate return, represented by some known return for cash, λ is the risk aversion rate and Σ is the covariance matrix, generated from historical data.

In order to establish these optimal weights it is required that the expected return and covariance matrix estimates, the risk free rate and the risk aversion rate are given, and thus used as inputs in the optimization procedure.

For reverse optimization however, as the name suggests, the optimization procedure is reversed. Of the two main inputs to an optimisation procedure, i.e. the expected return and covariance matrix estimates, it is assumed that the covariance matrix estimates are easier to estimate and get right. Therefore here, the expected returns need to be determined, while the other information is assumed to be given or known. The expected returns are thus "backed out" using the optimal portfolio weights, the covariance matrix estimates and the risk free and risk aversion rates as inputs. For this procedure, equation 3.53 can again be used and by rearranging the terms, the equation for the expected returns is given by equation 3.54.

$$(E[R] - r) = \lambda \Sigma x^* \quad (3.54)$$

The reverse optimization method makes use of the "state of equilibrium" assumption associated with the Capital Asset Pricing Model (CAPM). The "market", which is seen as the universe of investment opportunities, is assumed to be efficient and in a state of equilibrium (i.e. no arbitrage opportunities exist). This condition of efficiency is important since it allows for the assumption that assets are fairly priced. The equilibrium portfolio (market portfolio) is, as a result of its efficiency, assumed to be an optimal portfolio, and thus using reverse optimization, the equilibrium expected returns are determined using the market capitalisation weights [11].

This method of calculating expected return estimates allows for the use of expected returns implied by the market conditions. However its drawback is that it makes explicit use of the CAPM assumptions, which may not be appropriate in practice. The expected returns generated in this manner may be more appropriate if a better or more realistic choice of "optimal portfolio" were used.

Black-Litterman Model

While the Black-Litterman model is not a method by which to explicitly estimate expected returns, it does provide a method by which expected return estimates may be adjusted in keeping with more intuitive or subjective views on the returns estimates. These views may also arise from additional in-depth research into the companies for which accurate expected returns are required to be estimated.

While the Black-Litterman model can be implemented to adjust various types of expected returns, Black and Litterman used the model to adjust the expected returns of assets. This was done so that an investor would be able to obtain better results when trying to create an optimal portfolio. Mean-Variance optimisation is typically used as the optimisation procedure to establish these optimal portfolios.

However, the use of Mean-Variance optimisation, with its inherent problems relating to its sensitivity to its inputs, as discussed previously, has resulted in substantial time and effort being dedicated in trying to establish the correct method of estimating expected returns and constraining the optimisation in order to obtain reasonable solutions. As discussed by He and Litterman [10], in practice, many investors find that the value added from trying to obtain these reasonable solutions does not provide adequate benefit. Black and Litterman make use of equilibrium returns as a starting point for the expected returns. These

equilibrium returns are based on assumptions associated with the Capital Asset Pricing Model (CAPM) and are derived by the process of reverse optimisation, as described earlier. These returns reflect the market perception of the assets and since these perceptions may not be entirely appropriate in the light of additional research or information, adjusting them according to the investor's specific views increases their value in the optimisation procedure by making the estimates more appropriate.

An important part of the Black-Litterman model is defining the views and the confidence levels that an investor has in their views. The views may be stated in absolute terms, by specifying an asset's expected return exactly, or in relative terms, by specifying the extent to which one asset may out or under-perform another asset. In the specification of the view, the investor must also specify the level of confidence that they have in the view. The extent of the view and the degree of confidence in the view then ultimately affects the optimal solutions obtained, since the weight in an asset increases as an investor becomes more bullish on the view and also more confident in their view. Also the greater the level of confidence in a view, the closer the asset's adjusted return will be to the view, while the views with lower confidence levels will result in the asset's adjusted return being closer to the initial expected return estimate.

The Black-Litterman model enables an investor to only specify returns on which they have a view, since it allows for all the other expected returns to be adjusted away from their values in a manner consistent with their underlying covariances and the views being expressed [7]. While the investor need not have a view for every asset, the number of views is required to be less than the number of assets as a result of the mathematics involved.

As shown in Idzorek [11], the Black-Litterman formula for the adjusted expected returns is shown in equation 3.55.

$$E[R] = [(\tau\Sigma)^{-1} + P'\Omega^{-1}P]^{-1}[(\tau\Sigma)^{-1}\Pi + P'\Omega^{-1}Q] \quad (3.55)$$

Where, for a case with n assets and k investor views, $E[R]$ is a vector of the adjusted asset returns ($n \times 1$), Σ is the covariance matrix of returns ($n \times n$), τ is the scalar which measures the uncertainty of the covariance matrix, P is the matrix identifying the assets in which the investor has additional views ($k \times n$), Π is the vector of neutral or market equilibrium returns ($n \times 1$), Ω is the diagonal covariance matrix representing the uncertainty of each view ($k \times k$), Q is a vector of the views ($k \times 1$).

In order to understand how the formula works and how the adjusted expected returns are obtained, some of the terms of the equation are explained briefly.

The Black-Litterman model assumes that the error terms associated with the views are random, independent and normally distributed, with a mean of zero. The variance of each of these error terms is indirectly proportional to the level of confidence of the view. These variances form the diagonal element of Ω with its off diagonal elements equal to zero.

The assets' views expressed in the Q vector are matched to the relevant assets by the P matrix. The P matrix consists of k , ($1 \times n$) row vectors, where for an absolute view the respective row vector consists of a value of 1 for the asset on which a view is being taken and zeros for all the other assets. For a relative view the respective row vector consists of a value of 1 for the out-performing asset and a value of -1 for the under-performing asset, while all the other assets have a value of zero, and the row thereby sums to zero. Depending on the papers being consulted, for cases where there are more than one out or under-performing asset, their value may either be equally weighted, i.e. equal to 1/number of out or under-performing assets or they may be weighted according to their market capitalizations as a percentage of the total market capitalization of the out or under-performing assets [11]. The scalar, τ , should in general equal the average

value of the variance of error terms divided by the variance of the view portfolio.

From equation 3.55 it can be seen that the first part of the equation, $[(\tau\Sigma)^{-1} + P'\Omega^{-1}P]^{-1}$, is responsible for adjusting the the assets' underlying covariances to incorporate the uncertainty of the views which are being imposed, and can be considered as an effective covariance matrix. The second part of the equation, $[(\tau\Sigma)^{-1}\Pi + P'\Omega^{-1}Q]$, is responsible for combining the view with the initial equilibrium returns, by adding the implied weight associated with the view to the implied equilibrium weight to obtain an effective weight for each asset. The combination of these two parts of the equation, i.e. the effective covariance matrix multiplied by the effective weight, is an effective reverse optimization process, from which an adjusted expected return for each asset is obtained.

Thus, the Black-Litterman model is a method whereby an assets' initial equilibrium returns are adjusted by imposing views in a manner consistent with the asset's underlying covariances

Chapter 4

Application of Risk Budgeting in a Pension Fund Environment

For the purposes of this study, the processes required to demonstrate how risk budgeting would be implemented within the pension fund environment have been carried out. It is the objective of the risk budgeting process to allocate a pension fund's assets to its chosen asset managers in a manner that allows for the Fund to optimally utilise its risk budget to achieve the highest possible active return.

For the purposes of this thesis, optimisation has been utilised to obtain the various optimal solutions available for the differing levels of relative risk. Furthermore, various situations have been assessed in this chapter so as to accommodate as many plausible real life situations as possible, and to establish how the different situations impact on the optimal solutions obtained. The results of the various steps in the risk budgeting process and the different situations assessed will be discussed in this chapter.

4.1 Description of the Process and the Data used

The first part of the Risk budgeting process requires an Asset Liability Modelling (ALM) exercise be carried out. One of the main objectives of an ALM study is to determine the pension Fund's asset allocation, by contemplating the trustees risk tolerance. Part of this process is also to find an appropriate benchmark that could be considered representative of, or at least a good estimate of the Fund's liabilities. This choice may be determined by considering the extent of the expected returns that would be required to cover the liabilities.

In practice the ALM process would require that a detailed liability modelling exercise be carried out. This would entail taking into account mortality, morbidity and many other statistics in order to obtain an estimate of the true liabilities. Since this exercise falls outside the scope of this thesis, a simplistic ALM exercise has been adopted with assumptions made about the liabilities. It has been assumed that the pension fund trustees are mainly concerned with minimising price volatility and thereby minimising capital loss. For this required capital protection the minimum risk asset class would be short dated, low credit risk cash.

However, it has been further assumed that the trustees are more risk tolerant and are willing to introduce further risk into the strategy for a higher expected return. This may be achieved by investing in the generally

more risky asset classes, like equity and bonds. As such, the industry asset class indices have been assumed to be appropriate asset class benchmarks for the Fund. Also the universe of allowed investments has been assumed to be domestic only, thus excluding international benchmarks from the total Fund benchmark. The benchmarks chosen are as follows: Equity - FTSE/JSE All Share Index with 50% downweighting in Resources shares (ALSI50R), Bonds - BESA All Bond Index (ALBI) and Cash - Alexander Forbes Money Market Index (AFMM). These asset class benchmarks have been chosen because they are regarded as good barometers of the asset classes concerned.

The monthly return data for these benchmarks, for the 3 year period ending October 2004, was obtained from RisCura. These returns were used in an optimisation process in order to arrive at an optimal total Fund asset class allocation. The results of the optimisation and the subsequent choice of asset allocation is discussed in the "Asset Allocation Results" subsection, i.e. 4.2.1, later. Also it was decided to utilise the Mean Variance optimisation technique for the various steps requiring optimisation in the risk budgeting process. While some, possibly better, types of optimisation techniques do exist, mean variance optimisation was chosen because it tends to be the most widely recognised and understood optimisation technique. The Mean-Variance optimisations were performed using the Solver add-in to Microsoft Office Excel.

The subsequent step implicit in the Risk budgeting process would require that the Fund decide on an investment strategy. The Fund would need to decide whether it would try to achieve its objectives using passive or active management strategies. Assuming that the Fund would be willing to incur some risk to achieve excess returns, it was assumed for the purposes of this thesis that an active management strategy would be pursued.

Once the investment strategy is established, the next step is to find the asset manager allocation, i.e. the proportions in which the Fund's assets would be split among the various asset managers. From hereon, the Risk budgeting process comes into effect, because the Fund's decision to allocate to active managers will offset the Fund's implied risk constraints established during the ALM.

In practice, the next key decision would be to decide whether to give managers balanced mandates or specialist mandates. This decision would also impact the list of suitable asset managers available to the Fund. For the purposes of this thesis, both the balanced and the specialist mandate scenarios will be looked at. Varying situations of each of these mandate scenarios have been assumed as plausible real-life situations, and have been analysed to show how these differing situations influence the level of mismatch and selection risk and return that is introduced into the Fund's overall investment strategy. The different situations and their results are discussed in section 4.2.2 and section 4.2.3 for the balanced and specialist mandates respectively.

Before mandates are allocated, in order to narrow down the list of all potential asset managers to a final select few, an in-depth asset manager selection process is usually carried out. In the absence of this manager selection process, it was assumed that the list of chosen asset managers was narrowed down to three. These three asset managers will henceforth be referred to as Manager A, Manager B and Manager C.

For each manager, their monthly specialist equity, bond and cash portfolio returns were obtained from RisCura, for the 3 year period ending October 2004. These returns would be used for the specialist mandate cases, while for the balanced mandate cases, the three asset class returns of each manager were combined to obtain a balanced portfolio return for that asset manager. The active returns for each manager's specialist or balanced portfolios were calculated relative to their respective benchmarks.

In the final step in the risk budgeting process, the active return for each manager is used in an optimization

to establish the optimal Fund's allocation to these managers, given their respective risk and return characteristics. Within each mandate type the differing scenarios have been used in order to demonstrate the active and mismatch risks and returns that are incorporated in the strategy by choosing active managers who may have differing weights to the Fund's benchmark or entirely different benchmarks. The risk budgeting optimisation formula that splits the risk and return into the mismatch and selections components was used to show the impact the scenarios had on the levels of mismatch and selection risk and return introduced.

The entire (simplified) process looked at here can be summarized, showing the incremental risk introduced by each of the various steps, thereby ultimately making up the pension fund's risk budget. This is shown in Figure 4.1

The graphs illustrate the rand values (in thousands) of the cashflows paid to members through time, denoted by time periods 1 to 15.

The first graph shows the cashflows of the true liabilities. While this is generally the starting point for an ALM study, these cashflows are actually only truly established as they occur, and are subject to real life and the specifics of the fund. Some of the risks can be insured, and so the profile can be altered somewhat at a cost.

The second graph shows the cashflows generated by the estimation of the liabilities. 'a' shows the incremental risk introduced into the process at this early stage as a result of the mortality, morbidity, etc assumptions used in the estimation. The extent of 'a' should be small provided that the best estimation was obtained. It may be even smaller if re-insurance is introduced into the fund.

The third graph shows the cashflows of a risk neutral benchmark for the estimated liabilities. Due to the flexibility of the swap and fixed interest markets, their cashflows are very similar to the estimate and the risk introduced by this process, as indicated by 'b', is potentially negligible.

The cashflows of the fourth graph, show the result of the asset allocation selected by the trustees and could differ significantly from that of the risk neutral benchmark. This strategic benchmark may include more risky asset classes like equities and bonds and thus the risk introduced, 'c', could be expected to be vary depending on the trustees risk tolerance.

The final graph shows the cashflows obtained from the final investment strategy implemented by the trustees. If this strategy were purely passive its cashflows could be expected to be very similar to that of the strategic benchmark. On the other hand, should an active management strategy be implemented, as is done for this thesis, it is crucial to ensure that the risk introduced at this stage, 'd', be minimised by optimising the manager allocations such that the mismatch components of risk and return introduced are kept low.

Each step in the process introduces further risk relative to the original or true liability structure of the fund. The risk budgeting process attempts to quantify and classify these risks, as well as allocate the total risk to buckets of risks in each step to an extent that the trustees are comfortable. This will in each case be determined by the nature of the specific risk, and the extent to which it can be estimated and profited from. In this thesis we have simplified some of these steps.

Monthly returns have been used throughout the process in order to calculate the inputs for the optimization process, it was thus decided that the results would also be represented using monthly numbers. While in practice a Fund would generally look at annualised estimates, there is no need to do so here as a comparison

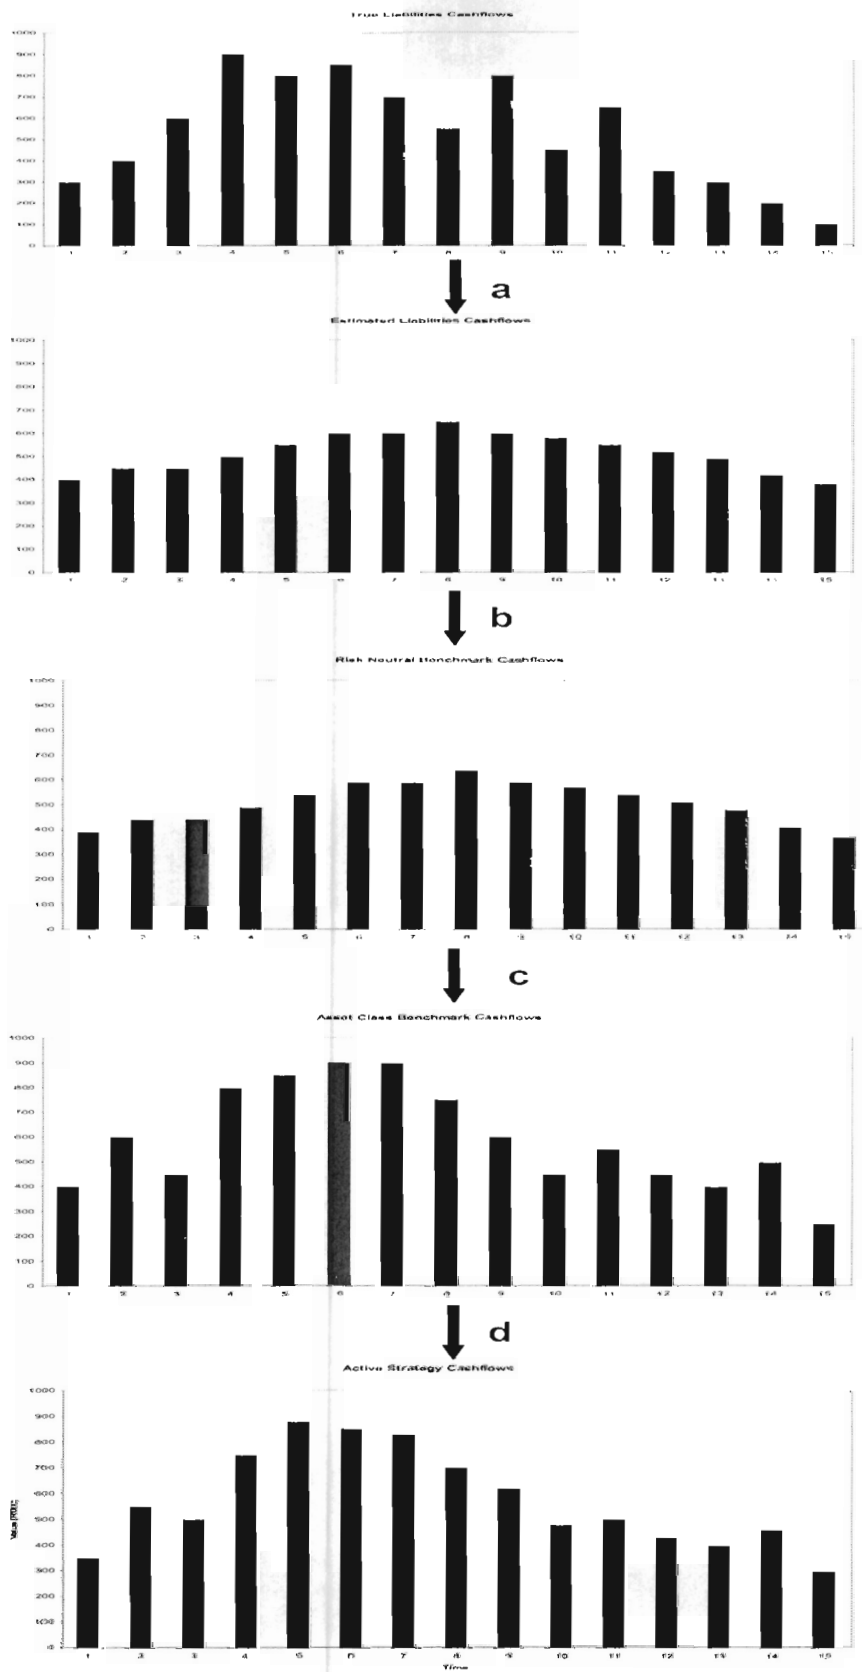


Figure 4.1: Differences in Cashflows between True Liabilities and Investment Strategy

can be made on the monthly estimates. The resulting monthly estimates could, if required, be changed to represent annual estimates.

4.2 Analysis of Risk Budgeting Results

Before getting started with the actual risk budgeting process, detailed analysis was done to assess the effects of the various types of covariance matrix and expected return estimators on Mean-Variance optimization. This initial analysis was performed so that a choice could be made, from the various estimators assessed, as to which would be most suitable and stable to use in the application of the Risk budgeting process. This analysis and the discussion of the results thereof have been included in the appendix. The expected return and covariance matrix estimators chosen for use in the risk budgeting process are discussed below.

From this analysis it was decided that weighted average returns, with a Black-Litterman adjustment would be used as the expected return estimator. The weighted returns were chosen because it would be more appropriate to give more credence to the most recent returns rather than the older returns. For example, a manager could have undergone major structural changes over the period being analysed, which could thus nullify the existence of any exceptional or poor returns that may have existed previously. The linearly weighted return methodology with the constant decay factor of $\frac{1}{36}$, shown by equation A.1 and used for Scenario 2a, was utilized for the entire risk budgeting process. The weights derived from this were applied across all asset class benchmarks and for all the individual manager returns used.

The Black-Litterman model has, here again, been adapted to apply to the adjustment of the expected returns for asset managers. The inclusion of the Black-Litterman adjustment allowed for the inclusion of a view which may be based on qualitative or other subjective information, thereby providing a better estimation for the manager's all round skill. The various views used for the different situations looked at will be discussed as they arise.

The Ledoit covariance matrix was chosen as the preferred covariance matrix estimator for the Risk Budgeting process. This was as a result of it being able to address the dimensionality problem associated with the Sample covariance matrix estimates and for being a generally more stable estimator. Also, as established by the analysis performed, the optimization is more sensitive to the expected return inputs and thus it was decided that it would not be necessary to weight the returns used for the covariance matrix estimate, since the expected return estimates were to be weighted in any case.

As explained in the previous chapter, the function used in the optimization to determine the optimal allocation to the chosen managers is broken down into its mismatch and selection risk and return components. This function, as shown in equation 3.18, has been replicated below.

$$MaxU_A = [(h^T X - h_b^T)r_k + h^T \alpha_s] - \lambda_A [(h^T X - h_b^T)\Sigma_k (h^T X - h_b^T)^T + h^T \Sigma_s h]$$

Where α_s is the vector of managers' expected selection alphas, Σ_s represents the covariance matrix of selection alphas, λ_A is the risk-aversion factor to active risk, r_k is the vector of expected returns for the asset class benchmarks, Σ_k represents the covariance matrix of asset class benchmark returns, X is the matrix of managers' asset class benchmark weights, h_b is the vector of Fund's benchmark weights and h is the vector of Optimal Managers' weights, arrived at and solved for during the optimization process

Please note that in the following sections that describe and discuss the results of the asset allocation and manager allocation steps of the risk budgeting process, the results are shown as monthly estimates, and

thus show smaller risk and return results than would be expected had the results been annualised.

4.2.1 The Asset Allocation Results

Using the input estimators chosen, and mentioned above, an optimal asset allocation was required. This sections discusses the results obtained and the ultimate optimal solutions chosen. As mentioned previously, the Fund's underlying strategic asset class benchmarks were chosen to be the ALSI50R as the equity benchmark, the ALBI as the bond benchmark and the AFMM as the cash benchmark.

The inputs used for the optimization are shown in Table 4.1 and Table 4.3, while the optimal results obtained from the optimization are shown in Figure 4.2 by the efficient frontier and the area graph, showing the various optimal asset class allocation solutions.

ALSI50R	ALBI	AFMM
<i>1.59%</i>	<i>1.20%</i>	<i>0.83%</i>

Table 4.1: Expected Return estimates for Asset Allocation

The Black-Litterman views applied to these expected return estimates are reflected in Table 4.2. These are the same views as applied to and discussed for Scenario 3a.

	View	Confidence
Equity outperforms Bonds	<i>0.40%</i>	<i>60.00%</i>
Equity outperforms Cash	<i>0.70%</i>	<i>70.00%</i>
Bonds outperform Cash	<i>0.30%</i>	<i>60.00%</i>

Table 4.2: Black-Litterman Views applied for Asset Allocation

For the Ledoit covariance matrix estimator, the single factor market model was again used as the estimator towards which the sample covariance matrix estimates would be shrunk, as seen with Scenario 4.

	ALSI50R	ALBI	AFMM
ALSI50R	0.002546	-0.000056	-0.000019
ALBI	-0.000056	0.000425	0.000008
AFMM	-0.000019	0.000008	0.000003

Table 4.3: Covariance Matrix estimates for Asset Allocation

The use of these inputs in the optimization gives rise to the optimal asset class allocation options and their respective risk and expected return estimates shown in Figure 4.2.

The most appropriate asset class allocation should be chosen from the various optimal solutions obtained. This allocation would serve as the Pension Fund's strategic benchmark, the basis upon which the Fund's investment strategy would be built.

For the purposes of this study, an optimal solution that included an allocation to all 3 asset classes was chosen to represent the Fund's strategic asset class allocation. This allocation choice and its accompanying portfolio risk and return levels are shown in Table 4.4.

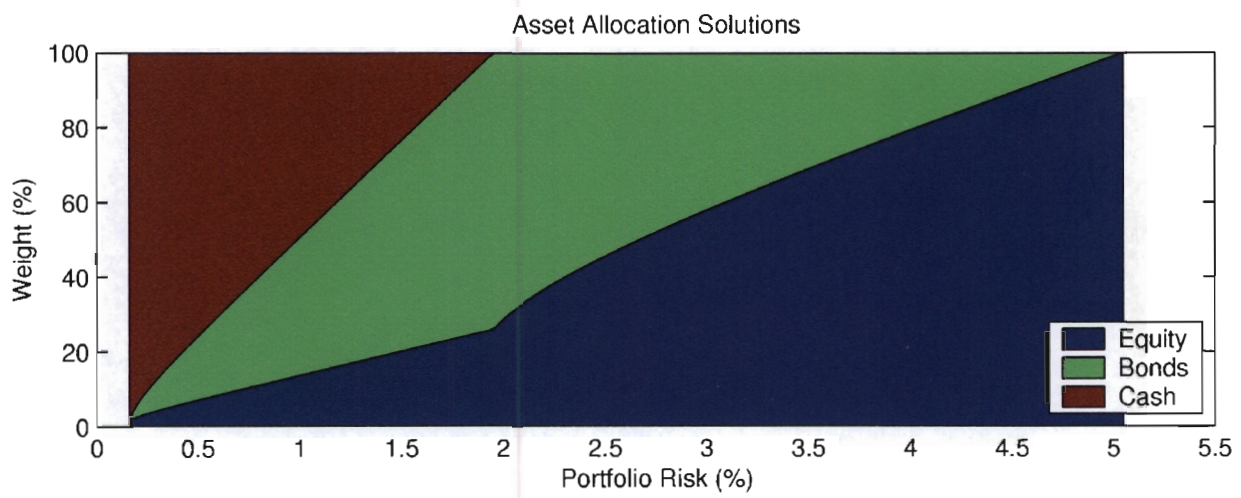
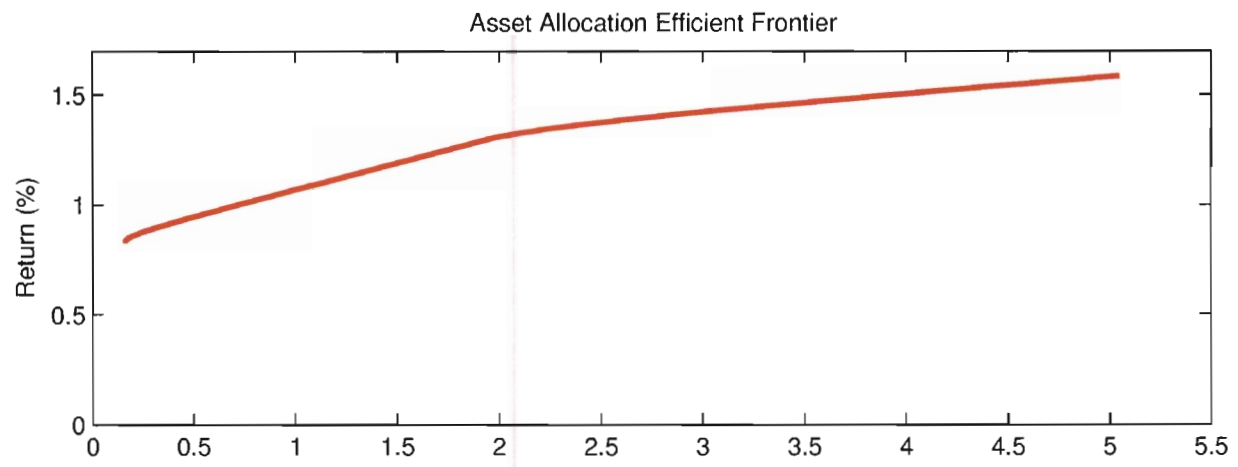


Figure 4.2: Efficient Frontier and Optimal Asset Allocation Solutions for Asset Allocation

Portfolio Risk	Portfolio Return	Equity Allocation	Bond Allocation	Cash Allocation
<i>1.82%</i>	<i>1.27%</i>	<i>24.01%</i>	<i>68.83%</i>	<i>7.16%</i>

Table 4.4: The Fund's chosen Optimal Asset Class Allocation & Its Expected Risk and Return Characteristics

Having decided on the optimal asset class allocation, the next step is to decide how to allocate mandates to the asset managers so as to best achieve the objectives implicit in the choice of the strategic benchmark. As mentioned earlier during the outline of the process, it was decided to investigate the effect of allocating both active balanced, as well as active specialist mandates to a pre-approved selection of three asset managers.

The percentage allocation of the total Fund's assets to the managers would be decided on via an optimization process, so as to minimize the mismatch and selection active risk included into the strategy with the incorporation of the managers' active management styles. The same asset class allocation (as shown in Table 4.4) will be used as the Fund's strategic benchmark allocation for both the balanced mandate and specialist mandate situations that are assessed.

4.2.2 Balanced Mandate Results

This section assumes that a pension fund would consider allocating balanced mandates to all its chosen asset managers. In order to establish the optimal allocation to each manager given the balanced portfolio that they propose the Fund invest in, five different situations were assessed to establish the levels of mismatch and selection active risk and return that are introduced by each. This analysis will enable a pension fund to allocate its assets in a manner that would yield the best possible return for the amount of additional risk that it is willing to take on.

As mentioned earlier, in order to obtain a balanced fund return for each of the chosen three asset managers, their specialist asset class returns were combined. Similarly, to obtain a balanced benchmark return for each manager, the relevant asset class benchmark returns were combined. The allocations used to combine the managers' specialist returns and the managers' benchmark returns have been varied to produce the different situations. The asset allocations used will be shown when discussing the relevant situation in more detail.

In the calculation of each manager's balanced return, for every month of the analysis period, the same asset allocation was used to calculate that month's balanced return. This results in the assumption that the managers' asset allocations were rebalanced monthly to the chosen allocation. Similarly the managers' benchmark asset allocations and the Fund's benchmark asset allocations were also assumed to be rebalanced monthly.

The expected return estimate for each manager was calculated from the monthly weighted active returns for each manager. First, each manager's monthly specialist asset class returns and their respective asset class benchmarks returns were weighted according to the weights derived from using the linear weighting methodology described in detail in the appendix. These weighted manager and manager benchmark asset class returns were then combined, using the different asset allocations for each situation, to obtain the monthly weighted balanced manager and manager benchmark returns. The monthly weighted active manager returns were calculated from these returns, by taking their geometric differences. The sum of these monthly weighted active returns yield the expected return estimates.

The geometric differences were calculated as shown in equation 4.1.

$$Active\ Return = \frac{(1 + MR)}{(1 + BR)} - 1 \quad (4.1)$$

Where MR represents the manager's return percentage and BR represents the benchmark return percentage.

The expected return estimates were then also adjusted by imposing the Black-Litterman views. The same Black-Litterman views were used for every balanced mandate situation, and are shown in Table 4.5.

	View	Confidence
Manager A outperforms Manager B	<i>0.40%</i>	<i>80.00%</i>
Manager A outperforms Manager C	<i>0.70%</i>	<i>70.00%</i>
Manager B outperforms Manager C	<i>0.30%</i>	<i>50.00%</i>

Table 4.5: Litterman Views taken for all the Balanced Mandate Situations

The covariance matrix estimates for managers was derived from the monthly original balanced active returns for each manager, as discussed earlier. These were used to calculate the Ledoit covariance matrix estimates.

As an overview, before the results of each situation are discussed, the efficient frontiers obtained for the five balanced mandate situations were plotted on the same graph so that their respective optimal solutions' active risk and return profiles could be compared. This can be seen in Figure 4.3.

It is immediately obvious from Figure 4.3 that the differences in the construction of the situations result in different optimal solutions for the most part. The reasons for the optimal solutions differences will become clearer during the analysis and discussion of each situation's results.

In the following situations where the managers asset allocations and their benchmark asset allocations have been varied, please note that the allocations used have been arbitrarily chosen for illustrative purposes in order to show their impact on the final optimal results. The allocations could well be entirely different in practice, but, for explanatory purposes, the allocations used were chosen so as to provide differing plausible situations. It is important to note that the results obtained for each situation are thus dependent on the allocations used, and could be entirely different should the allocations be drastically altered.

Situation A

For Situation A, the manager asset allocations and their benchmark asset allocations were chosen to be exactly the same as the Fund's benchmark asset allocation (see Table 4.4). All three managers and their benchmarks were assumed to have the same allocations as those shown in Table 4.6. Also it has been assumed that the managers' benchmarks are composed of the same underlying asset class benchmarks as the Fund's.

Situation A	Equity Allocation	Bond Allocation	Cash Allocation	U'lying Bench
Managers allocation	<i>24.01%</i>	<i>68.83%</i>	<i>7.16%</i>	
Managers Benchmark allocation	<i>24.01%</i>	<i>68.83%</i>	<i>7.16%</i>	<i>same</i>

Table 4.6: Manager and Manager Benchmark Asset Allocations for Situation A

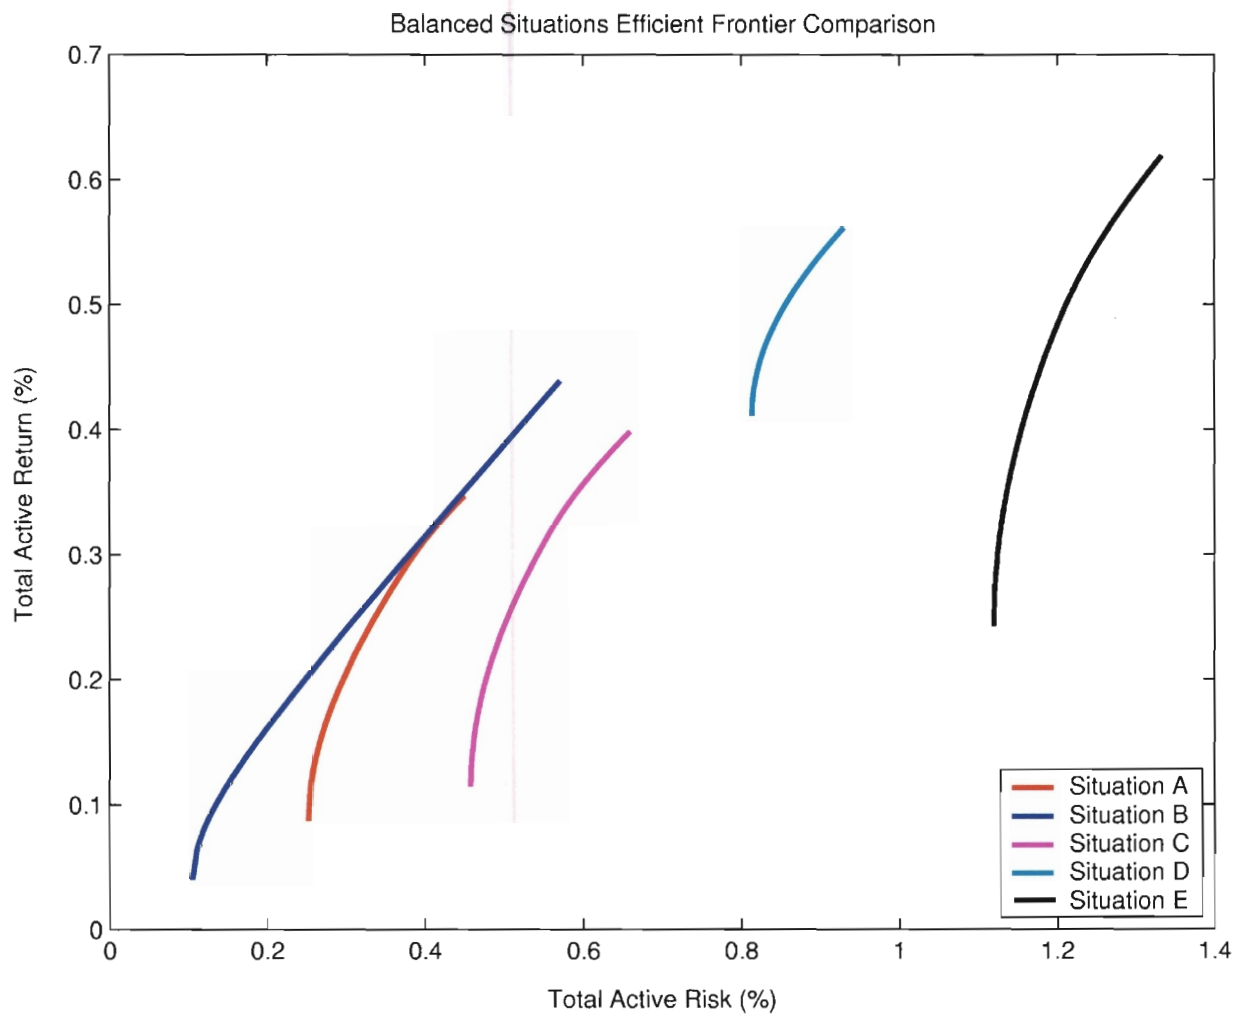


Figure 4.3: Efficient Frontier Comparison of the Balanced Mandate Situations

The expected return and covariance matrix estimates used in the optimization are shown in Table 4.7 and Table 4.8 respectively.

Manager A	Manager B	Manager C
<i>0.35%</i>	<i>0.12%</i>	<i>-0.13%</i>

Table 4.7: Expected Return estimates for Situation A

	Manager A	Manager B	Manager C
Manager A	0.00002030	0.00000421	-0.00000216
Manager B	0.00000421	0.00001780	-0.00000059
Manager C	-0.00000216	-0.00000059	0.00001819

Table 4.8: Covariance Matrix estimates for Situation A

The active risk and return characteristics of the optimal solutions available for Situation A have been shown by the Efficient Frontier in Figure 4.4, while each solution's optimal manager weightings are shown by the area graph below it. The third graph shows the solutions' total active risk breakdown, by percentage, into its component mismatch and selection risks, while the fourth graph shows the solutions' breakdown, by percentage of total active return, into mismatch and selection return. All these graphs have been drawn using total active risk as their x-axis values, so that they may be comparable for each level of total active risk. The results for the other four balanced mandate situations will also be represented in this manner.

The optimal manager allocations for Situation A's lowest risk level are 30%, 30%, and 40% to Manager A, Manager B, and Manager C respectively. As the risk levels increase, Manager A's allocation increases to 100% at the highest risk level, while the allocations of both Manager B and Manager C decrease to 0% at this same risk level.

Since Situation A uses exactly the same manager benchmark asset allocations and benchmark constituents as the Fund it would not introduce any mismatch risk or return into the Fund's investment strategy. This situation will however still incorporate selection risk and return into the Fund's strategy, as a result of the managers' active investment philosophies. The selection risk and return generated for this situation would account for only the managers' stock selection skills since the asset allocation effects are omitted since the managers' and their benchmarks' allocations are the same.

It is important to remember that while this situation does not contain any mismatch risk relative to the Fund's strategy, the selection risk introduced is still an additional risk that the Fund is required to manage. The Fund may choose to do so by imposing restrictions on exposures or by setting risk limits in the mandate that it ultimately gives to the managers. The managers are then compelled to carry out their investment decisions within those guidelines in an effort to contain the levels of active risk introduced into the overall Fund strategy.

Situation B

For Situation B, like Situation A, the manager and manager benchmark asset allocations were chosen to be exactly the same as the Fund's benchmark asset allocation (see Table 4.6). All three managers were assumed to have these same asset allocations. The only difference between Situation B and Situation A is that Situation B uses the FTSE/JSE All Share Index (ALSI) as its managers' equity benchmark instead

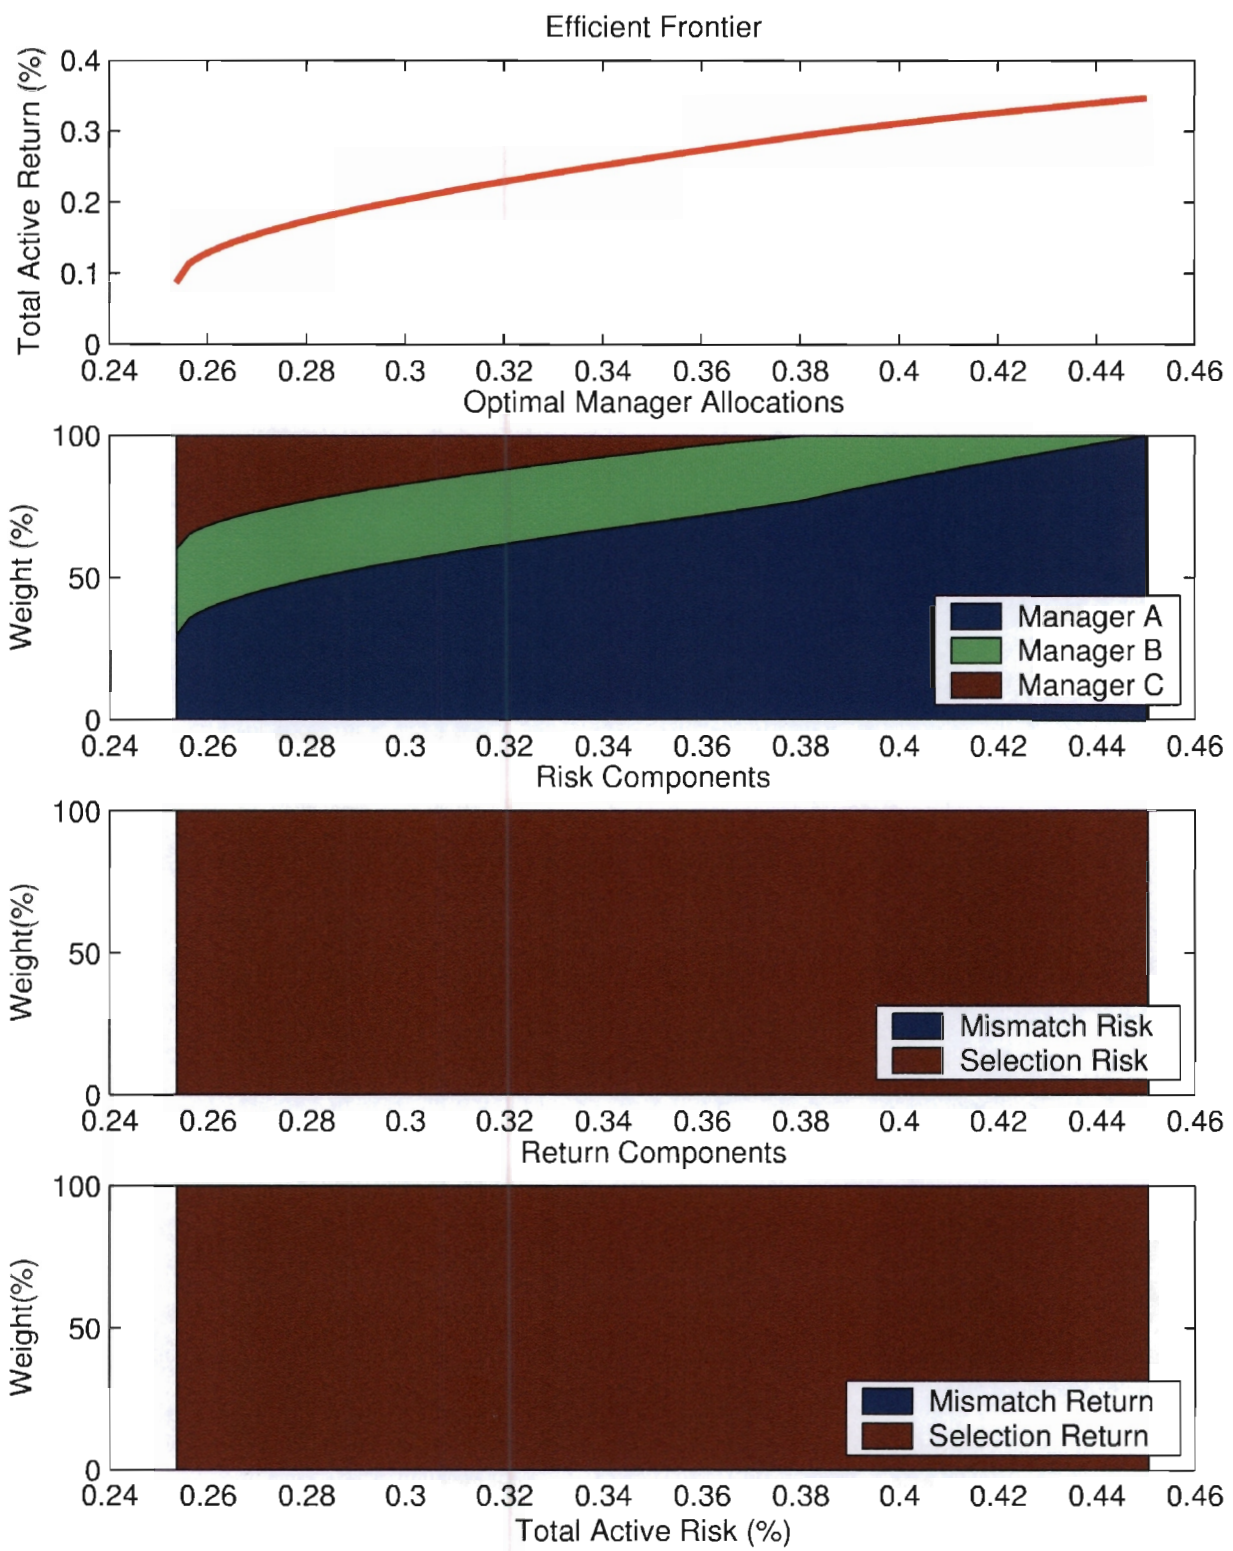


Figure 4.4: Situation A's Optimal Solutions and their Active Risk and Return Component Breakdown

of the ALSI50R. Its bond and cash underlying benchmarks, i.e. ALBI and AFMM, are the same as those used for the Fund, and also Situation A.

In practice, a Fund would not necessarily only choose an asset manager where the manager's internal benchmark exactly matches the Fund's. Should this kind of limitation exist, and assuming that only a small number of managers with the required benchmark exist, it would be impractical since it would severely reduce the overall number of asset managers available from which the Fund could choose. Further, this limitation could also possibly unnecessarily prejudice the Fund's return potential by excluding a large number of potentially excellent asset managers. With this in mind it was decided to introduce a different equity benchmark as a means of establishing the extent to which the optimisation results could be affected.

Tables 4.9, 4.10 and Figure 4.5 show the inputs and the subsequent results obtained for Situation B.

Manager A	Manager B	Manager C
0.44%	0.21%	-0.04%

Table 4.9: Expected Return estimates for Situation B

	Manager A	Manager B	Manager C
Manager A	0.00003269	0.00001522	-0.00000493
Manager B	0.00001522	0.00002842	-0.00000391
Manager C	-0.00000493	-0.00000391	0.00000257

Table 4.10: Covariance Matrix estimates for Situation B

From Figure 4.3, it can be seen that the risk and return ranges of Situation B's optimal solutions are wider than those of Situation A, as a result of the inclusion of a different equity benchmark to the Fund. Also on comparison of the optimal manager allocations for Situation A and Situation B, it can be seen that there are very distinct differences, as a result of the differences in the expected return and covariance matrix estimates for the situations.

This is expected given that we have already established that the optimisation is very sensitive to changes in its inputs, particularly its expected return estimates. The managers' benchmark would, in this situation, have different monthly and weighted average returns to those in Situation A. As a result, the managers' active returns, used to calculate the expected return inputs for the optimization, are very different to those in Situation A, and thus give rise to different expected return estimates.

For Situation B, the manager allocations at the lowest risk level (much lower than that of Situation A) are 11%, 9% and 80% to Manager A, Manager B and Manager C respectively. Manager A's allocation increases to 100% as the risk increases to its maximum, while the other two managers' allocations decrease to 0%. Also for Situation B, Manager C has a significantly greater weight and Manager B a significantly lower weight at each risk level than they have in Situation A.

For Situation B the introduction of the different equity benchmark affects just the selection component of active risk and return, since the mismatch component is only affected by changes in the asset class allocations. The underlying benchmark change thus does not affect the total active risk and total active return's component composition, and it remains at 100% in the selection component of each.

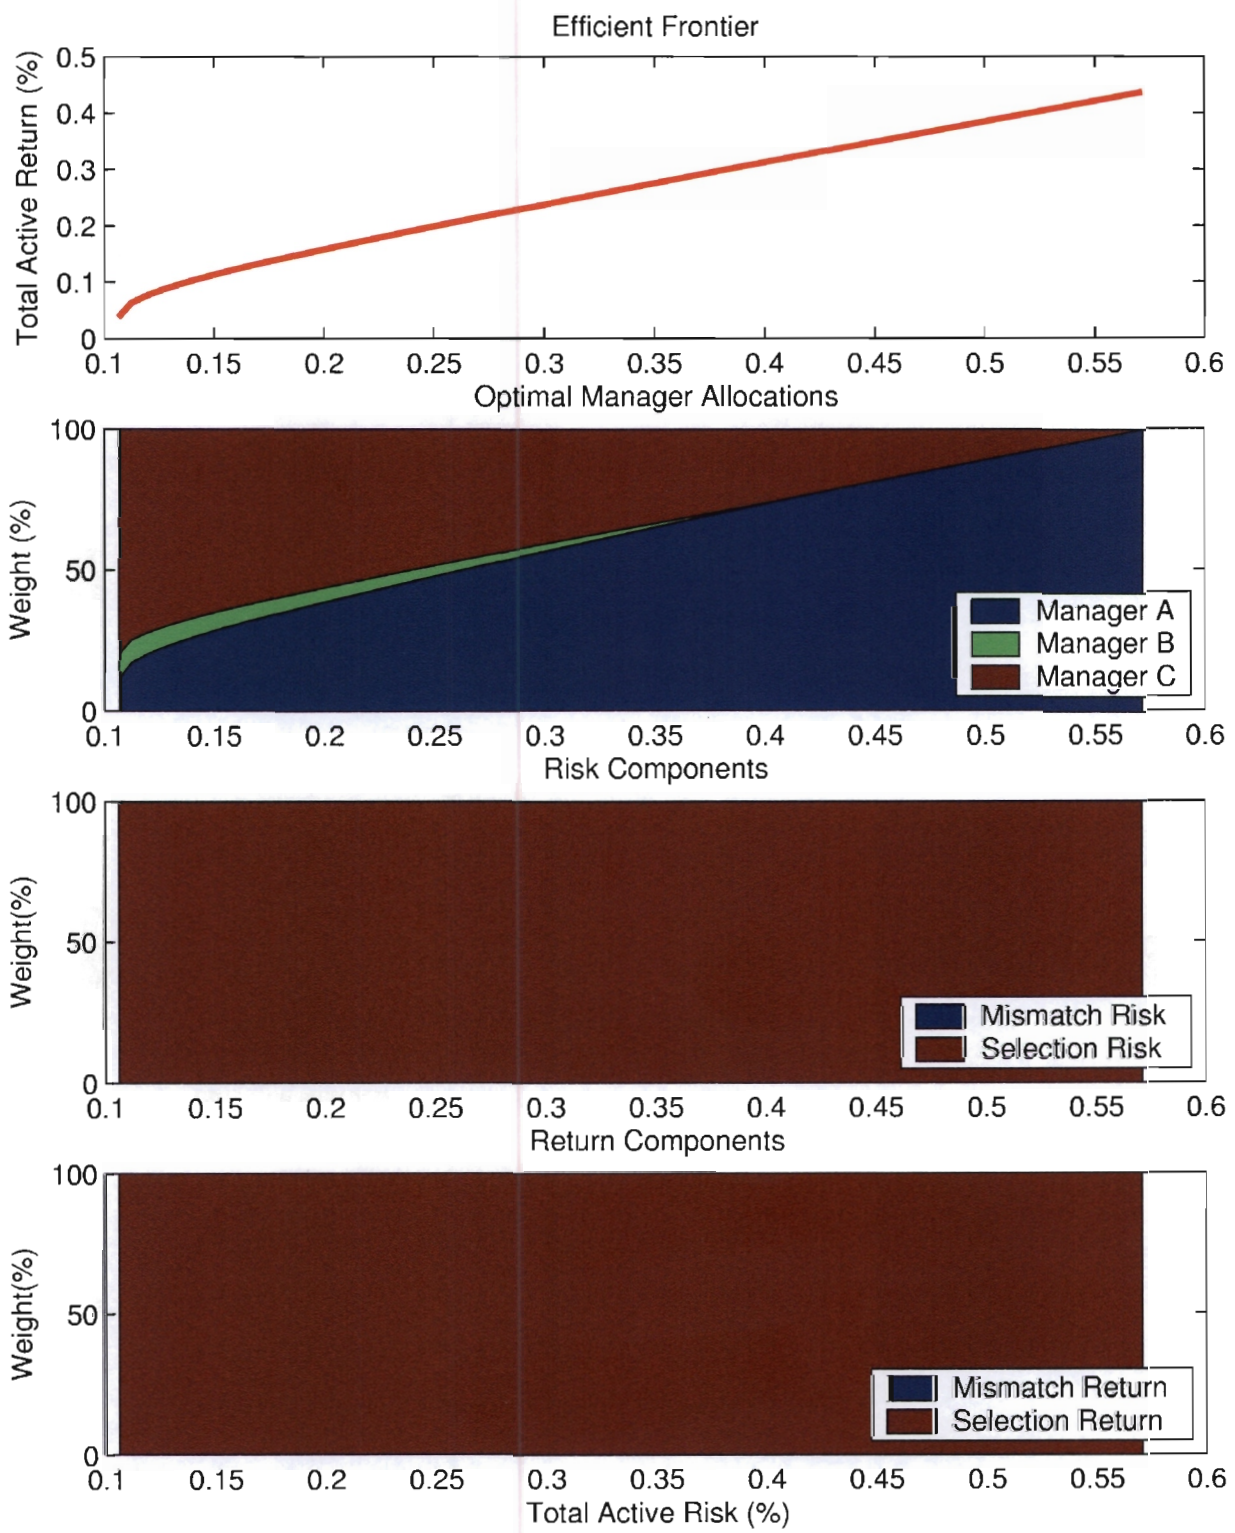


Figure 4.5: Situation B's Optimal Solutions and their Active Risk and Return Component Breakdown

Situation C

For Situation C, the manager and manager benchmark asset allocations were chosen to be different to the Fund's benchmark asset allocation, but the managers were assumed to have the same asset allocations as their benchmarks. All three managers were assumed to have the same asset allocations as shown in Table 4.11. Also, since the effect of having a different underlying benchmark to the Fund has been seen in Situation B, for the rest of the situations it has been assumed that the managers' benchmarks are composed of the same underlying asset class benchmarks as the Fund's.

Situation C	Equity Allocation	Bond Allocation	Cash Allocation	U'lying Bench
Managers allocation	30.00%	60.00%	10.00%	
Managers Benchmark allocation	30.00%	60.00%	10.00%	same

Table 4.11: Manager and Manager Benchmark Asset Allocations for Situation C

For this situation, it has been assumed that the Fund has allocated mandates that have a different manager benchmark asset allocation to the Fund's benchmark. This type of mandate will introduce a component of mismatch risk and return into the Fund's strategy, in addition to the selection risk and return that exists as a result of the managers' active investment styles.

Tables 4.12, 4.13 and Figure 4.6 show the inputs and the subsequent results obtained for Situation C.

Manager A	Manager B	Manager C
0.39%	0.13%	-0.14%

Table 4.12: Expected Return estimates for Situation C

	Manager A	Manager B	Manager C
Manager A	0.00003064	0.00000458	-0.00000463
Manager B	0.00000458	0.00002078	-0.00000208
Manager C	-0.00000463	-0.00000208	0.00002775

Table 4.13: Covariance Matrix estimates for Situation C

From Figure 4.3 it is apparent that Situation C's optimal solutions all have higher risk levels than those of Situation A, and in fact they share no common risk level. At Situations C's lowest risk level solution, which is already higher than Situation A's highest risk level, the manager allocations are 27%, 37% and 37% to Manager A, Manager B and Manager C respectively. At Situation C's highest risk level solution the optimal allocation is 100% to Manager A.

In keeping with its higher risk levels, Situation C also has higher return estimates than Situation A for each of its optimal solutions. This is expected given that the Fund would expect to adequately compensated, in the form of higher returns from Situation C's solutions, since these solutions introduce higher active risk levels into the Fund's overall strategy.

Thus far all three of the situations that have been looked at have had 100% allocation to Manager A at their respective highest risk levels. This would indicate that Manager A is the most risky asset manager of the three used in these analyses. This is also evident from its higher variance estimates in the covariance matrices. Manager A is however providing good returns for the Fund, as evident by its high expected return

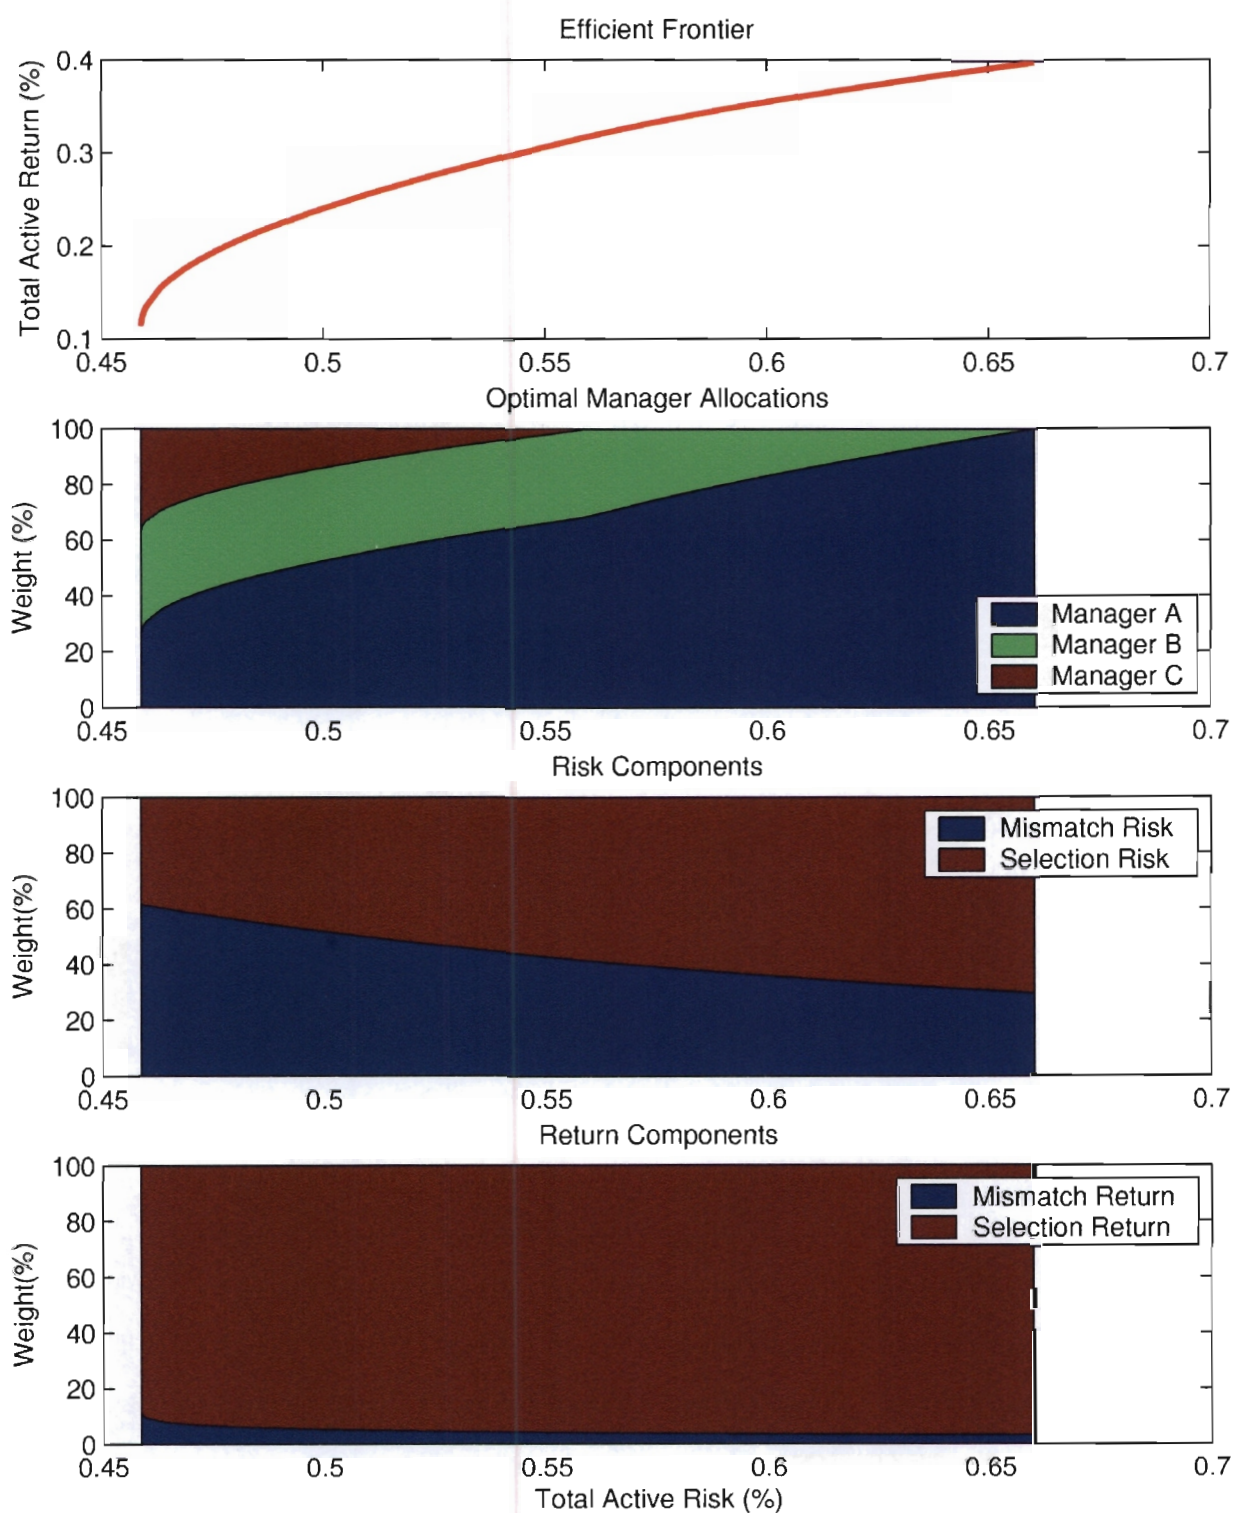


Figure 4.6: Situation C's Optimal Solutions and their Active Risk and Return Component Breakdown

estimates. From this it can be concluded that Manager A is providing better risk-adjusted returns at the higher risk levels relative to the other managers, thereby explaining its increasing allocations at the higher risk levels of each situation.

Also, when looking at the results shown in Figure 4.6, it is immediately obvious that Situation C is different from Situation A and Situation B in that its total active risk and return comprise both the mismatch and selection components. The mismatch risk and return components are introduced as a result of the difference in asset allocation between the Fund's benchmark and the managers' benchmark.

Looking at the risk components, it can be seen that mismatch risk is a significant portion of the total active risk for each of the optimal solutions, showing that the allocation differences between the managers' and the Fund's benchmarks should be an important consideration in a Fund's mandate construction.

For the return components however, the mismatch return component is far smaller than the selection return component. The results show that the differences in manager benchmark and Fund benchmark allocation in Situation C does not affect the Fund's returns as much as the selection return component does.

Situation D

For Situation D, the manager and manager benchmark asset allocations were chosen to be different to the Fund's benchmark asset allocation, and in addition the managers' asset allocations were assumed to be different to their benchmarks. All three managers were assumed to have the same asset allocations as shown in Table 4.14. Also it has been assumed that the managers' benchmarks are composed of the same underlying asset class benchmarks as the Fund's.

Situation D	Equity Allocation	Bond Allocation	Cash Allocation	U'lying Bench
Managers allocation	45.00%	50.00%	5.00%	
Managers Benchmark allocation	30.00%	60.00%	10.00%	same

Table 4.14: Manager and Manager Benchmark Asset Allocations for Situation D

For Situation D, like with Situation C, it has been assumed that the Fund has allocated mandates that have a different manager benchmark asset allocation to the Fund's benchmark. The difference between Situation C and Situation D however, is that it for Situation D it has been assumed that the managers' asset allocations are different from their benchmark asset allocations.

This difference in manager and manager benchmark asset allocations will bring in an element of the managers' asset allocation skills as well as their stock selection skills in attaining their active returns. In practice the managers' asset allocations would be continually changing, in keeping with their views and market conditions, and all managers would not typically have the same allocations. For this situation however, while not entirely practical, it is hoped that it will still be able to provide some idea of the returns generated from the managers' asset allocation decisions.

Tables 4.15, 4.16 and Figure 4.7 show the inputs and the subsequent results obtained for Situation D.

Manager A	Manager B	Manager C
0.55%	0.24%	-0.09%

Table 4.15: Expected Return estimates for Situation D

The expected return estimates for Situation D (in Table 4.15) are somewhat higher than those for Situation C (in Table 4.12), indicating that the change in the managers' allocations relative to their benchmark allocations has contributed positively to the managers' active returns.

	Manager A	Manager B	Manager C
Manager A	0.00007372	0.00003214	0.00005734
Manager B	0.00003214	0.00007551	0.00007968
Manager C	0.00005734	0.00007968	0.00019307

Table 4.16: Covariance Matrix estimates for Situation D

In Figure 4.3 it is apparent that Situation D's optimal solutions all have higher risk and return levels than those of Situation C, and in fact they share no common risk or return level. The change in the managers' asset allocations relative to their benchmarks has thus resulted in higher risk solutions with the compensating higher return levels.

Another distinct difference between Situation D and Situation C is that Situation D's solutions have allocated 0% to Manager C at every risk level. The optimal manager allocations are some combination of Manager A and Manager B for each risk level. The solutions range from an allocation, at the lowest risk level, of 47% to Manager A and 53% to Manager B, and as the risk level increases, Manager A's allocation increases to 100% at the highest risk solution.

On comparison of the results for Situation D and Situation C it is noticeable that Situation D's mismatch risk percentage is lower than Situation C's. This may be as a result of there being only two managers included in the optimal solutions. The mismatch risk that would have been introduced with the third manager is excluded, since its allocation is 0% for every level of risk. Also the selection components are greater due to the large difference between the managers' asset allocation and the managers' benchmark asset allocation generating a greater deviation from the benchmark's return, i.e. a greater active return.

Situation E

For Situation E, the manager and manager benchmark asset allocations were again chosen to be different to the Fund's benchmark asset allocation. The difference between Situation D and Situation E is, for Situation E each manager's asset allocation was assumed to be different to their benchmarks. The managers' benchmark asset allocations were all assumed to be the same, in keeping with the assumption that the Fund allocated to each manager the same mandate within which to apply their respective skills. The managers and managers' benchmark asset allocations are shown in Table 4.17. It has also been assumed that the managers' benchmarks are composed of the same underlying asset class benchmarks as the Fund's.

Situation E	Equity Allocation	Bond Allocation	Cash Allocation	U'lying Bench
Manager A allocation	<i>50.00%</i>	<i>45.00%</i>	<i>5.00%</i>	
Manager B allocation	<i>40.00%</i>	<i>55.00%</i>	<i>5.00%</i>	
Manager C allocation	<i>48.00%</i>	<i>50.00%</i>	<i>2.00%</i>	
Managers Benchmark allocation	<i>43.00%</i>	<i>50.00%</i>	<i>7.00%</i>	<i>same</i>

Table 4.17: Manager and Manager Benchmark Asset Allocations for Situation E

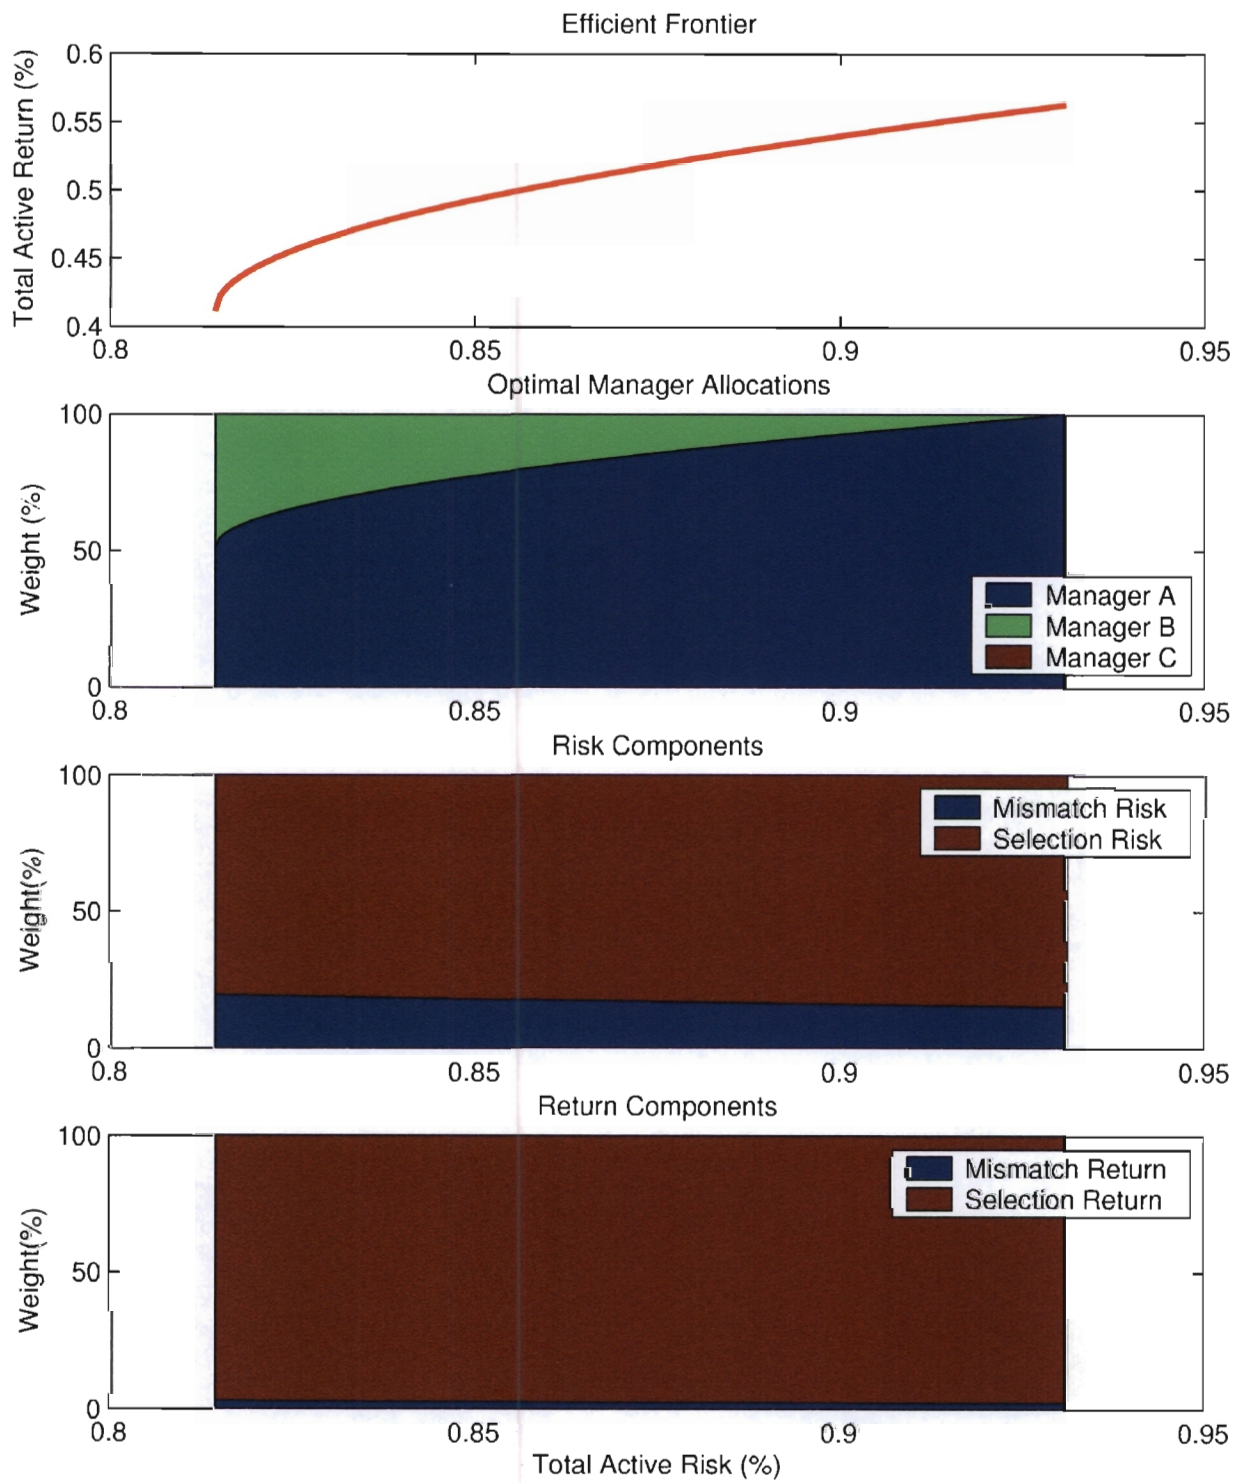


Figure 4.7: Situation D's Optimal Solutions and their Active Risk and Return Component Breakdown

Tables 4.18, 4.19 and Figure 4.8 show the inputs and the subsequent results obtained for Situation E.

Manager A	Manager B	Manager C
0.54%	0.17%	-0.12%

Table 4.18: Expected Return estimates for Situation E

	Manager A	Manager B	Manager C
Manager A	0.00006674	0.00000135	-0.00000258
Manager B	0.00000135	0.00003895	-0.00002102
Manager C	-0.00000258	-0.00002102	0.00009386

Table 4.19: Covariance Matrix estimates for Situation D

When looking at the efficient frontier comparison (Figure 4.3), Situation E's optimal solutions all have much higher risk levels and also a wider return range than the other situations and particularly Situation D. The managers' benchmark asset allocation, which is quite different from the Fund's benchmark asset allocation, is a major contributor to this. Since the managers' benchmark incorporates far more equity, which is considered the more risky asset class as a result of it's more volatile returns, and has less bonds and cash which are usually considered less risky; it is understandable that the solutions risk levels should be higher. In keeping with the concept of the risk/return tradeoff, to compensate for the solutions' higher risk levels the solutions return levels would also have to be higher. The managers' asset allocations which are bets around their benchmarks, still also contain much higher allocations to equity than the Fund's benchmark asset allocation, and thus also contribute to the higher risk solutions obtained.

The solutions' optimal manager allocations start at the situation's lowest risk level with an allocation of 21%, 51% and 27% to Manager A, Manager B and Manager C respectively, and as the risk increases Manager A's allocation increases to 100% at the situation's highest risk level, while the other two managers' allocations decrease to 0%.

On comparison of the results for Situation E and Situation D, it was noticed that the magnitude of the mismatch for risk and return are greater for Situation E. In Situation E, the mismatch risk is the largest component of the solutions' total active risk at every risk level solution. The mismatch risk percentage in this situation is far greater than that of Situation D because all three managers have received positive allocations for most of the optimal solutions. Also the extent of the differences between the managers' benchmark and the Fund's benchmark asset allocations is greater for Situation E than for Situation D thereby leading to a greater mismatch.

The mismatch return component for Situation E may also be similarly affected by the positive allocation to all three managers, and thus its greater mismatch return component than Situation D's is understandable.

Having discussed the impact of the different balanced mandate situations on the mismatch and selection components of risk and return and the varying optimal solutions that the situations give rise to, the specialist mandate situations and their impact on the optimisation results are now considered.

4.2.3 Specialist Mandate Results

Here it has been assumed that the pension fund would consider allocating only specialist asset class mandates to its chosen managers. For the specialist mandate option also, five different situations were looked at

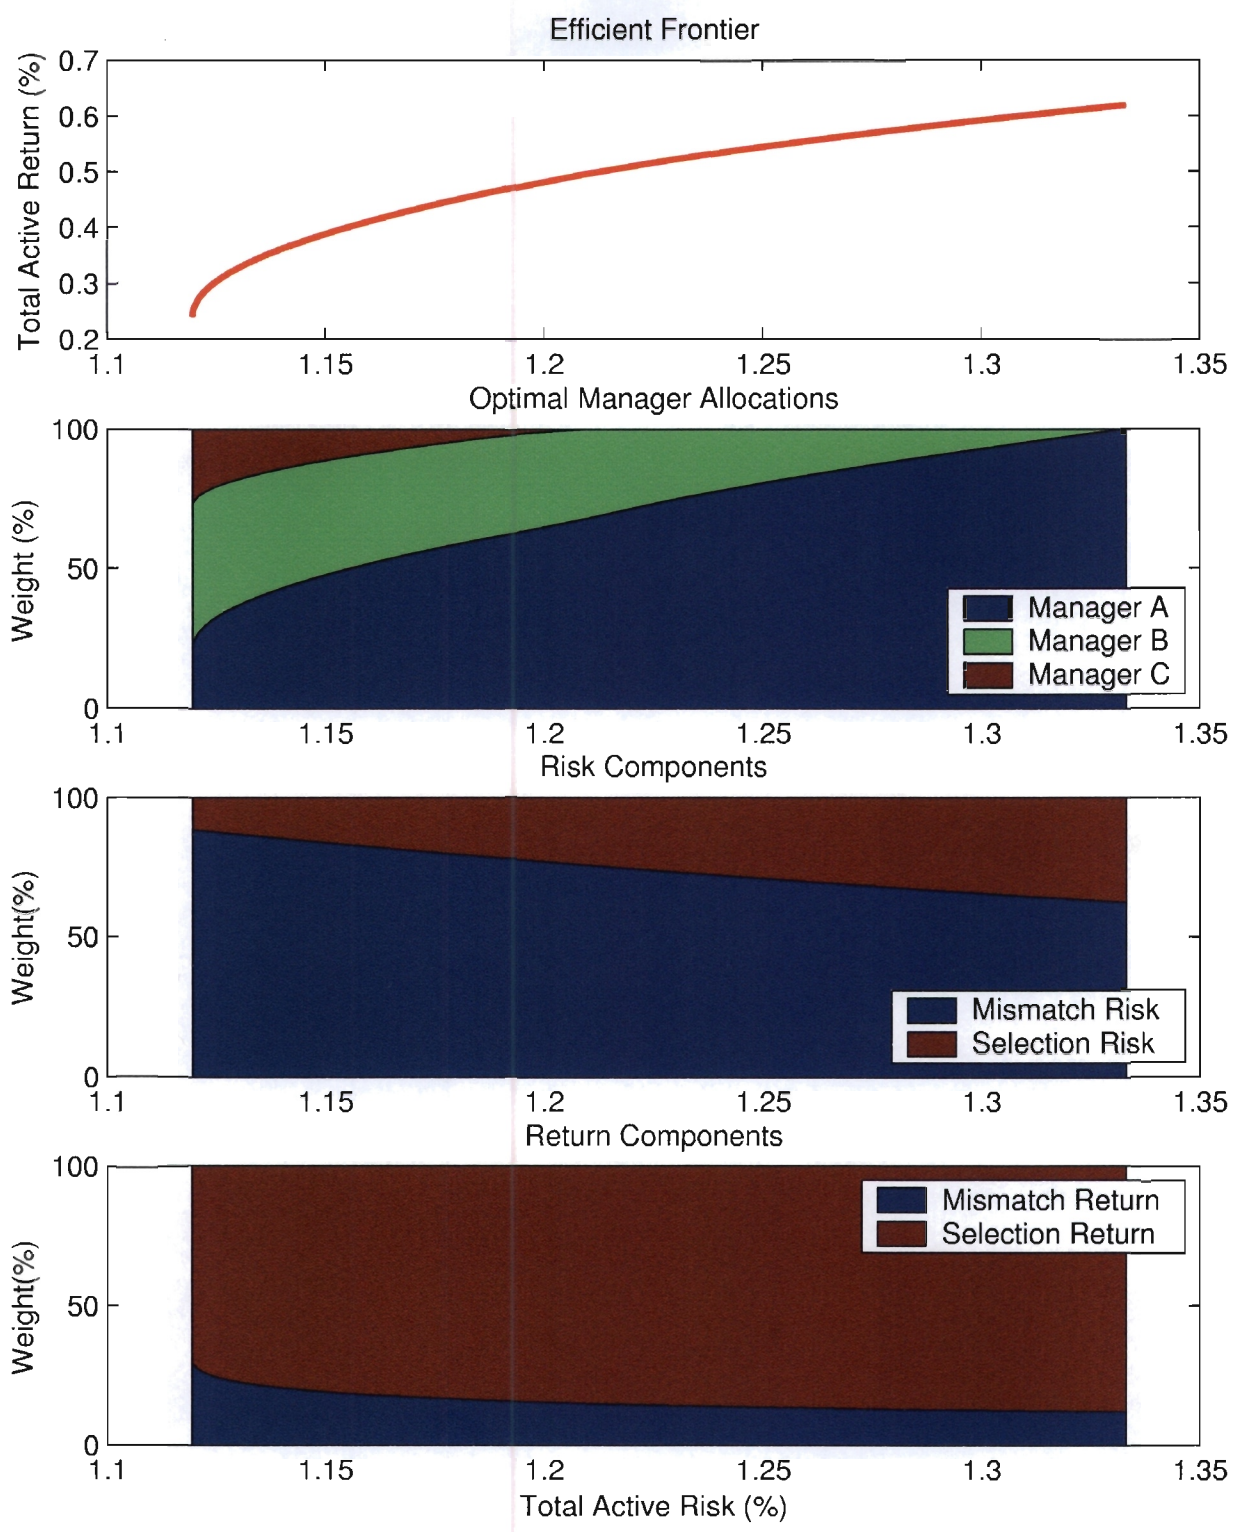


Figure 4.8: Situation E's Optimal Solutions and their Active Risk and Return Component Breakdown

to establish their impacts on the optimal asset manager allocation in terms of how they affect the mismatch and selection components of risk and return. The managers' specialist asset class returns were used in the optimisation. The managers' specialist asset class portfolio returns represent the managers' portfolio management skill for the respective asset classes, i.e. Equity, Bonds and Cash. The key difference between the specialist mandate situations and the balanced mandate situations is that the manager only receives an allocation if one of its specialist asset class portfolios receives a positive allocation via the optimization process.

Also should any of a manager's specialist asset class portfolios receives an allocation it implies that the manager has been allocated a mandate to manage a portfolio of assets in that asset class only. If a manager receives for example, an allocation to two of its specialist asset class portfolios, the manager is obligated to manage the money from each of these allocations entirely separately in their respective asset classes. The manager may not use or combine the money it receives for specialist portfolios to manage a balanced portfolio.

The expected return estimates for each manager's asset class portfolio was again calculated from the monthly weighted active returns for each manager. Each manager's monthly specialist asset class returns and their respective asset class benchmarks returns were weighted according to the weights derived from using the linear weighting methodology described previously. The monthly weighted active manager returns were calculated from these returns, by taking their geometric differences.

The sum of these monthly weighted active returns yield the expected return estimates. The expected return estimates were then also adjusted by imposing the Black-Litterman views. The views will be discussed during the detailed analysis of the situations' results.

The covariance matrix estimates for the managers were derived from the monthly original (unweighted) active returns for each manager. The geometric differences of the managers' and their respective benchmark's monthly returns resulted in the managers' monthly active returns, which were used to calculate the Ledoit covariance matrix estimates for the optimisation. In order to cater for the Managers' specialist portfolio names, a shorter naming convention was used. Manager A, Manager B and Manager C are represented by their associated letter, i.e. A, B and C which appears before the asset class type of the specialist portfolio, i.e. Equity, Bonds or Cash. For example, Manager A's specialist equity portfolio is called 'A-Equity' and similarly Manager C's specialist cash portfolio is called 'C-Cash'. This notation has been used throughout the discussion on the specialist mandate situations.

Since each manager's asset class portfolios are specialist portfolios, their benchmarks are also specialist asset class benchmarks, and not a combination of different asset class benchmarks. For example, for the managers' specialist equity portfolios, the portfolio benchmarks will be 100% equity. In the specialist mandate situations, depending on the optimal manager allocation chosen, the significant differences between the managers' portfolio benchmarks and the Fund's strategic benchmark allocation (as shown in Table 4.4) may introduce significant mismatch active risk and return into Fund's strategy. These mismatches will be assessed via the optimisation results.

The efficient frontiers obtained for the five specialist mandate situations were plotted on the same graph (Figure 4.9) so that their respective optimal solutions' active risk and return profiles could be compared.

From the comparison in Figure 4.9 it can be seen that the differences in the construction of the situations result in different optimal solutions.

For Situation 1, Situation 2 and Situation 3 it has been assumed that all the managers are skilled in all

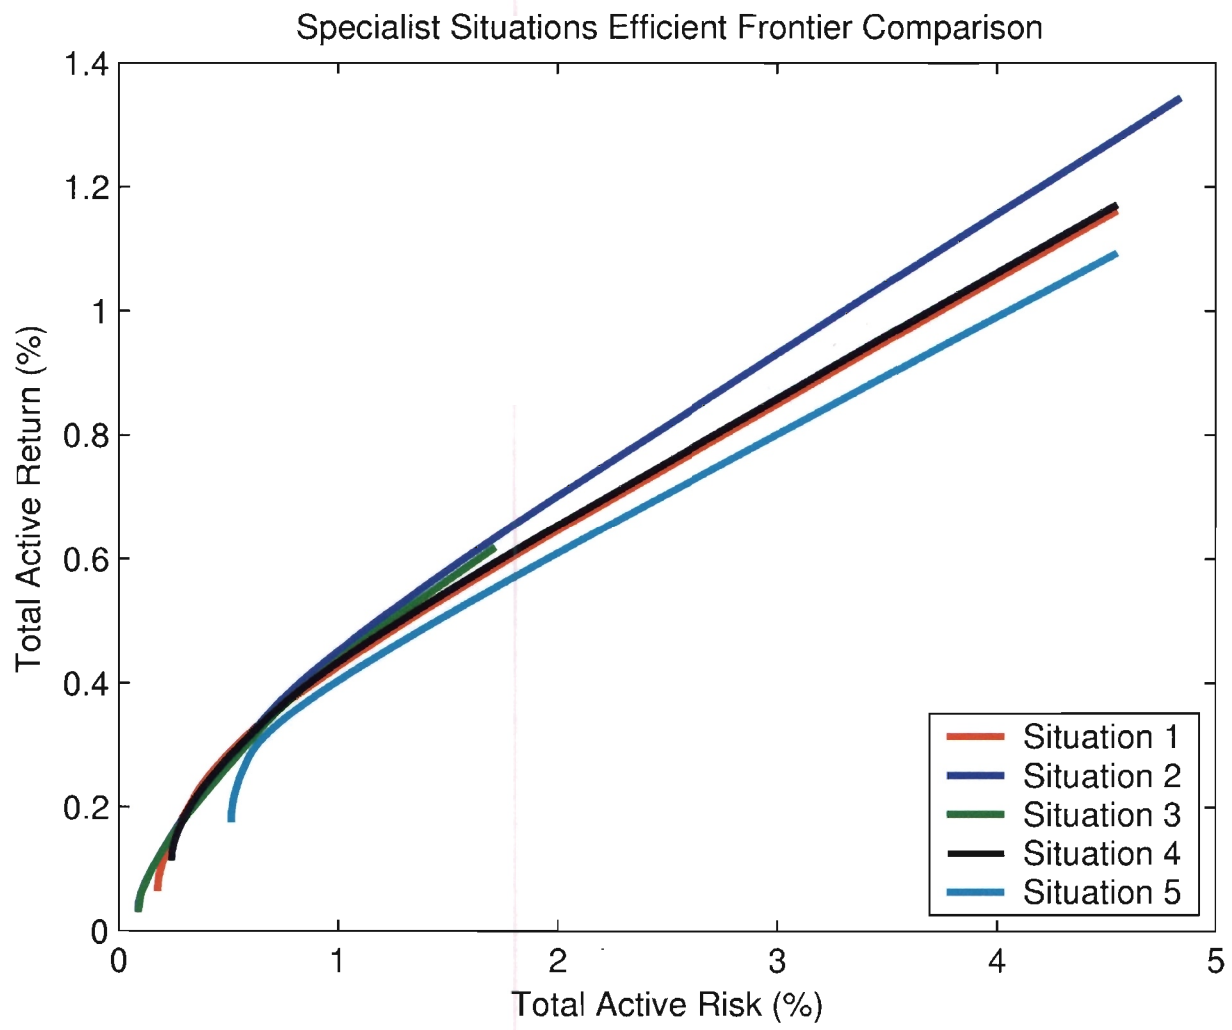


Figure 4.9: Efficient Frontier Comparison of the Specialist Mandate Situations

three asset classes. As such, each one of the three asset managers specialist asset class portfolios have been included as possible specialist portfolio choices for the Fund's strategy. The optimization for these situations is therefore required to provide an optimal allocation across nine specialist portfolios.

The Black-Litterman views used to adjust the expected return estimates for Situation 1, Situation 2 and Situation 3 are shown in Table 4.20.

	View	Confidence
A-Equity outperforms C-Equity	<i>0.50%</i>	<i>80.00%</i>
A-Equity outperforms B-Equity	<i>0.40%</i>	<i>60.00%</i>
B-Equity outperforms C-Equity	<i>0.40%</i>	<i>70.00%</i>
B-Bonds outperforms A-Bonds	<i>0.30%</i>	<i>60.00%</i>
B-Bonds outperforms C-Bonds	<i>0.35%</i>	<i>70.00%</i>
C-Cash outperforms A-Cash	<i>0.15%</i>	<i>70.00%</i>
C-Cash outperforms B-Cash	<i>0.10%</i>	<i>65.00%</i>
A-Equity outperforms B-Bonds	<i>0.60%</i>	<i>70.00%</i>
B-Equity outperforms B-Bonds	<i>0.30%</i>	<i>60.00%</i>

Table 4.20: Black-Litterman Views for Situation 1, Situation2 and Situation 3

For Situation 4, the three chosen asset managers are assumed to be skilled in only two of the three asset classes. The choice of the asset classes in which they are skilled is assumed to be have been established during the manager selection process. This optimization is therefore required to provide an optimal allocation across six specialist portfolios.

To incorporate further manager selection into the process, for Situation 5 each of the chosen asset managers is assumed to be skilled in the management of only one asset class. For this investigation it has been assumed that each manager is a specialist in a different asset class, so that no two managers obtain a mandate for the same asset class, and so that all three asset classes are covered. The optimization is thus required to provide an allocation across three specialist portfolios.

The situations have been broadly outlined here, but the assumptions used for each situation will be discussed in more detail with the individual results. Here again it is important to note that the results obtained for each situation are dependent on the assumptions made, and the results could be entirely different should the assumptions be changed. These particular situations have been chosen so as to provide some idea of situations that may arise in practice.

As with the balanced mandate situations, the results of the various specialist situations have been represented in a similar manner. The active risk and return characteristics of the optimal solutions available for the situations have been shown by an Efficient Frontier; while each solution's optimal manager weightings are shown by the area graph below it. The third graph shows the solutions' total active risk breakdown, by percentage, into its component mismatch and selection risks, while the fourth graph shows the solutions breakdown, by percentage of total active return, into mismatch and selection return. However, for the specialist mandate situations' a fifth graph has been introduced, showing the optimal asset class allocations, represented as an area graph. The total asset class allocations have been derived by summing each solution's relevant asset class portfolio weightings. All these graphs have been drawn using total active risk as their x-axis values, so that they may be comparable for each level of total active risk.

Situation 1

For Situation 1, each manager's asset class portfolio is considered a specialist portfolio and thus has its respective specialist asset class benchmark, i.e. equity, bonds or cash. For this situation it has been assumed that the specialist asset class portfolios have the same benchmarks as the underlying asset class benchmarks of the Fund's strategic benchmark. For equity the benchmark used is ALSI50R, for bonds the ALBI and for cash the AFMM.

The expected return and covariance matrix estimates used in the optimization are shown in Table 4.21 and Table 4.22 respectively. Please note that the covariances are shown in millions to accommodate the very small values. The optimal solutions for Situation 1 are shown, as described earlier, in Figure 4.10.

A-Equity	A-Bonds	A-Cash	B-Equity	B-Bonds	B-Cash	C-Equity	C-Bonds	C-Cash
0.84%	0.04%	0.02%	0.32%	0.06%	0.00%	-0.23%	0.01%	0.02%

Table 4.21: Expected Return estimates for Situation 1

The expected return estimates in Table 4.21 incorporate the Black-Litterman adjustment with the views as shown in Table 4.20.

	A-Equity	A-Bonds	A-Cash	B-Equity	B-Bonds	B-Cash	C-Equity	C-Bonds	C-Cash
A-Equity	332.19	0.20	-0.70	31.20	7.15	2.32	-85.71	-0.18	0.57
A-Bonds	0.20	2.80	0.06	-3.05	4.02	-0.31	11.81	0.35	-0.09
A-Cash	-0.70	0.06	0.33	-0.75	0.17	-0.05	1.99	0.06	0.02
B-Equity	31.20	-3.05	-0.75	149.00	5.61	1.36	-45.70	-2.96	0.05
B-Bonds	7.15	4.02	0.17	5.61	16.75	0.67	10.46	0.36	0.79
B-Cash	2.32	-0.31	-0.05	1.36	0.67	0.69	-4.10	-0.13	0.38
C-Equity	-85.71	11.81	1.99	-45.70	10.46	-4.10	290.06	1.63	-2.00
C-Bonds	-0.18	0.35	0.06	-2.96	0.36	-0.13	1.63	0.67	0.05
C-Cash	0.57	-0.09	0.02	0.05	0.79	0.38	-2.00	0.05	0.68

Table 4.22: Covariance Matrix estimates for Situation 1 ($\times 10^{-6}$)

It is interesting to note that while each manager's specialist asset class portfolios do receive a positive allocation in at least some of the optimal solutions generated, none of the optimal solutions, for this situation, have a positive allocation to each and every manager's specialist asset class portfolios. This shows that in this situation it is not optimal for the Fund to allow each manager to manage assets in every asset class, irrespective of the size of the allocations. This finding, while not without its limitations, contributes to the case for Fund's allocating specialist mandates rather than balanced mandates to their chosen asset managers. Of course this will also depend on the particular skills that have been identified for the managers.

Another interesting observation is that while the Fund's strategic benchmark is balanced, there exists optimal solutions for this situation where one or more of the asset classes receive a 0% allocation. For instance, some of the solutions with a total active risk level of less than 0.5% have varying allocations to all three asset classes, but as the risk increases beyond this risk level, and for the majority of the situation's solutions, the optimal allocations have a positive weighting to only equity and bonds. In fact the highest risk level solution has a 100% allocation to equity.

In keeping with these allocation changes it is also understandable then that the mismatch risk and return

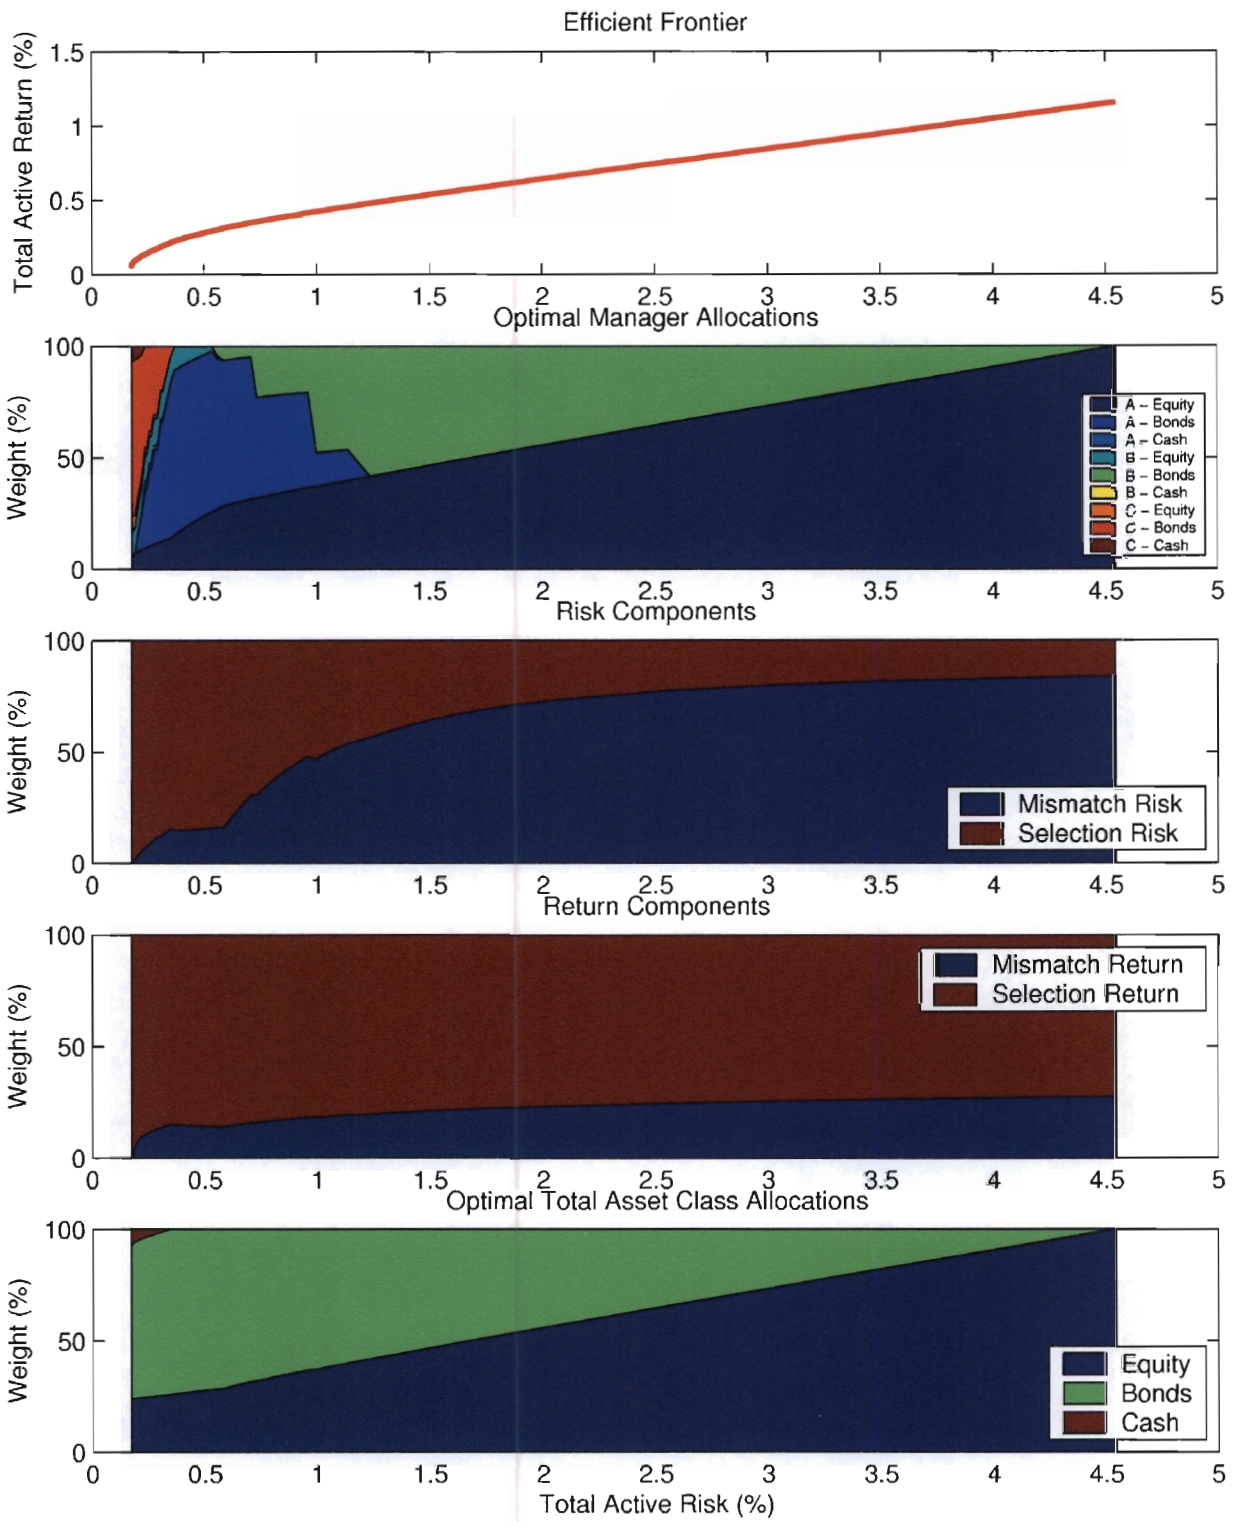


Figure 4.10: Situation 1's Optimal Solutions and their Active Risk and Return Component Breakdown

percentages increase as the solutions' risk levels increase. The mismatch risk and return increases as the solutions' allocations change to be only in equity and bonds and increases further as the bond allocation dwindles to 0%, while the equity allocation is 100%.

As the differences between the Fund's strategic benchmark allocation and the optimal solution's asset class allocations increase the mismatch component of risk increases drastically, while the mismatch component of returns increases far less drastically. In fact the returns are dominated by the selection component, even when the risk is dominated by the mismatch component.

Also when looking at the manager allocations it was noticed that for the most part the equity allocation and bond allocation is entirely to Manager A's specialist equity portfolio and Manager B's specialist bond portfolio respectively. This shows that the combination of just these two portfolios provides the best possible active return for the given level of active risk.

Situation 2

Situation 2 is similar to Situation 1 except that the managers' specialist equity portfolios have been assumed to have a different benchmark to the Fund's strategic underlying equity benchmark. The managers' equity portfolios' benchmark is assumed to be the ALSI, while the Fund's strategic equity benchmark is the ALSI50R. The managers' other asset class benchmarks have been assumed to be the same as the Fund's.

The expected return and covariance matrix estimates used in the optimization are shown in Table 4.23 and Table 4.24 respectively. The optimal solutions for Situation 2 are shown in Figure 4.11.

A-Equity	A-Bonds	A-Cash	B-Equity	B-Bonds	B-Cash	C-Equity	C-Bonds	C-Cash
1.02%	0.04%	0.02%	0.48%	0.06%	0.00%	-0.08%	0.01%	0.02%

Table 4.23: Expected Return estimates for Situation 2

As a result of the managers' different equity benchmark, it can be seen that the managers' specialist equity portfolios' expected return estimates have been affected relative to Situation 1's. Here again the expected return estimates in Table 4.23 incorporate the Black-Litterman adjustment with the views as shown in Table 4.20.

	A-Equity	A-Bonds	A-Cash	B-Equity	B-Bonds	B-Cash	C-Equity	C-Bonds	C-Cash
A-Equity	608.76	-9.47	-2.13	279.05	0.65	5.54	-96.11	-1.31	1.7
A-Bonds	-9.47	2.80	0.06	-12.40	4.19	-0.34	3.48	0.38	-0.09
A-Cash	-2.13	0.06	0.33	-2.12	0.17	-0.06	0.86	0.06	0.02
B-Equity	279.05	-12.40	-2.12	371.12	-0.10	4.43	-69.87	-4.34	1.08
B-Bonds	0.65	4.19	0.17	-0.10	16.75	0.72	4.05	0.39	0.84
B-Cash	5.54	-0.34	-0.06	4.43	0.72	0.69	-1.61	-0.14	0.41
C-Equity	-96.11	3.48	0.86	-69.87	4.05	-1.61	34.16	0.70	-1.20
C-Bonds	-1.31	0.38	0.06	-4.34	0.39	-0.14	0.70	0.67	0.05
C-Cash	1.70	-0.09	0.02	1.08	0.84	0.41	-1.20	0.05	0.68

Table 4.24: Covariance Matrix estimates for Situation 2 ($\times 10^{-6}$)

On comparison of Situation 1 and Situation 2's efficient frontiers (Figure 4.9) it is apparent that Situation 2's has wider return and risk range solutions than Situation 1. The change to the manager's specialist

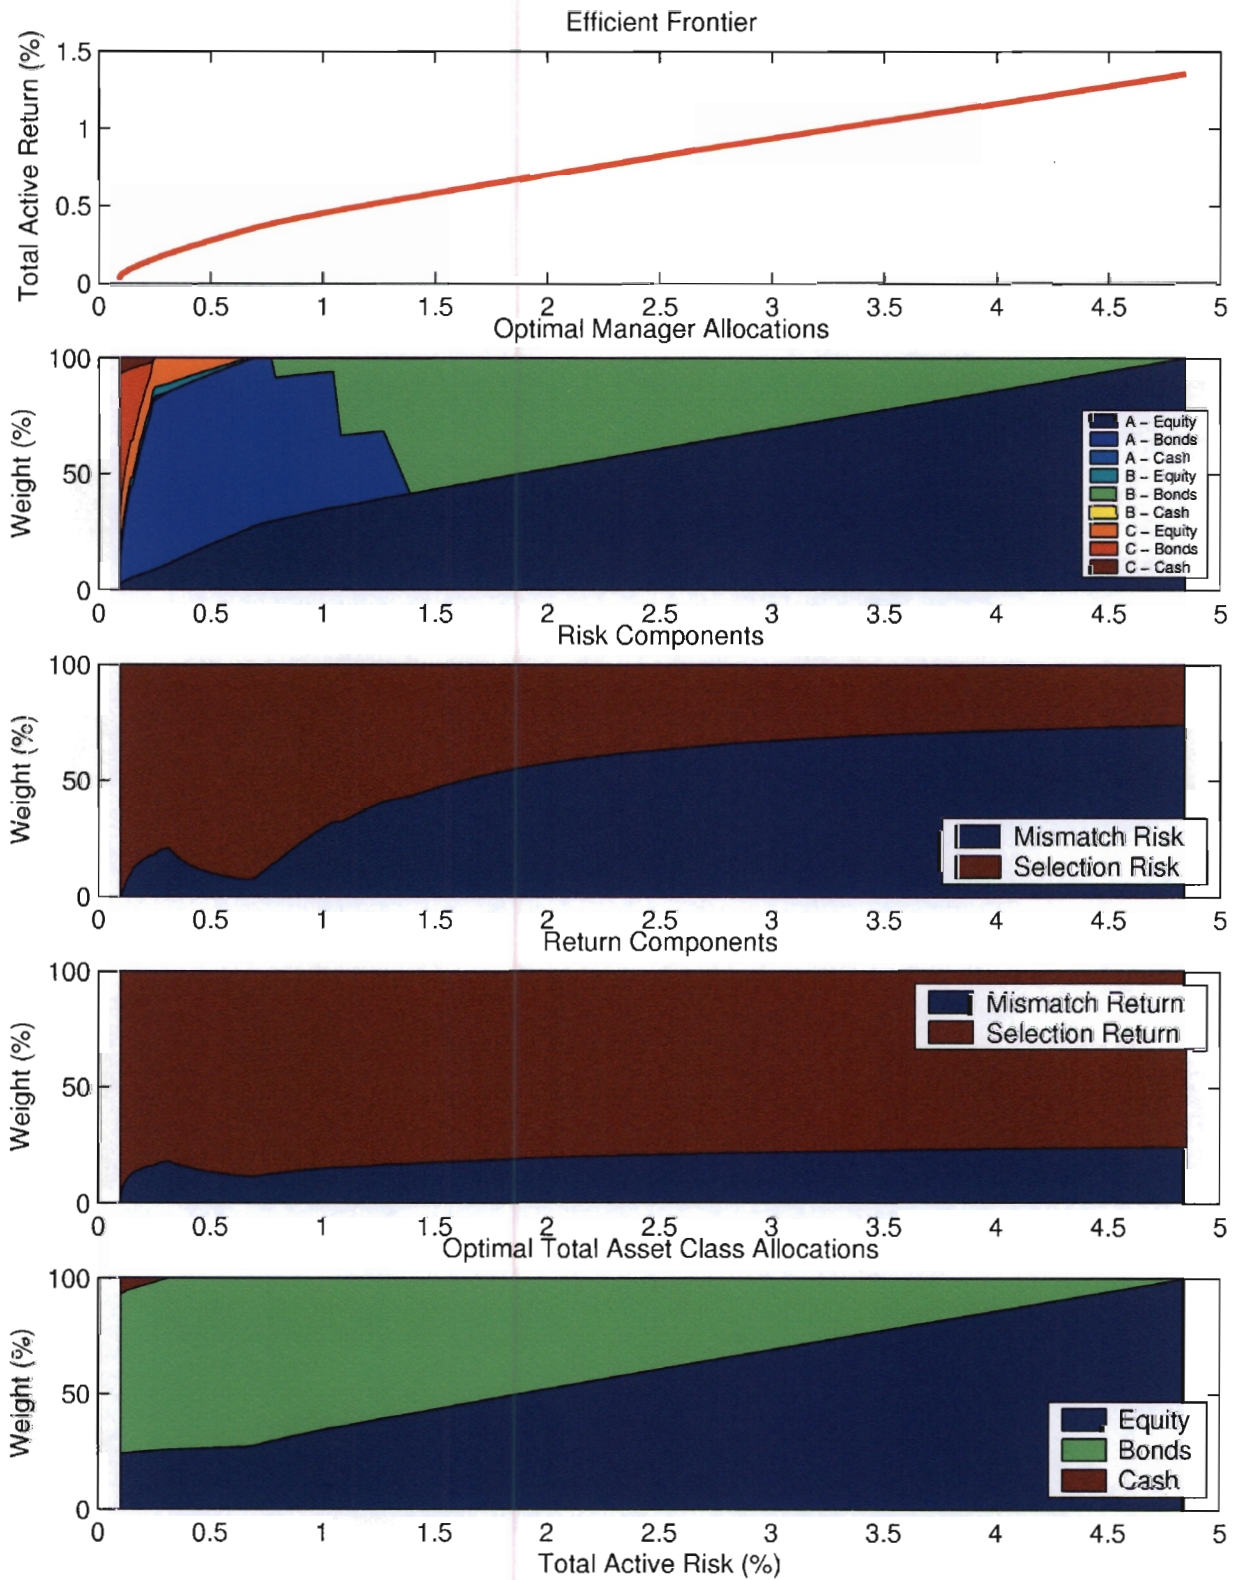


Figure 4.11: Situation 2's Optimal Solutions and their Active Risk and Return Component Breakdown

equity portfolios benchmark has an impact on the expected return estimates and the covariance matrix estimates, thus affecting the optimal solutions. It is also possible that the changes have occurred more as a result of the change in the expected return estimates rather than the changes to the covariance matrix estimates since we have earlier established that the optimization is more sensitive to changes in its expected return estimates.

From Figure 4.11 other noticeable changes in the results can be seen in the optimal manager allocations on the lower risk level solutions, i.e. lower than 0.5% risk. While the asset class allocations of Situation 1 and Situation 2, at these lower risk levels are somewhat similar, the allocations to the managers' specialist portfolios, making up the asset class allocation, are quite different. For example for Situation 2's solutions with a risk level of less than 0.5%, the allocation to Manager C's equity portfolio is far more noticeable than for Situation 1's solutions with similar risk levels.

Also, as with Situation 1, Situation 2's mismatch return and risk components increase as the solutions risk levels increase, again as a result of the allocation being split between bonds and equity only, for the most part. Here too, Manager A's specialist equity portfolio and Manager B's specialist bond portfolio are the portfolios of choice.

Situation 3

Situation 3 is similar to Situation 2 in that the managers' specialist equity portfolios have again been assumed to have a different benchmark to the Fund's strategic underlying equity benchmark. Here too, the managers' equity portfolios' benchmark is assumed to be the ALSI, while the managers' other asset class benchmarks have been assumed to be the same as the Fund's. However, the difference between these situations is that for Situation 3 some additional minimum asset class weighting constraints have been imposed on the optimization.

This has been done so that its effect on the optimal solutions may be assessed. Also given that the Fund's strategic benchmark is balanced so as to introduce some diversification into the strategy, it may not be prudent to allow an optimizer to allocate 100% of the Fund's assets to only one asset class, as is evident on Situation 1 and Situation 2's results. In practice the Fund may want to maintain some minimum amount of diversification by setting minimum asset class allocations.

The minimum asset class allocation applied to the optimization is shown in Table 4.25. These minimum allocations have been chosen arbitrarily but with the ultimate Fund's strategic benchmark allocation in mind.

	Allocation
Minimum Equity Allocation	<i>10%</i>
Minimum Bond Allocation	<i>50%</i>
Minimum Cash Allocation	<i>3%</i>

Table 4.25: Minimum Asset Class Allocations imposed on Situation 3's Optimisation

The expected return and covariance matrix estimates used in the optimization are shown in Table 4.26 and Table 4.27 respectively. The optimal solutions for Situation 3 are shown in Figure 4.12. As would be expected the expected return estimates for Situation 3 are the same as those for Situation 2.

From Figure 4.9, it is again immediately obvious that the limitation on the minimum allocations, introduced

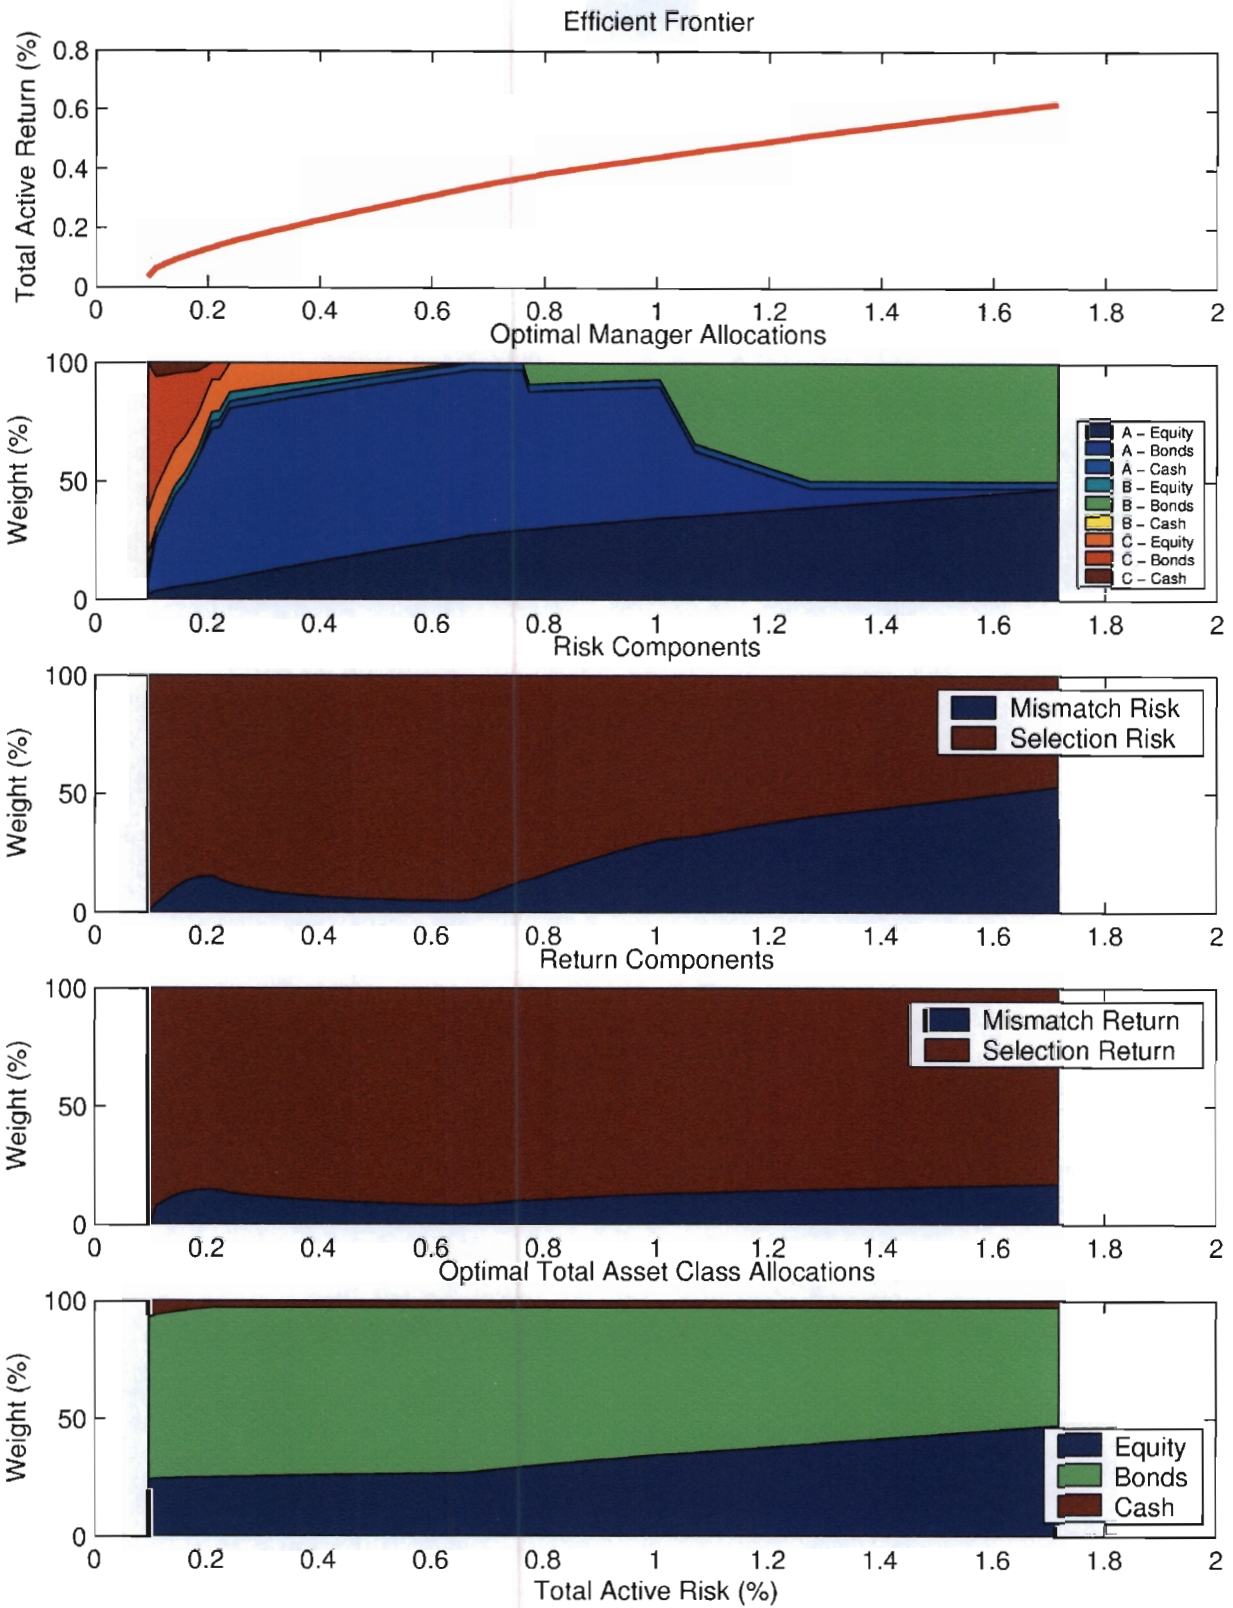


Figure 4.12: Situation 3's Optimal Solutions and their Active Risk and Return Component Breakdown

A-Equity	A-Bonds	A-Cash	B-Equity	B-Bonds	B-Cash	C-Equity	C-Bonds	C-Cash
1.02%	0.04%	0.02%	0.48%	0.06%	0.00%	-0.08%	0.01%	0.02%

Table 4.26: Expected Return estimates for Situation 3

	A-Equity	A-Bonds	A-Cash	B-Equity	B-Bonds	B-Cash	C-Equity	C-Bonds	C-Cash
A-Equity	608.76	-9.47	-2.13	279.05	0.65	5.54	-96.11	-1.31	1.7
A-Bonds	-9.47	2.80	0.06	-12.40	4.19	-0.34	3.48	0.38	-0.09
A-Cash	-2.13	0.06	0.33	-2.12	0.17	-0.06	0.86	0.06	0.02
B-Equity	279.05	-12.40	-2.12	371.12	-0.10	4.43	-69.87	-4.34	1.08
B-Bonds	0.65	4.19	0.17	-0.10	16.75	0.72	4.05	0.39	0.84
B-Cash	5.54	-0.34	-0.06	4.43	0.72	0.69	-1.61	-0.14	0.41
C-Equity	-96.11	3.48	0.86	-69.87	4.05	-1.61	34.16	0.70	-1.20
C-Bonds	-1.31	0.38	0.06	-4.34	0.39	-0.14	0.70	0.67	0.05
C-Cash	1.70	-0.09	0.02	1.08	0.84	0.41	-1.20	0.05	0.68

Table 4.27: Covariance Matrix estimates for Situation 3 ($\times 10^{-6}$)

for Situation 3, has resulted in its optimal solutions having significantly narrower return and risk ranges than those of Situation 2.

As a means of comparison, the optimal manager allocation solutions for Situation 2, which share the same risk levels as Situation 3 are shown in Figure 4.13. This allows the graphs to be compared on the same risk scale for ease of comparison.

By comparing the relevant graphs in Figure 4.11 and Figure 4.12 it is apparent that as a result of the minimum allocations imposed, the total active returns, total active risk and the asset class allocations of Situation 3 are different to those of Situation 2. These differences in turn have implications on the optimal manager allocations. It can be seen from Figure 4.13 that the optimal manager allocations for Situation 3 and Situation 2 are also different to each other for the same levels of risk.

Looking at the highest mismatch risk percentages of Situation 3 and Situation 2 at their respective highest risk level solutions, it is evident that while Situation 2's mismatch risk percentage is about 74%, Situation 3's mismatch percentage is about 53%. Here again the difference can be accounted for by the asset class allocation differences for the situations relative to the Fund's strategic benchmark allocation at these risk levels. While Situation 2 has an allocation of 100% equity, Situation 3 has an allocation of 47%, 50% and 3% to equity, bonds and cash respectively as a result of the minimum allocations imposed on its optimization process.

Situation 4

For this situation it has again been assumed that the managers' specialist asset class portfolios have the same benchmarks as the underlying asset class benchmarks of the Fund's strategic benchmark, i.e. ALSI50R, ALBI and AFMM for equity, bonds and cash respectively. Also no minimum asset class allocations have been imposed for this situation's optimization. However, Situation 4 differs from the previous three in that it has introduced a greater component of manager selection into the situation. This has been done by assuming that of the three chosen asset managers, each manager is only skilled in managing assets in two of the three asset classes underlying the Fund's strategic benchmark.

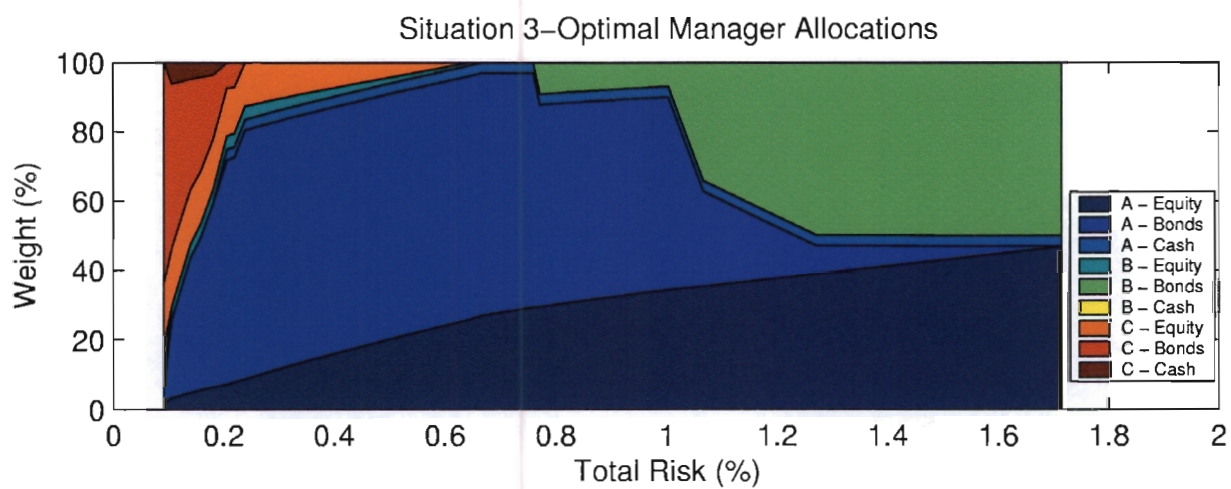
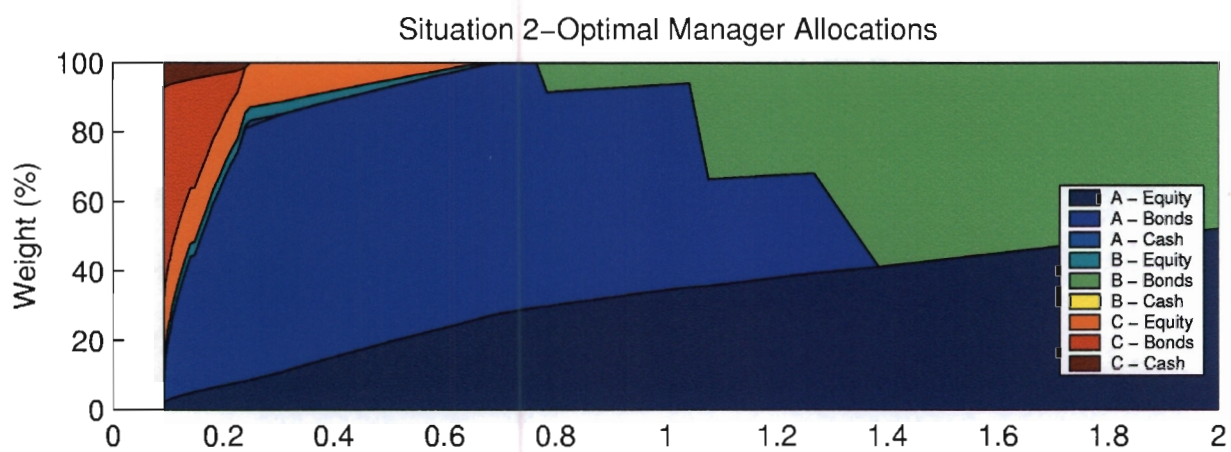


Figure 4.13: Situation 2's Optimal Manager Allocations relative to Situation 3's

The managers selected as candidates for each asset class are shown in Table 4.28.

	Choice 1	Choice 2
Equity	Manager A	Manager B
Bonds	Manager B	Manager C
Cash	Manager A	Manager C

Table 4.28: Situation 4's Manager Selection choices per Asset Class

Given the elimination of some of the managers' specialist asset class portfolios, the Black-Litterman adjustment views were accordingly altered and the views imposed are shown in Table 4.29.

	View	Confidence
A-Equity outperforms B-Equity	<i>0.40%</i>	<i>60.00%</i>
B-Bonds outperforms C-Bonds	<i>0.35%</i>	<i>70.00%</i>
C-Cash outperforms A-Cash	<i>0.15%</i>	<i>70.00%</i>
A-Equity outperforms B-Bonds	<i>0.60%</i>	<i>70.00%</i>
B-Equity outperforms B-Bonds	<i>0.30%</i>	<i>60.00%</i>

Table 4.29: Black-Litterman Views for Situation 4

The expected return estimates, adjusted according to the Black-Litterman views, and covariance matrix estimates used in the optimization are shown in Table 4.30 and Table 4.31 respectively.

A-Equity	A-Cash	B-Equity	B-Bonds	C-Bonds	C-Cash
<i>0.85%</i>	<i>0.02%</i>	<i>0.35%</i>	<i>0.07%</i>	<i>0.01%</i>	<i>0.02%</i>

Table 4.30: Expected Return estimates for Situation 4

The optimization performed for this situation would be required to provide an optimal solution by providing an allocation to these six specialist portfolios. The results for Situation 4 are shown in Figure 4.14.

When comparing the efficient frontiers of Situation 4 and Situation 1 in Figure 4.9, it is apparent that there seems to be very little difference between the risk and return ranges of these two situations' solutions. Situation 4's efficient frontier does however lie slightly above that of Situation 1, showing that it has slightly higher return estimates for each level of risk. At first this may be disconcerting, but when the effects of the additional manager selection are taken into account it becomes more understandable.

The two managers selected as managers for each asset class, were chosen as the best two of the three managers available. These choices were also further enhanced by the Black-Litterman views expressing the specialist asset class portfolio preferences. Assuming that the manager selection was done correctly and accurately, it makes sense that the optimization process for Situation 4 need not worry about taking the weakest manager for each asset class into account. This will allow for a better optimal allocation amongst the strongest two managers for each specialist asset class without being distorted by allocations to the weakest managers' specialist asset class portfolios.

For equity and bonds, from Situation 4's optimal manager allocations, it is clear that the preferred specialist asset class portfolios receive the majority of the asset class's allocation, i.e. A-Equity and B-Bonds. For cash though, of the allocation given to cash assets, 100% of the allocation was to A-Cash rather than C-Cash. Even though these specialist portfolios have the same expected return estimates, the variance of

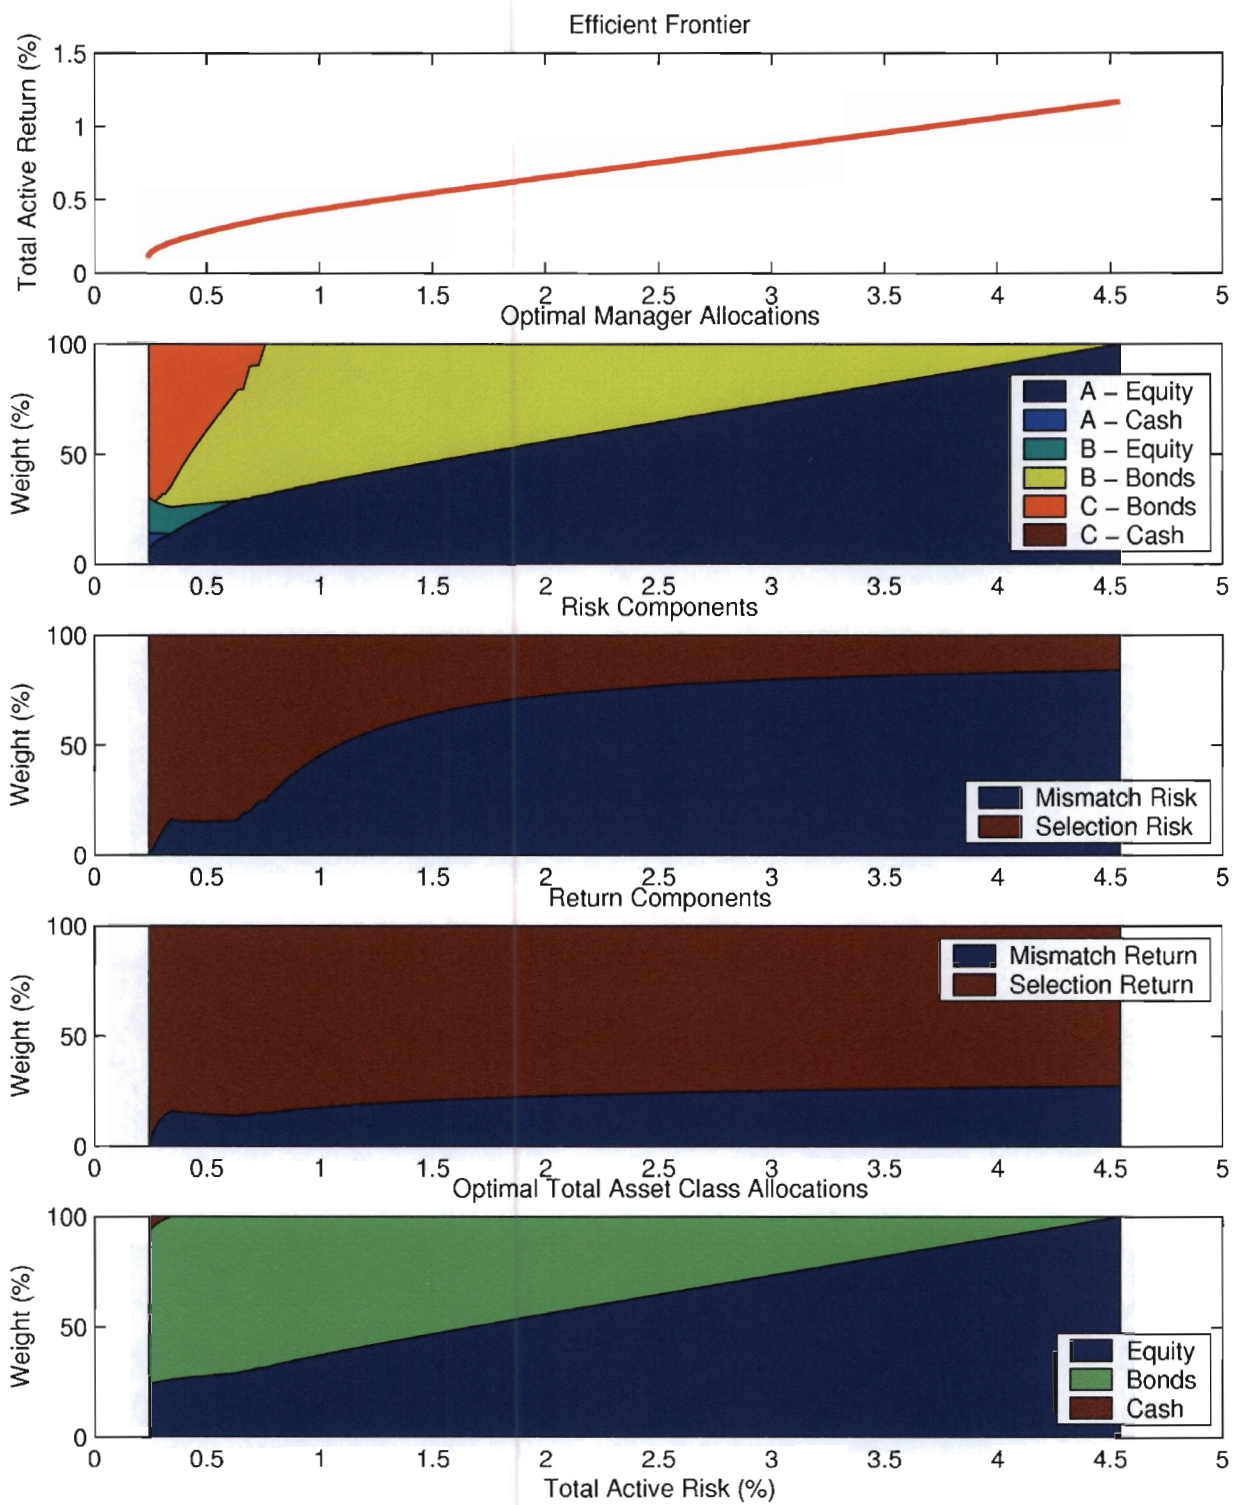


Figure 4.14: Situation 4's Optimal Solutions and their Active Risk and Return Component Breakdown

	A-Equity	A-Cash	B-Equity	B-Bonds	C-Bonds	C-Cash
A-Equity	332.19	-0.82	32.01	6.32	-0.24	0.64
A-Cash	-0.82	0.33	-0.84	0.17	0.06	0.02
B-Equity	32.01	-0.84	149	5.39	-3.17	0.08
B-Bonds	6.32	0.17	5.39	16.75	0.38	0.84
C-Bonds	-0.24	0.06	-3.17	0.38	0.67	0.05
C-Cash	0.64	0.02	0.08	0.84	0.05	0.68

Table 4.31: Covariance Matrix estimates for Situation 4 ($\times 10^{-6}$)

A-Cash is lower than that of C-Cash and thus the optimizers result is understandable.

Also for the majority of the risk level solutions, the bond and equity asset classes share the allocation, and again the bond allocation decreases to 0%, while equity increases to 10% at the highest risk level solution. Accordingly it is noticed that the mismatch components of risk and return increase as the risk levels of the solutions increase and the asset allocation becomes more biased to one asset class rather than some combination of the three.

Situation 5

Situation 5 is similar to Situation 4, except that the specialist asset class portfolios have been further limited to one manager per asset class, rather than two. This situation thus assumes further reliance on an initial manager selection process in order to narrow down the specialist portfolio choices. Here also the managers' asset class benchmarks are assumed to be the same as the Fund's underlying asset class benchmarks and no minimum allocations have been imposed on the optimization.

The manager selected for each asset class is shown in Table 4.32.

	Chosen Asset class portfolio
Equity	Manager A
Bonds	Manager B
Cash	Manager C

Table 4.32: Situation 5's Manager Selection choice per Asset Class

Again given the elimination of some of the managers' specialist asset class portfolios, the Black-Litterman adjustment views were accordingly altered and the views imposed are shown in Table 4.33.

	View	Confidence
A-Equity outperforms B-Bonds	0.40%	60.00%
A-Equity outperforms C-Cash	0.60%	70.00%
B-Bonds outperforms C-Cash	0.20%	70.00%

Table 4.33: Black-Litterman Views for Situation 5

The expected return estimates and covariance matrix estimates used in the optimization are shown in Table 4.34 and Table 4.35 respectively.

As a result of the increases manager selection assumptions the optimization performed for this situation is

A-Equity	B-Bonds	C-Cash
0.77%	0.07%	0.02%

Table 4.34: Expected Return estimates for Situation 5

	A-Equity	B-Bonds	C-Cash
A-Equity	332.19	5.83	0.61
B-Bonds	5.83	16.75	0.87
C-Cash	0.61	0.87	0.68

Table 4.35: Covariance Matrix estimates for Situation 5 ($\times 10^{-6}$)

required to provide an optimal solution by providing an allocation to just these three specialist asset class portfolios. The results are shown in Figure 4.15.

From Figure 4.9 it can be seen that Situation 5's efficient frontier lies below and to the right of Situation 4's. This implies that for a Situation 5 optimal solution with the same risk level as a Situation 4 optimal solution, Situation 4's solution can expect a higher return. This would suggest that while Situation 5's solutions are optimal, Situation 5 is inefficient relative to Situation 4.

This can be attributed to the fact that with Situation 5, because there is only one managers' specialist portfolio per asset class, the optimizer is compelled to allocate to that managers' specialist portfolio if it requires an allocation to the asset class. The asset classes' expected returns are thus limited by the manager's skill in that asset class and also the return benefit that could be achieved by varying the allocations to managers, when there is a choice of asset managers per asset class, has been removed from this situation.

Also for this situation, while the Fund does have some diversification in that it has allocated its funds to three different managers, when looking within an asset class: the Fund will not enjoy any diversification benefits. This is another very relevant type of risk that the Fund would have to be cognizant of in order to make a sound investment strategy decision.

Looking at the results in Figure 4.15, it is obvious that the optimal manager allocation and the optimal asset class allocation graphs will be identical due to there being only one manager's specialist portfolio per asset class. While the asset class allocation initially contains all three asset classes at the lowest risk level solutions, the mismatch components of risk and return increase significantly as the solutions' risk levels increase and as the asset class allocation changes to being ultimately 100% in equity.

4.3 Discussion of Results

For the purposes of this study a strategic benchmark allocation was chosen for the Fund and is shown in Table 4.4. The corresponding expected monthly risk and return characteristics of this chosen allocation, also shown in Table 4.4, amounts to an annualized expected risk and return estimate of 6.291% and 16.323% respectively.

The expected return value shows the average return for the Fund's strategic benchmark. The risk value gives us an indication of the percentage by which this average return may differ, by indicating the range, above or below the average, within which the Fund's strategic benchmark return could vary. Since the risk value encapsulates the risk of a one-standard deviation event, we could expect the Fund's strategic

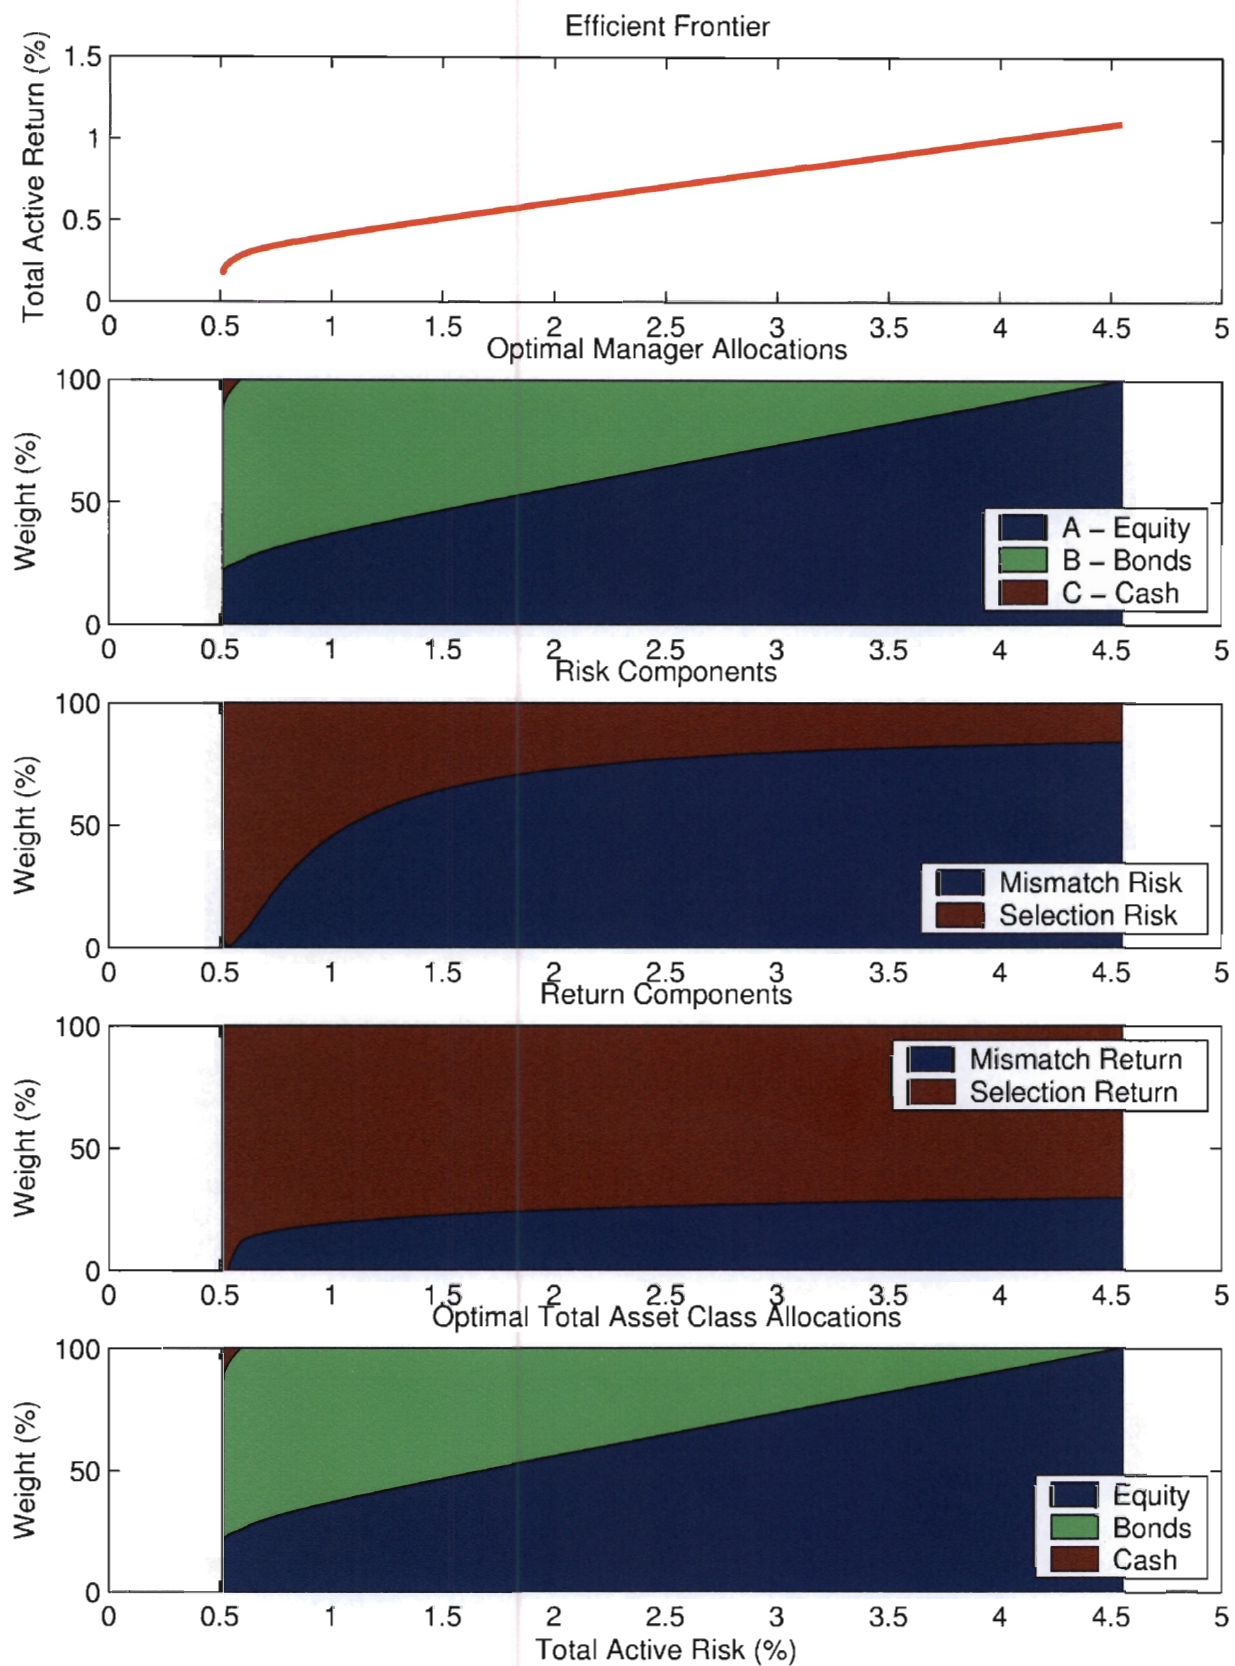


Figure 4.15: Situation 5's Optimal Solutions and their Active Risk and Return Component Breakdown

benchmark return to lie within this range 68% of the time.

Should the Fund choose to invest exactly in its strategic benchmark, it could expect to achieve a return of 16.323% per annum on average, but the Fund could expect its return to differ from this average by 6.291% or less, 68% of the time. This could result in the Fund's return estimate varying between 10.032% and 22.614%.

It is important to realize that this return band does not encapsulate the Fund's best and worst expected return estimates since it only deals with one standard deviation events. Should an event occur outside the one standard deviation range, it may result in an entirely different return estimate.

Given that it has been assumed here that the Fund chooses to pursue an investment strategy that incorporates active asset managers, the Fund's total risk and return estimates may be affected by the active risk and return introduced by the managers' active investment styles. The extent to which the Fund's total risk and return may be affected will depend on the ultimate mandate type and the subsequent manager allocation chosen.

While some assumptions were required for the purposes of this study, of the various balanced mandate and specialist mandate situations assessed in the previous section, all the situations were chosen to be plausible situations, that in principle, mirror real-life, practical Pension Fund situations.

In all the balanced and specialist situations assessed it is immediately obvious that, the Fund's incorporation of active managers into its investment strategy introduces active risk and return into the Fund's strategy. This is evident from the positive active risk and return levels for each of the various optimal solutions generated.

Furthermore, when looking at the make-up of the active risk and active return generated, in some situations the active risk and return comprised both the mismatch and selection components, while at other times it comprised just the selection component. As mentioned previously, the mismatch component is introduced when the managers' benchmark asset allocation differs from the Fund's benchmark asset allocation. The selection component is introduced immediately as a result of managers' active investment styles and is thus 100% of the active risk when there are no benchmark asset allocation mismatches.

For the situations' optimal solutions where both the mismatch and selection components are evident, it was noticed that even when an optimal solution's mismatch component dominates its active risk, the selection component always dominates its active return. This implies that the difference in asset allocation between the Fund's benchmark and the managers' benchmark has a greater effect on active risk than it does on active return.

Looking at the optimal manager allocations for the various situations, for most of the balanced situations there was always an optimal solution in which all the chosen managers received an allocation to manage a balanced portfolio, while for the specialist situations, none of the optimal solutions showed an allocation to all three specialist asset class portfolios of all three chosen asset managers. This observation may contribute to the argument for the allocation of specialist mandates rather than balanced mandates to a Fund's chosen asset managers, since even though the managers' specialist asset class returns were used to create the managers' balanced returns the specialist mandate situations did not yield any solution that allocated money to all three specialist portfolios for any of the chosen asset managers.

However, it must be recognized that this observation is specific to the balanced and specialist mandate situations looked at here. Even though the same specialist asset class returns were used to create the

managers' balanced returns, the asset allocations arbitrarily chosen to combine these returns into balanced returns may play a significant role in the outcome of the optimization and thus the resultant optimal solutions generated. The chosen asset allocations may have already been implicitly biased towards the specialist asset class in which the managers demonstrated the most skill, thereby resulting in the optimizer generating positive allocations to all these managers.

When considering the role of preliminary manager selection in the Fund's investment process, for the balanced mandate situations the manager selection process was limited to merely narrowing down the number of all potential asset managers, while for the specialist mandate situations slightly more selection was introduced in some of the situations. The preliminary manager selection process was assumed to have yielded the three asset managers used in the situations.

For the balanced mandate situations all three asset managers were used in every one of the five mandate situations looked at. For the specialist mandate situations however, in three of the situations all three specialist asset class portfolios for all three chosen asset managers were used and for Situation 4 and Situation 5 the list was further narrowed down, as explained previously.

By comparing the results for Situation 4 and Situation 5 it was noticed that Situation 5's solutions, based on a choice of just one specialist portfolio per asset class, may be sub-optimal relative to Situation 4's since they had lower active return levels for similar active risk levels. This may imply that narrowing down the choice of specialist portfolios too much, and in this case to just one managers' specialist portfolio per asset class, may be counter-productive. Giving the optimizer a choice of managers for a particular asset class may provide better active return and risk estimates for optimal solutions by incorporating the added benefit obtained from potential manager and style diversification.

The manager diversification could be obtained by just sharing the allocation amongst different asset managers, within or even across asset classes, while the style diversification would depend on the particular investment styles of the managers' specialist portfolios chosen within an asset class. For example if a Value manager's specialist equity portfolio was chosen together with a Growth manager's specialist equity portfolio as the portfolio choices for the specialist equity mandates, the optimisation may provide a better optimal solution in terms of its solutions' risk and return characteristics by giving both portfolios an allocation rather than being restricted to just one of them. Similar style diversification benefits may be obtained in each of the three asset classes being used.

Irrespective of the active risk and return level solutions obtained for the various situations it is important to realize that the active risk and return of the chosen manager allocation solution will provide a tolerance band around the risk obtained from the asset allocation process. Here the Fund's asset allocation risk amounted to 6.291% annualized. The higher (lower) the active risk and return levels of the chosen solution, the greater (less) the Fund's appetite is for deviating from its asset/liability matched risk and return profile, thereby affecting the Fund's overall expected risk and return levels.

As an example, let's look at one particular manager allocation solution, i.e. Situation 1's optimal solution that showed a 100% allocation to Manager A's specialist equity portfolio. This solution has a monthly total active risk of 4.5% (or 15.7% annualized) with a corresponding expected total active return of 1.2% (or 14.9% annualized). This level of annualised active risk is considered extremely high and is usually associated with very aggressive type mandates, usually used as a satellite portfolio rather than a core portfolio in a Core-Satellite Strategy. Further, since this solution has an allocation to just Manager A, the information ratio calculated using the risk and return numbers yield an information ratio of 0.9. This implies that Manager A can generate a return of 0.9 per unit of risk. This value is extremely high, given that in general

the average manager is assumed to be able to provide an information ratio of 0.5, and implies that Manager A is an above average or highly skilled manager.

When this solution's risk and return characteristics are combined with the Fund's underlying strategic asset allocation's risk and return characteristics, we see that the Fund's total risk amounts to 16.9% and its total expected return amounts to 17.523%. Given that this solution is considered a high risk solution, it may not be practical as a solution for a Pension Fund based on the assumption that the Fund is generally risk averse.

Chapter 5

Conclusion

The aim of this thesis to give an overview of the various steps in the risk budgeting process and decisions needed to ultimately allocate pension fund assets in the most optimal manner to investment managers. The thesis has demonstrated this approach for a variety of plausible situations that a fund may be confronted with.

Given that the detailed asset-liability modeling exercise is beyond the scope of this thesis, a simplistic assumption was made with respect to the possible outcome of such an exercise, and hence we define the needs of the pension fund in relatively simple risk and return terms. The asset class benchmarks to be allocated to each asset class were chosen to be the All Share Index with a 50% down-weighting to the resources sector, the All Bond Index and the Alexander Forbes Money Market Index for equities, bonds and cash respectively. Following the assumption for the asset class benchmarks, it was required to establish an optimal asset class allocation benchmark for the fund, in keeping with its particular risk tolerance. The Mean-Variance optimisation technique, using Microsoft Excel's Solver optimizer, was chosen to perform this task. The results of this exercise provided a range of optimal asset allocation solutions for different risk levels that the fund may be targeting.

After making an assumption of the pension fund trustees' risk tolerance relative to the liabilities, it was determined that a monthly expected risk level of 1.82% (or 6.29% annualised) with an expected return of 1.27% (or 16.32% annualised) was acceptable. This choice of risk tolerance gave rise to an optimal asset allocation of 24.01% to equity, 68.83% to bonds and 7.16% to cash which was used as the fund's strategic benchmark allocation for the following steps in the risk budgeting process. It has been noted however that the solutions presented in this thesis are particular to the specific situations, choice of techniques and inputs and procedures utilized for this analysis.

After the asset class benchmark and its specific allocations have been chosen, the next steps would be to decide on the investment strategy and on how to allocate the fund's assets optimally to its asset managers in order to get the best possible returns while keeping within the confines of the risk budget. Generally, in practice, not enough attention is paid to this part of the process to the detriment of the fund and its members by yielding a sub-optimal resultant structure. It is important to ensure that the investment strategy decided on and the subsequent manager allocations are given the same amount of consideration.

The strategy decided on should best suit the needs of the fund and the manager allocations made should be done in an optimal manner. Allocation should be carried out by keeping track of the mismatch and selection risk introduced into the process, as demonstrated in this thesis, and by ensuring that risk is kept

to the minimum level possible given the fund's pre-specified risk tolerance levels.

While for this analysis only the active balanced and specialist management strategies were explored, it may be necessary, given the fund's risk budget and certain strategies' cost implications with respect to management fees, to look at passive management strategies also, or even some combination of active and passive management strategies. This would in turn impact on the choice of asset managers and thereafter the subsequent allocation made to the chosen managers. The results obtained in this thesis could also be improved by using a better optimisation technique and improved estimates for the covariance matrix and expected returns.

In this thesis, a variety of plausible situations were analysed in order to assess the optimal manner in which to allocate the fund's assets to the chosen asset managers. Using the risk budgeting objective function described, the different levels of mismatch and selection risk (and return introduced by each situation) was assessed for a variety of risk levels and the optimal asset manager allocation for each established. By approaching the task of allocating the fund's assets to its chosen asset managers in this manner, the fund could expect that its investments will achieve the greatest return possibility relative to their indicated tolerable risk level. In so doing the fund should be able to prevent becoming significantly under-funded and thereby not being able to meet its liabilities.

Chapter 6

Glossary

Active Return (or Relative return) The difference between the actual return of a portfolio and the corresponding benchmark return.

Active Risk This is a measure of the volatility of a portfolio relative to a benchmark, which gives insight into how much the return of the portfolio will vary over time relative to the corresponding benchmark. This risk arises out of having different assets in the portfolio relative to the benchmark, comprised of market factors and specific factors related to the uniqueness of underlying asset bets.

Alpha Sometimes also called active performance, alpha is the expected residual return of a portfolio independent of market or benchmark performance. In true form however alpha differs from active performance due to a 'Beta' or market affect, which is subtracted from the active return. The two calculations then correspond only when the portfolio has no incremental sensitivity to either the market or benchmark, viz. $\beta = 1$.

Asset Allocation The % weighting or exposure of assets in an investment portfolio among different asset classes (equities, bonds, property, cash and overseas investments) Also known as Investment Mix. This mix is usually the largest determinant of absolute risk or risk relative to the liabilities of the fund, guidelines for the structure generally being ascertained through an Asset Liability Study (ALM).

Benchmark A relevant reference portfolio for active or passive fund management. This benchmark will capture the required performance and risk characteristics of the portfolio and indicate the universe of available yet appropriate investments.

Bet Size (Bet) The Bet Size indicates the impact of the positioning of the portfolio in terms of exposure to a variety of assets and portfolio characteristics / classifications relative to the benchmark. Bet size then arises wherever exposures differ and can be underweight (negative bet) or overweight (positive bet)

Capital Asset Pricing Model (CAPM) The simplest version states that the expected excess return on securities will be exactly in proportion to their systematic risk coefficient, or beta. The CAPM implies that total return on any security is equal to the risk-free return, plus the security's beta, multiplied by the expected market excess return.

Efficient Frontier The set of portfolios, one for each level of expected risk, with maximum expected returns. Investors would select the appropriate efficient portfolio given their risk/renewal requirements. Different efficient frontiers exist for portfolios with different constraints e.g. legislation etc.

- Excess Return** Return relative to the risk free (or commonly cash type) return. Risk free assets are determined by looking at government (or no (low) risk of default) assets of appropriate term and nature.
- Information Ratio** Commonly, the ratio of annualized expected relative return to relative risk. Information ratios are most often used to measure a manager's performance in terms of both risk and return relative to a benchmark or other measure, e.g. the inflation rate. It shows the expected relative return achieved per unit of active risk.
- Modern Portfolio Theory (MPT)** The theory of portfolio optimization, which accepts the risk/reward tradeoff for total portfolio return as the crucial criterion. Derived from Markowitz's pioneering application of statistical decision theory to portfolio problems, optimisation techniques and related analysis are increasingly applied to investments.
- Monte-Carlo Simulation** Iterative numerical process of determining outcomes to a random process. Allows one to assess risk using pre-modeled distributions under a variety of scenarios. Particularly useful when selecting the 'best of' strategy or portfolio i.e. optimization, or assessing worst case risks and scenarios for any given portfolio. Used extensively in optimization and ALM models.
- Optimisation** Finding the best solution among all the solutions available for consideration within a problem using an interactive or linear programming technique to consider the many potential solutions available. In portfolio construction, given a certain view on return and the cost of risk in pursuit of return, optimization finds the fine balance of the correct portfolio where return is maximized for every level of risk. This is achieved through selection of optimal portfolios given a risk aversion or utility function.
- Passive Management** Managing a portfolio to match (not exceed) the return of a predefined benchmark, consensus or market index. This form of management is generally low cost, efficient and allows stable relative performance. Passive management generally involves fairly complex modelling, however is generally an objective replicable process. Choice of index or benchmark is crucial and carries with it high benchmark risk. Also know as indexation.
- R-Square** A statistic usually associated with regression analysis, where it describes the fraction of observed variation in data captured by a model. R-square varies between 0 and 1 with a higher value usually indicating a model with higher explanatory power.
- Residual Return (Specific Return)** Return independent of the benchmark or market performance. An asset's residual return equals historic alpha or its excess return minus beta times the benchmark excess return.
- Residual Risk (Specific Risk)** The risk or tracking error (annualized standard deviation) of the residual return.
- Risk Aversion / risk tolerance / risk appetite** An investor's attitude to risk is influenced by the impact of risk on the finances of the investor, as well as age, wealth, dependents etc. This quantifies the ability and desire of an investor to take risk and is molded by the consequences of the risks actually occurring. Best quantified by looking at the maximum loss an individual can bear, financially and emotionally.

- Risk Free Return** The return achievable with absolute certainty, including negligible risk of default. In the U.S. market, short maturity treasury bills exhibits effectively risk free returns. In South Africa cash instruments exhibit most of these characteristics but do have an element of default risk.
- Sharpe Ratio** A statistical measure which attempts to show the performance of a portfolio's return in risk adjusted terms. It is calculated by dividing the portfolio's excess return (over the risk-free rate) by the risk, i.e. standard deviation of portfolio returns. The higher the Sharpe Ratio, the better the portfolio's return in risk adjusted terms. A Sharpe Ratio higher than one can be considered to be very good; while a ratio below 0.1 shows that the portfolio has been poorly rewarded for the risk taken.
- Sortino Ratio** Similar to the 'Sharpe Ratio', except it uses downside deviation for the denominator, whereas Sharpe uses standard deviation. Also reference point for return comparison and downside deviation can be changed to any appropriate minimum return requirement e.g. cash returns (as in Sharpe Ratio), inflation, zero (capital protection) etc.
- Standard Deviation** A statistical measure of the dispersion of a set of numbers around the mean. If the standard deviation is small, the frequency of distribution is concentrated within a narrow range of values. For a normal distribution, about two thirds of the observations will fall within one standard deviation of the mean, approximately 95 percent of all outcomes fall within approximately 2 standard deviations, and approximately 99 percent of all outcomes fall within approximately 2.5 standard deviations. Standard deviation is a commonly used measure of risk because the higher the standard deviation the higher the uncertainty of the return. It therefore helps explain the distribution of potential returns and allows one to create confidence intervals for return observations
- Stock Selection** The selection of an individual security within an asset class or sub sector of the asset class. For example, stock selection in relation to equity investments is made after analysing the financial standing, future earnings prospects and valuation of the shares of the company concerned. Along with asset allocation, stock selection is a key way in which investment managers add value.
- Systematic Return** The part of the return dependent on the benchmark or market return. We can break excess returns into two components: systematic and residual. The systematic return is the beta times the benchmark excess return. See residual return.
- Systematic Risk** The risk (annualized standard deviation) of the systematic return.
- Tracking Error** A measure of the closeness with which a portfolio follows a representative market index (benchmark). Technically the tracking error is represented by the standard deviation of the differences in return between the portfolio and the index (benchmark). Tracking error measures the likelihood, based on historical data or an equity risk model, of actual returns differing from index (benchmark) returns.

Appendix A

Analysis of inputs used in the implementation of the optimisation

A.1 The Optimisation Technique

It was decided to utilise the Mean Variance optimisation technique for the various steps in the risk budgeting process. While some, possibly better, types of optimisation techniques do exist, mean variance optimisation was chosen because it tends to be the most widely recognised and understood optimisation technique.

The formula used for the optimisation, as explained in the previous chapter and shown by equation 3.6, can be seen again below.

$$\begin{aligned} \text{Max } U &= \mu^T x - \lambda x^T \Sigma x \\ \text{subject to } e^T x &= 1 \end{aligned}$$

An asset class allocation optimization aims to establish the optimal weights required for each asset class, represented above by the vector (x). In order to establish these weights, the inputs, i.e. the expected return (μ) and the covariance matrix (Σ) are a key component of the optimization, as explained previously.

The asset class allocation optimization was used to establish the effect of the inputs on the optimization results. Various types of the input estimators, as discussed in chapter 2, were utilized so that the effect on the optimal weights generated by the optimization could be demonstrated and the best estimators chosen for the application of the risk budgeting process.

The inputs used in the differing scenarios are as follows:

	Covariance Matrix Estimator	Expected Return Estimator
Scenario 1	Sample covariance matrix	Average return
Scenario 2	Weighted Sample covariance matrix	Weighted Average return
Scenario 3	Sample covariance matrix	Average return with Black Litterman adjustment
Scenario 4	Ledoit covariance matrix	Average return

Table A.1: Breakdown of Inputs used in Optimisation Scenarios

The following 2 sections will describe the differences obtained while performing the asset class allocation optimization using these different scenarios.

The Mean-Variance optimisations were performed using the Solver add-in to Microsoft Office Excel.

A.2 The Covariance Matrix

The effect of the covariance matrix estimator on the optimization was established by varying only the type of covariance matrix estimator used, and then assessing the resulting optimal solutions that were generated. The sample covariance matrix estimator was used as a base case against which to assess the effect of some of the different estimators mentioned in the previous chapter.

Three of the scenarios, as outlined in Figure A.1, are relevant to assessing the impact on the optimization by changing the covariance estimator. The relevant scenarios that will be further discussed here are Scenario 1, Scenario 2 and Scenario 4.

The Excel Data Analysis Toolpak was used to generate the Sample Covariance matrices. The Observation Weighted Sample Covariance Matrix was calculated from its formula, using the relevant Excel functions. The Ledoit Covariance matrix was generated using code as obtained from the Ledoit homepage [12]. This code was implemented using Matlab functionality to generate the Ledoit covariance matrix.

Please note that for all the efficient frontier comparisons, only the efficient part of the curve is used, i.e. when $\lambda \geq 0$.

A.2.1 The Sample Covariance Matrix and Average Return case output

Scenario 1 will be used as the base case against which to compare the other scenarios.

Scenario 1

Here the covariance matrix and expected returns inputs used for the mean variance optimisation are, the sample covariance matrix and average returns respectively.

The average returns over the 3 year period ending October 2004 for the three different asset class benchmarks are given in Table A.2 and the sample covariance matrix generated from these asset class benchmark returns is shown in Table A.3.

ALSI50R	ALBI	AFMM
<i>1.29%</i>	<i>1.05%</i>	<i>0.87%</i>

Table A.2: Table of estimated Expected Return Values for Scenario 1

	ALSI50R	ALBI	AFMM
ALSI50R	0.002476	-0.000093	-0.000020
ALBI	-0.000093	0.000413	0.000009
AFMM	-0.000020	0.000009	0.000003

Table A.3: The estimated Covariance Matrix for Scenario 1

Based on the various risk tolerance levels, many optimal weighting solutions are obtained using these inputs. The respective risk and return numbers for each of these optimal solutions are plotted to obtain an efficient frontier, as shown in the first graph in Figure A.1. The second graph in Figure A.1 shows an area plot of each of the optimal solutions obtained for the scenario's various risk levels.

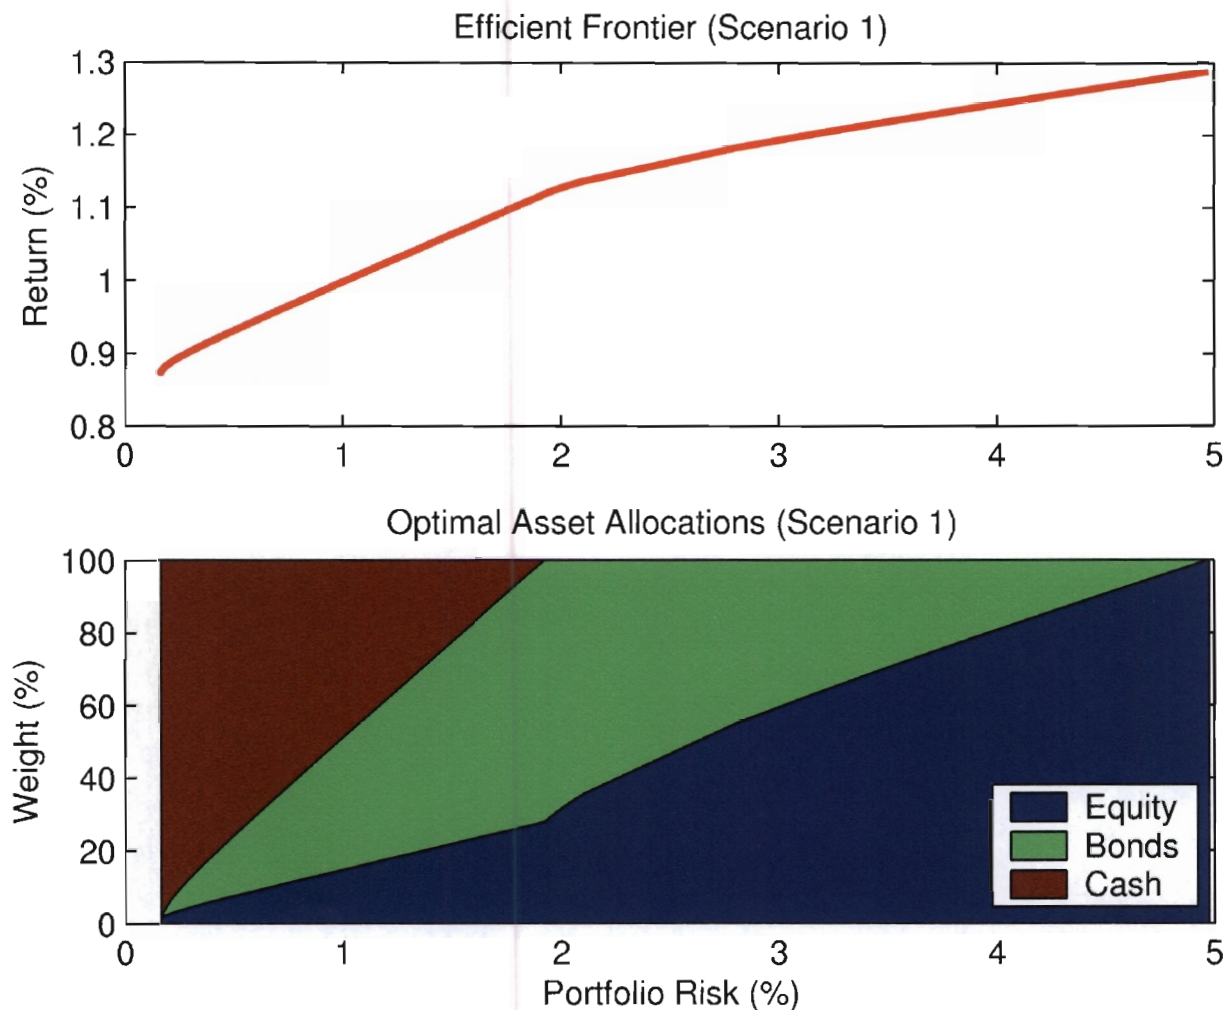


Figure A.1: Efficient Frontier and Optimal Asset Allocation Solutions for Scenario 1

In keeping with the assumption that rational investors expect to be rewarded for the extra risk that they take on, it can be seen from efficient frontier above that the expected return increases as the level of risk increases. Also, as evident in area graph, the optimal weightings change from being 100% in Cash at the low risk level to 100% in Equity at the highest risk level.

A.2.2 The Covariance Matrix output comparisons

Scenario 2

This scenario differs from Scenario 1 in that Weighted returns and the Observations Weighted Sample covariance matrix have been used for the optimization.

The Scenario 2 cases shown in this section, will demonstrate the combined effect of the observations weighted

covariance matrix and the weighted expected return inputs.

For the weighted returns further distinction is made regarding the type of weighting used. There are 4 different types of scenarios, discussed below, which will be compared to each other as well as to the base case, Scenario 1.

Scenario 2a makes use of a linear weighting methodology where the weights are decreased by the same proportion every month. The monthly weights were calculated using the following formula.

$$W_{t-1} = W_t - \frac{1}{36} \tag{A.1}$$

such that $W_{36} = 100\%$ and $t = 1, \dots, 36$

The factor used to decrease the preceding monthly weights by, $\frac{1}{36}$, was decided on arbitrarily as a result of the number for data points used, and any other value could be used in its place. This weighting factor results in the weight being at 50% in month 18 of 36.

Scenario 2b, Scenario 2c and Scenario 2d make use of the exponential smoothing methodology. This methodology was described earlier and the formula shown by Equation 3.52 is replicated below.

$$E[R_{t+1}] = \alpha * X_t + \alpha(1 - \alpha)X_{t-1} + \alpha(1 - \alpha)^2X_{t-2} + \alpha(1 - \alpha)^3X_{t-3} + \dots$$

Different values for α were used for each of the scenarios. For Scenario 2b an 'optimal' α , 3.35%, was calculated using the equity benchmark historical returns as a base for the forecast values. The solution for the 'optimal' α value was calculated using Microsoft Excel Solver so that the sum of the squared error terms was minimized. Similarly, the 'optimal' α , 11.38%, for Scenario 2c was calculated using the bond benchmark historical returns as a base. For Scenario 2c an arbitrary value for α , 40%, was chosen so that the weightings it generated were significantly different to the other two cases.

Figure A.2 shows the resulting optimal weighting solutions of the four Scenario 2 cases and the base case, Scenario 1. The efficient frontiers of each scenario have been plotted on the same graph for comparative purposes and the area plots of each scenario show the optimal weighting changes for each of the scenarios' respective risk levels.

All of the graphs in Figure A.2 have been drawn on exactly the same risk scale, so that their values may be directly comparable. The optimal weighting solutions of the various Scenario 2 cases and the base case, shown by the area graphs, can be compared for each risk value in the range. Also, the expected returns of the optimal weighting solutions of each scenario can be compared on the efficient frontier graph.

As an example, for a chosen risk value of 3%, the optimal allocations and expected returns of Scenario 1 and the Scenario 2 cases have been compared in Table A.4.

At 3% Risk	Equity	Bonds	Cash	Expected Return
Scenario 1	59%	41%	0%	1.20%
Scenario 2a	63%	37%	0%	1.50%
Scenario 2b	62%	38%	0%	1.00%
Scenario 2c	71%	29%	0%	1.90%
Scenario 2d	82%	18%	0%	3.00%

Table A.4: Comparison of Optimal Allocations and Expected Returns at 3% Risk level

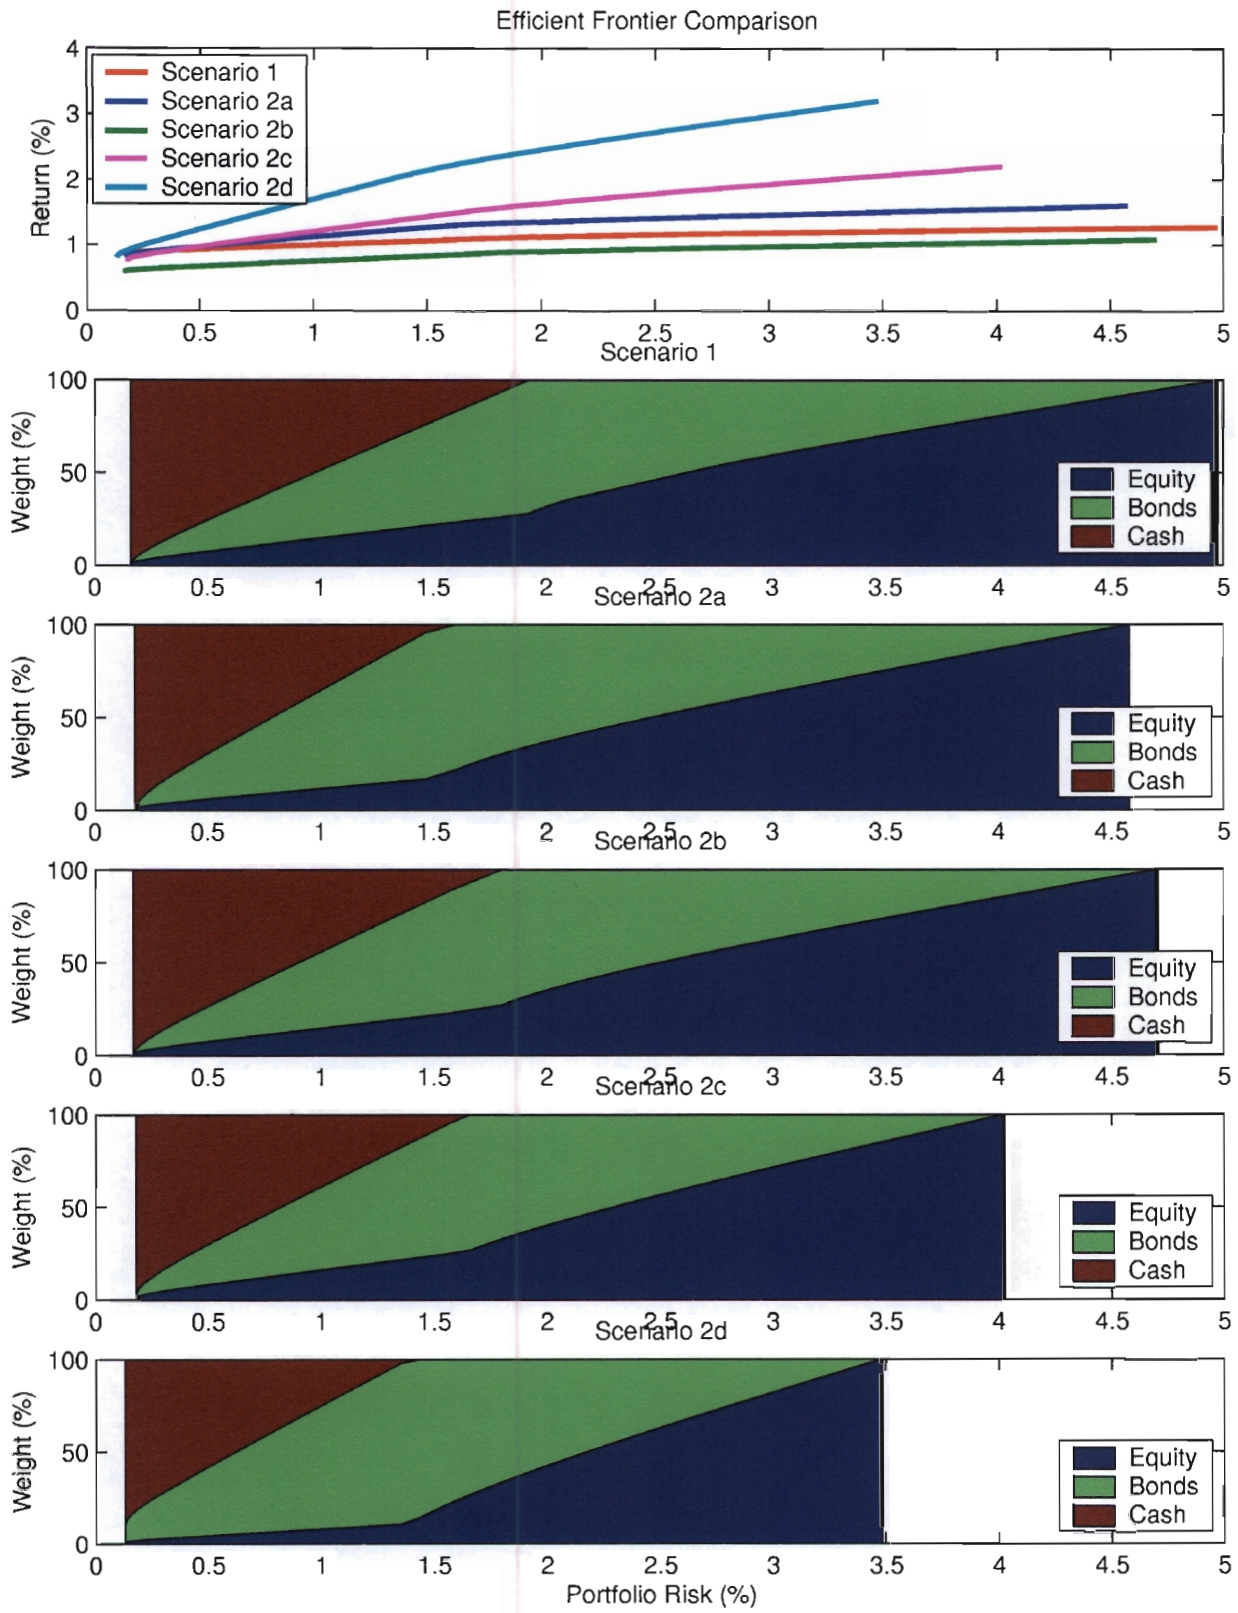


Figure A.2: Comparison of the Efficient Frontiers and Optimal Asset Allocation Solutions of the Scenario 2 cases

The observations weighted covariance matrix and weighted expected return estimators are supposed to be an improvement on the sample covariance matrix and standard average returns respectively. This is because they limit the impact older data has on the optimization results by applying smaller weights to the data. From the graphs in Figure A.2 it is immediately obvious that the Scenario 2 cases have a somewhat smaller risk range and varying expected return ranges to the base case, Scenario 1.

With the data set used, the more recent return history is less volatile than the older return history from earlier in the analysis period. This can be seen in Table A.5, which shows the actual historical annualised volatilities for each asset class benchmark for the different periods.

	ALSI50R	ALBI	AFMM
1year to Oct 04	<i>10.41%</i>	<i>4.34%</i>	<i>0.14%</i>
2years to Oct 04	<i>16.45%</i>	<i>4.77%</i>	<i>0.68%</i>
3years to Oct 04	<i>17.24%</i>	<i>7.04%</i>	<i>0.58%</i>

Table A.5: Annualised Risk (Volatility) of the asset class benchmarks

Table A.5 shows that for the equity and bond benchmarks, as the amount of historical data used was increased, the volatility of the returns increased, while for the cash benchmark there was an increase in volatility over the 1 year to 2 year period and then a decrease in the volatility from the 2 year to 3 year period. The 2 year volatility for cash captures the decreasing returns from when the money market was at the height of the interest rate cycle, where interest rates were at almost their highest rates, to the current low returns. This is as a result of the many interest rate decreases that we have experienced in South Africa over the last year and a half. The 3 year volatility for cash captures the upward trend in returns to the interest rate peak in 2002 and then equally the downward trend in returns to the current lows, thereby keeping the overall volatility over the period lower, due to the lower average return over the period.

The decrease in return volatility over the recent history would limit the range of risk possibilities for the weighted covariance matrices. On assessment of the differing weights used for Scenario 1 and the Scenario 2 cases, the results shown in Figure A.2 make more sense. Figure A.3 shows the weightings for each point in the analysis period, for Scenario 1 and the Scenario 2 cases.

The weights used in this comparison were normalized to ensure that the sum of each scenario's weights was 100%. Scenario 1 has the highest weighting for the initial data point, November '01, while Scenario 2d has the lowest. Scenario 1 has a constant weighting to each data point while the weights for Scenario 2d decrease quickly from the highest weighting of 40% in October '04, to practically a zero weight in November '01. Given the change in the benchmark return volatilities as shown in Table A.5, it is understandable that with Scenario 1's constant weighting throughout the period and its subsequent higher weighting to the older data points that it would have the highest risk range. On the other hand, Scenario 2d has the majority of its weighting concentrated in the most recent history, and with the benchmark's lowest return volatility over this period: it would thus have the lowest risk range.

Similar explanations can be given to show why each of the Scenario 2 cases' risk ranges are lower than that of Scenario 1. Also the Scenario 2 cases' weights may be compared to one another to explain how the various weighting options have affected their risk ranges, given the benchmark return volatilities.

The effects on the risk ranges of the weighted scenarios demonstrates that the weightings applied to the Scenario 2 cases has diminished the effect of the older data on the optimization results, as would have been expected.

Weightings Comparison




Figure A.3: Weightings for Scenario 1 and the Scenario 2 cases.

However, as a result of the combined effects of the covariance matrix estimator and the average return estimates, it is not entirely clear to what extent the weighting of the covariance matrix estimates has affected the optimal solutions as opposed to the weighting of just the return estimates. The total effect of weighted inputs relative to un-weighted inputs is clearly visible from the graphs in Figure A.2 though.

While the weighting of the data is beneficial in reducing the effect of possibly irrelevant data, as mentioned earlier when describing the estimator, a major downfall of the weighted estimator is that it allocates the same decay factor across all asset classes involved in the optimisation. This may be sub-optimal given that 3 different asset classes were used in each optimization and that it may be better to allocate different decay rates to each of them. For example, the effects of the different weightings that would be optimal for different asset classes can be seen when comparing Scenario 2b and Scenario 2c. Scenario 2b uses the 'optimal' for equity while Scenario 2c uses the 'optimal' for bonds, and from the results shown in Figure A.2, it can be seen that the optimizations provide very different optimal asset class allocations and risk and expected return ranges.

Scenario 4

In this scenario, the expected return estimator was kept the same as in the base case, i.e. average returns, while the covariance matrix estimator was changed to the Ledoit Covariance matrix. For the Ledoit covariance matrix, the market model (single factor model) was used as the biased estimator towards which the sample covariance matrix (unbiased estimator) would be shrunk, as per the Ledoit methodology.

This methodology aims to pull the out lying error estimates that may occur with a sample covariance matrix

towards more central values, thereby decreasing the error implicit in the covariance matrix estimation. Thus the Ledoit covariance matrix provides a better estimate, which lies somewhere between the two extremes, than would be obtained by using either the sample covariance matrix or some other more structured, biased estimator (in this case the single factor market model).

Figure A.4 shows the Efficient frontiers of Scenario 1 and Scenario 4 as well as the optimal weighting solutions of each scenario at their specific risk levels. These optimal solutions have again been represented in an area plot.

In this comparison, the use of the Sample covariance matrix and the Ledoit covariance matrix produce very similar optimal solutions. This can be seen in the area plots of the two scenarios and also from the Efficient Frontier comparison. The Efficient Frontier comparison shows that the risk and return spreads of each efficient frontier are very similar. The Efficient Frontier of Scenario 1 does however seem to lie just slightly above that of Scenario 4 over most of the graph.

This result may at first glance be very disconcerting given the argument that the Ledoit covariance matrix should be a superior estimate. However, when the specifics of these particular optimizations are looked into, the results become clearer.

The optimisations performed here, used 3 asset classes with 3 years of historical return data. One of the major drawbacks of using sample covariance matrix estimators is the dimensionality problem, as explained earlier. In order to provide a good estimate the sample covariance matrix requires $n * (n + 1) / 2$ parameters to be estimated. Where n is large, this becomes very problematic, but in this case $n = 3$ and thus only six parameters are required to be estimated. Also there is sufficient historical data available with which to perform this estimation.

In this case then, the estimation error which is generally problematic with the sample covariance matrix estimator is not a major drawback. Hence, the Ledoit covariance matrix estimator which aims primarily to reduce estimation error will provide very similar results to those of the sample covariance matrix, as was experienced here.

However, while the dimensionality problem may have been somewhat averted for the purposes of this study, it remains a very significant and real problem when dealing with sample covariance matrix estimators. Therefore, the Ledoit covariance matrix will remain the covariance matrix estimator of choice, given its ability to reduce estimation error and its resultant more stable variance and covariance estimates.

A.3 Expected Returns

The effect of the expected return estimates on the optimization was established by varying the type of estimates used, and then assessing the resulting optimal solutions that were generated. The average return estimate was used as a base case against which to assess the effect of some of the different estimates mentioned in the previous chapter.

Three of the scenarios, as outlined in Table A.1, are relevant to assessing the impact on the optimization by changing the expected return estimates. The relevant scenarios are Scenario 1, Scenario 2 and Scenario 3.

The inputs and results of the optimization for Scenario 1 have already been discussed in the covariance matrix section, and will therefore be omitted here.

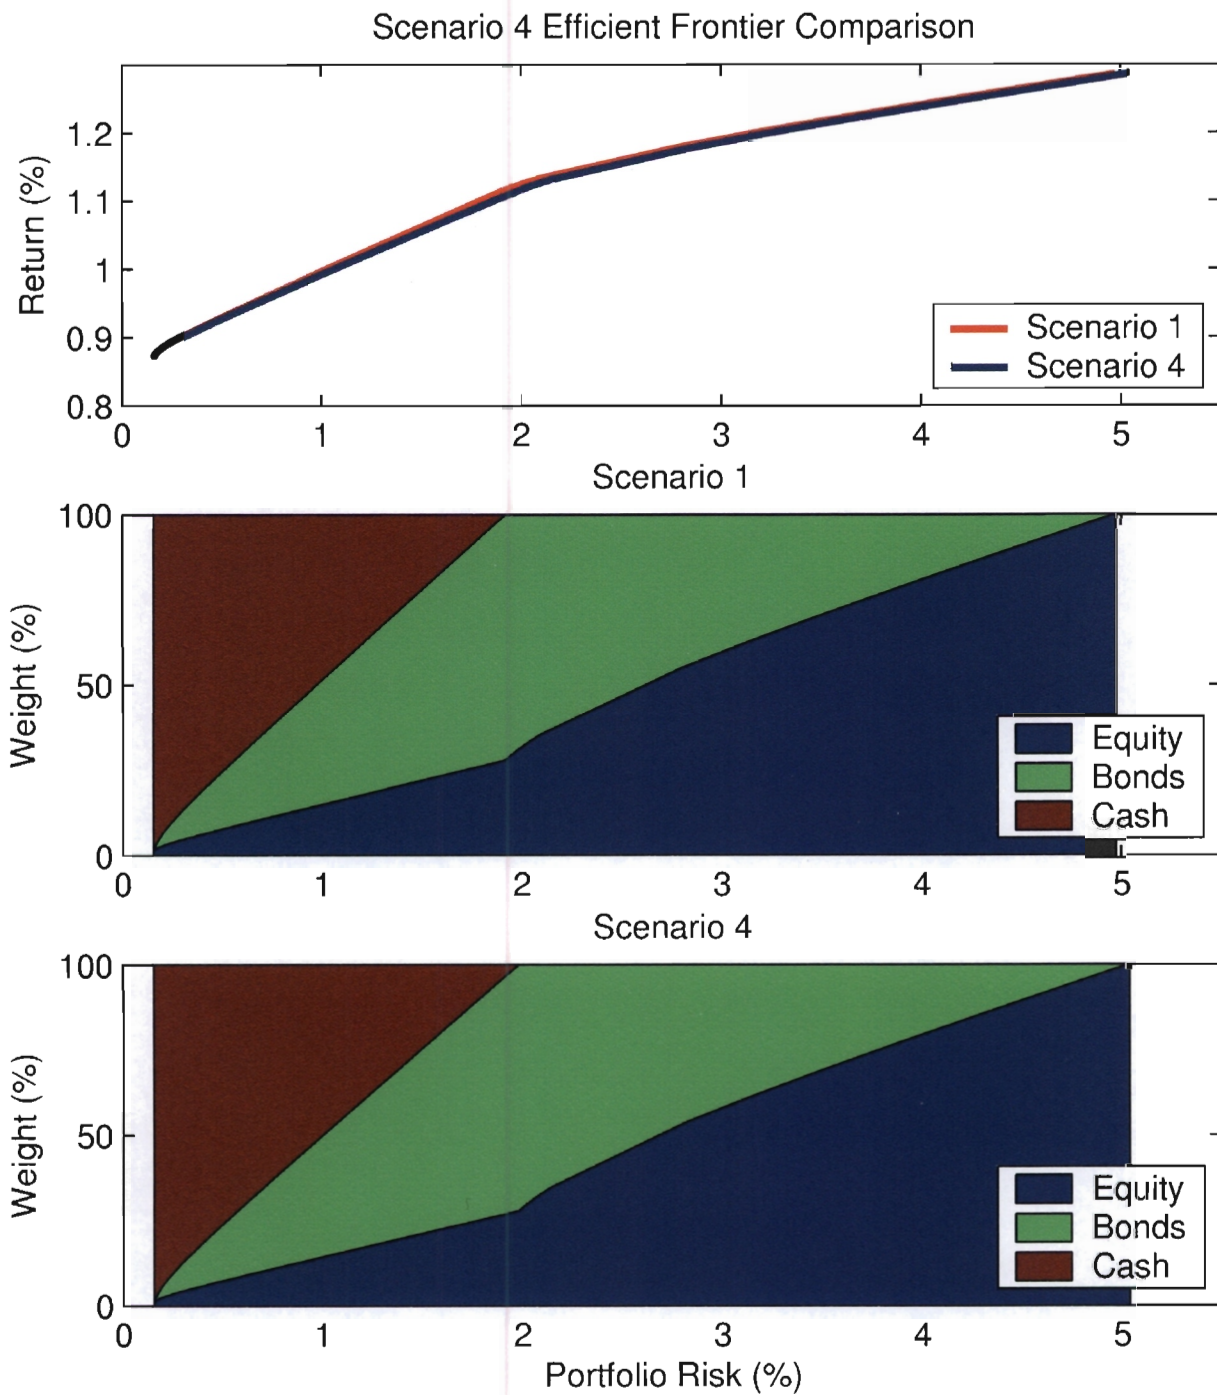


Figure A.4: Comparison of the Efficient Frontiers and Optimal Asset Allocations solutions for Scenario 4

A.3.1 The Average Return output comparisons

Scenario 2

As mentioned previously, Scenario 2 utilises weighted returns and the weighted sample covariance matrix as the expected return and covariance matrix estimators.

Some of the following discussion will build on the previous discussion about the Scenario 2 cases, but will focus more on the differences in the Expected return estimates. The weighted expected return estimates generated for each of the Scenario 2 cases as well as Scenario 1 are shown in Table A.6.

	ALSI50R	ALBI	AFMM
Scenario 1	<i>1.29%</i>	<i>1.05%</i>	<i>0.87%</i>
Scenario 2a	<i>1.61%</i>	<i>1.20%</i>	<i>0.83%</i>
Scenario 2b	<i>1.09%</i>	<i>0.80%</i>	<i>0.59%</i>
Scenario 2c	<i>2.21%</i>	<i>1.24%</i>	<i>0.74%</i>
Scenario 2d	<i>3.21%</i>	<i>1.90%</i>	<i>0.66%</i>

Table A.6: Comparison of the Expected Return estimates for the Scenario 2 cases

With a first look at the values here, no clear pattern can be determined for the weighted scenarios relative to the un-weighted scenario. Some expected return estimates are higher than those of Scenario 1 while some are lower. To obtain a clearer understanding of the reasons for the estimates that were obtained, one would need to consider the monthly changes in the underlying benchmarks' returns, relative to the weightings applied for Scenario 1 and the various Scenario 2 cases.

In order to do this, graphs of the monthly weightings for Scenario 1 and the Scenario 2 cases have been drawn relative to the monthly returns of each of the asset class benchmarks. The graphs are Figure A.5, Figure A.6 and Figure A.7, showing the weights relative to the Equity, Bonds and Cash benchmark returns respectively.

To make the comparison easier, a few specific points in the analysis period have been chosen so that the exact weightings of the Scenario 2 cases at those points can be compared. These are shown in Table A.7. The weights here will give us some better idea of the values shown in the following graphs.

	Oct 04	Oct 03	Oct 02	Nov 01
Scenario 1	<i>2.78%</i>	<i>2.78%</i>	<i>2.78%</i>	<i>2.78%</i>
Scenario 2a	<i>5.41%</i>	<i>3.60%</i>	<i>1.80%</i>	<i>0.15%</i>
Scenario 2b	<i>4.74%</i>	<i>3.15%</i>	<i>2.09%</i>	<i>1.44%</i>
Scenario 2c	<i>11.53%</i>	<i>2.70%</i>	<i>0.63%</i>	<i>0.17%</i>
Scenario 2d	<i>40.00%</i>	<i>0.07%</i>	<i>0.00%</i>	<i>0.00%</i>

Table A.7: Scenario 2 Weights at yearly end points of the analysis period

The weighted expected return estimated for Scenario 2d, for example, can be explained by looking at what its monthly weights were for every monthly period, relative to the Equity benchmark's return in that month. In October '04 alone, Scenario 2d has about a 3.5 times higher weighting to Equity than does Scenario 2c, which has the next highest weighting at that point. The majority of estimated expected return for Scenario 2d is determined by the four most recent data points, i.e. July '04 to Oct '04. Also by April '03 the weights

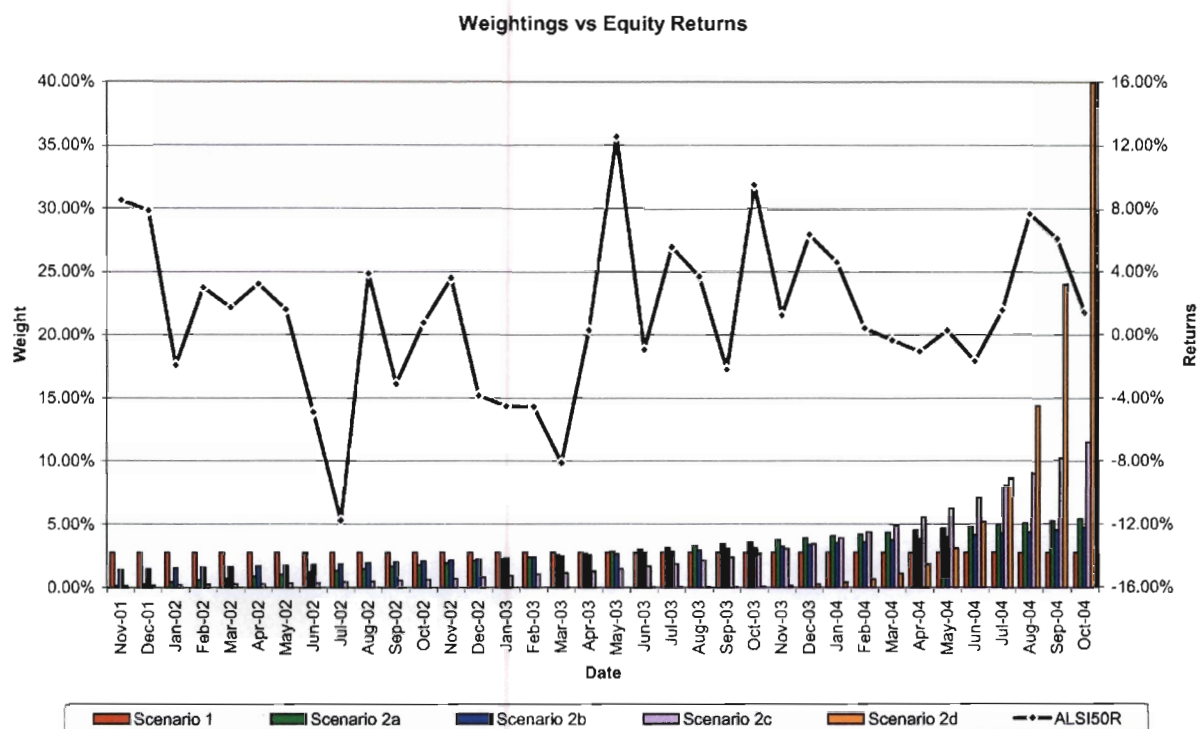


Figure A.5: Scenarios monthly weights relative to the Equity benchmark returns (ALSI50R)

for Scenario 2d are practically zero. This means that very little of the period over which equity had the worst returns is taken into account when calculating the expected return for Scenario 2d. As a result of the combined effects of these weightings, Scenario 2d has the highest expected Equity return estimate.

As another example, look at Scenario 1 relative to Scenario 2a. Scenario 1 has a constant equal weighting to all the observations in the analysis period, While Scenario 2a's weightings decreases by a constant factor. For the period November '01 to April '03, Scenario 1 has a higher weighting than Scenario 2a at every point, in May '03 they have more or less the same weight, and then from June '03 Scenario 2a has an increasingly higher weighting than Scenario 1. Over the November '01 to April '03 period, during which Equity had large negative returns, the higher Scenario 1 weights would thus incorporate more of the negative effects, while over the May '03 to October '04 period where Equity had mostly positive returns, Scenario 2a's higher weights would bias it towards the higher equity returns. It is therefore understandable that the expected return estimate for Scenario 2a is higher than that of Scenario 1 for Equity.

As a means of explaining the differences in the Bond expected return estimates, it was decided to look at Scenario 2c relative to Scenario 1. Looking at the monthly weighting differences again, it is noticed that Scenario 1 has a higher weighting than Scenario 2c for the period November '01 to October '03. Over this period Bonds have a short period of significantly negative returns, while for most of the rest of the analysis period they have positive returns. This extremely negative period from December '02 to about March '02 will affect the Scenario 1 estimate more than it would the Scenario 2c estimate, since Scenario 2c's weighting was close to zero by then. From April '02 to October '03, the higher weighting in Scenario 1 would influence its estimate more positively than it would Scenario 2c. Scenario 2c's estimate would be more influenced by the returns from November '03 to October '04, since its weightings were much higher

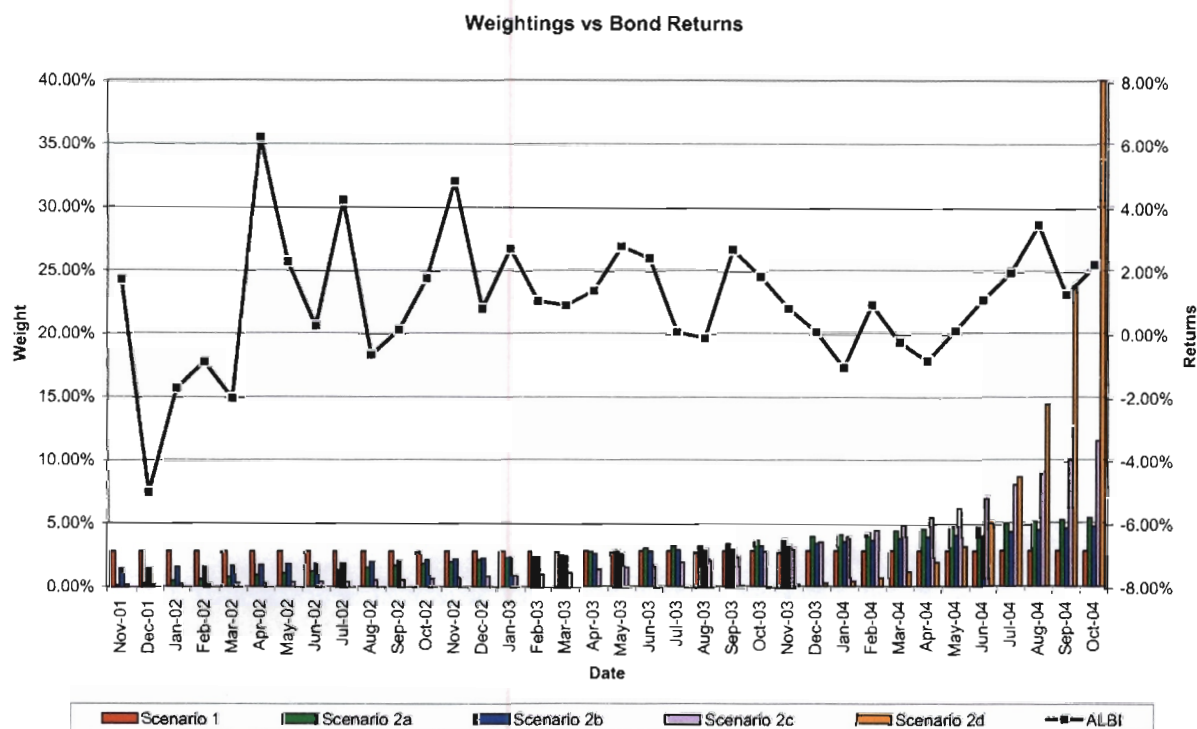


Figure A.6: Scenarios monthly weights relative to the Bond benchmark returns (ALBI)

in a time when bonds had mostly positive returns.

Scenario 2c's expected bond return is higher than Scenario 1's because its weightings have excluded it from experiencing the negative returns to the extent that Scenario 1 would have. Also even though Scenario 1 had a consistent weighting to the positive returns than Scenario 2 did, the significantly higher weightings that Scenario 2c had to the most recent positive observations far out-weighted Scenario 1's consistent weighting which was also dragged down by the negative returns earlier on in the period. Another key reason for this result would be that Scenario 2c's weightings were specifically chosen so that its expected return estimate was fitted best to the past returns. Since over most of the history bonds had positive returns, its expected return estimate would be more optimistic than Scenario 1's which gave more weight to the negative returns.

Using similar arguments to those for the previous two asset class comparisons, one can make sense of the cash expected return estimates generated for the different Scenario 2 cases. Scenario 2d's weights start at 40% (in October '04) and decrease exponentially to close to 0% (in October '03). Over this period the Cash returns, while always positive, were at lower levels than previously. This is as a result of the structurally different macro-economic environment in which South Africa currently finds itself. Interest rates have been steadily decreasing for this entire period due to the Reserve Bank's inflation targeting. Scenario 1, however, with its equal weighting of the entire history, would have a bias to higher returns due to the returns being at higher levels prior to October '03. This would result in Scenario 1 having the higher expected cash return, The Scenario 2 cases discussed so far have used weighted estimators for both the expected returns and the covariance matrix. This however does not allow one to assess the individual effects of each weighted input estimator. It was thus decided to run another Scenario 2 case, Scenario 2e, where just the expected return




Figure A.7: Scenarios monthly weights relative to the Cash benchmark returns (AFMM)

estimate has been weighted, while the un-weighted sample covariance matrix was used.

For Scenario 2e, the linear weighting methodology of Scenario 2a was used. The Expected return estimate for Scenario 2a and Scenario 2e are thus exactly the same. The sample covariance matrix was used for its covariance matrix estimates and they would therefore be the same as those of Scenario 1. Scenario 2e will thus allow us to assess the effect of just the weighted expected return estimates on the optimization.

For the purposes of the analysis, Scenario 2e will be compared against Scenario 1 to show the effect the weighted expected return estimates have on the optimization, and it will also be compared against Scenario 2a so that the effect of the weighted covariance matrix estimates will be observable. The differing asset class allocations and the efficient frontiers of the above mentioned cases can be seen in Figure A.8.

When assessing the results it was noticed that the risk ranges of Scenario 2e and Scenario 1 were exactly the same, while the expected returns for Scenario 2e were, for the most part, higher than those of Scenario 1. This shows that using weighted returns to generate the expected return estimates has no impact on the optimization's optimal solutions risk ranges, but rather has a significant impact on the optimization's expected return outcomes.

When comparing Scenario 2e to Scenario 2a, it was immediately obvious that the only major difference between the two was that Scenario 2a had a smaller risk range than Scenario 2e. It can thus be concluded that the use of weighted returns to generate the covariance matrix estimates has a direct impact on the optimizations optimal solutions risk ranges. In this case, as a result of the lower volatility in the more recent history, as discussed previously, the risk range for the weighted covariance matrix case, i.e. Scenario 2a, resulted in a smaller risk range than the un-weighted covariance matrix case, i.e. Scenario 2e. This

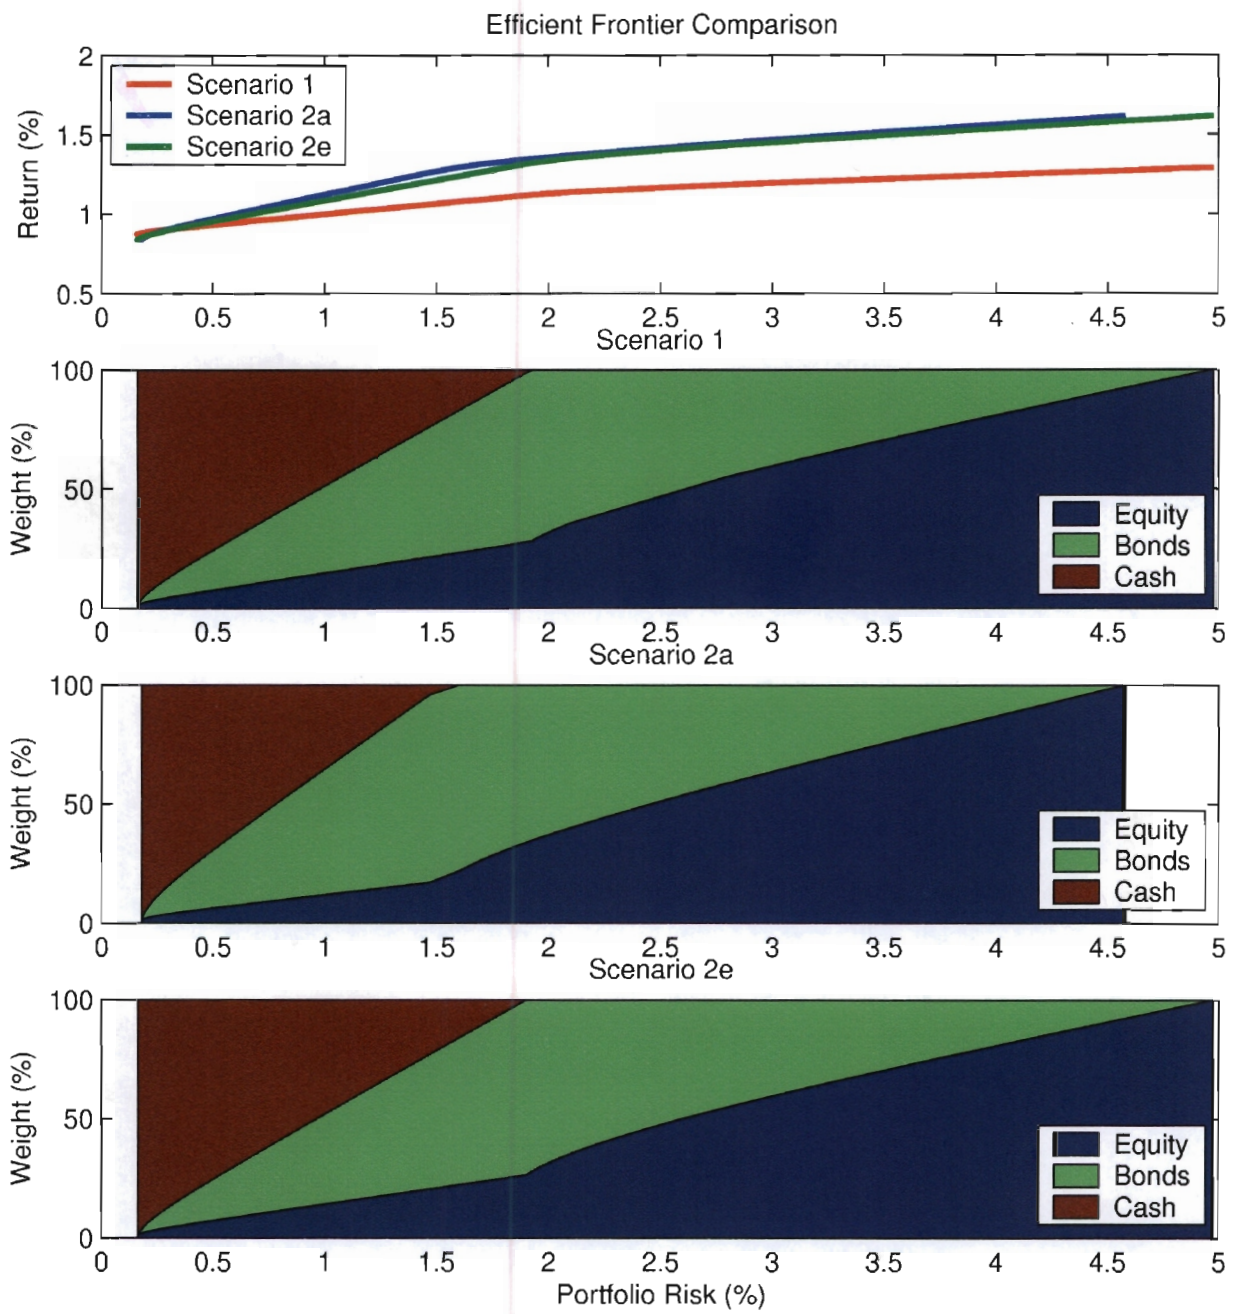


Figure A.8: Comparison of the Efficient Frontiers and Optimal Asset Allocations Solutions for Scenario 2e

has also had an effect on the asset class allocations of the two cases. For example, at a risk level of 1.5%, Scenario 2a has a significantly lower cash allocation and a higher bond allocation than Scenario 2e.

Also, apart from the noticeable difference in the risk ranges of Scenario 2e and Scenario 2a, there were only very slight differences with respect to the expected return estimates. The efficient frontier for Scenario 2a with both inputs weighted lay almost on top of or slightly above Scenario 2e's efficient frontier, showing further that the weighted covariance matrix has a very slight effect on the expected return outcomes.

This observation shows us that the mean variance optimization is much more sensitive to changes in the expected return estimates than to changes in the covariance matrix estimates.

Scenario 3

Scenario 3 used the sample covariance matrix as the covariance matrix estimator, and the expected returns have been estimated using average returns with a Black Litterman adjustment. The comparison between Scenario 1 and Scenario 3 will allow for an assessment of the effect of a Black-Litterman adjustment on the optimization results.

As discussed previously, the Black-Litterman adjustment (B-L adjustment) allows for a subjective overlay to the expected return estimates. While the Black-Litterman model is normally used to adjust the expected returns of assets, for the purposes of this investigation we have applied the Black-Litterman model's principles to the adjustment of the expected returns of asset classes. Here, Scenario 3a and Scenario 3b are cases that have used different Black-Litterman adjustments to illustrate its effect on the optimization. The views represent monthly out-performance numbers, since the entire optimization is based on monthly numbers

Scenario 3a This case takes the following 3 views into account to be used as the subjective overlays. The views are based on the principles of the long run risk-return tradeoff, i.e. that equity should outperform bonds and cash given that it is seen as the riskier asset class relative to the other two, and also bonds should outperform cash, given that it is riskier than cash. The return views are additive from the worst performing asset class to the best performing asset class, so that they represent a risk premium of sorts from one asset class to the next best performing asset class, i.e. Equity Return ζ Bond Return ζ Cash Return.

	View	Confidence
Equity outperforms Bonds	<i>0.40%</i>	<i>60.00%</i>
Equity outperforms Cash	<i>0.70%</i>	<i>70.00%</i>
Bonds outperform Cash	<i>0.30%</i>	<i>60.00%</i>

Table A.8: Black-Litterman Views for Scenario 3a

Scenario 3b In this case the return views have been altered in keeping with the outperforming asset classes. For Scenario 3b it was assumed that bonds were the best performing asset class relative to the other two, as was witnessed in the markets 2 years ago. Also, it has been assumed that cash out-performs equity.

These views have resulted in the following expected return estimates for the respective optimizations, as shown in Table A.10.

On comparison with the expected return estimates for Scenario 1, it can be seen that Scenario 3a's views has resulted in an increased expected return for equity while the other asset class returns remain almost

	View	Confidence
Bonds outperform Equity	<i>0.70%</i>	<i>60.00%</i>
Cash outperforms Equity	<i>0.40%</i>	<i>70.00%</i>
Bonds outperform Cash	<i>0.30%</i>	<i>60.00%</i>

Table A.9: Black-Litterman Views for Scenario 3b

	ALSI50R	ALBI	AFMM
Scenario 1	<i>1.29%</i>	<i>1.05%</i>	<i>0.87%</i>
Scenario 3a	<i>1.39%</i>	<i>1.05%</i>	<i>0.87%</i>
Scenario 3b	<i>0.94%</i>	<i>1.10%</i>	<i>0.87%</i>

Table A.10: Expected Returns estimates for Scenario 3, adjusted by B-L Views

unchanged. A possible reason for the unchanged estimates is that the views on bonds and cash were already in keeping with the view already expressed by the original data, and thus a further change was not made.

When comparing Scenario 3b's estimates relative to Scenario 1's, it was noticed that Scenario 3b has a very much decreased equity expected return and an increased bond expected return, while the cash expected return was mostly unchanged. The unchanged cash expected return estimate may be as a result of the Black-Litterman adjustment being made based on the Confidence level of the view and the asset classes volatility. Although the cash view expressed was different to that reflected by the original data, the volatility of cash was far less than the other two asset classes, and its required adjustment was made to the equity expected return due to equity's higher volatility.

The optimal asset class allocations and the efficient frontiers for the Scenario 3 cases and Scenario 1 are shown in Figure A.9.

The curve of the Scenario 3a solutions is steeper than that of the Scenario 1 curve. This tells us that Scenario 3a's optimal asset class allocation solutions provide a higher expected return for every level of risk. The additional optimism that was placed on the expected equity returns by the Black-Litterman adjustment, resulted in a higher equity return estimate being used in the optimisation. As mentioned previously, the optimization is sensitive to changes in the expected return estimates and thus we see a significant change to the efficient frontier.

An increase in the expected return estimates from the Black-Litterman adjustment, for a specific level of risk, would imply a lower risk tolerance level. It was thus noticed that the Scenario 3a optimal solution of 100% equity allocation was achieved at the same risk level but at a lower risk tolerance level (τ) than the Scenario 1 solution.

Scenario 3b, on the other hand, assumed that bonds were the out-performing asset class. From the results it was noticed that the efficient frontier for Scenario 3b has a gentler slope than that of Scenario 1 and a much narrower risk range. This implies that for every risk level the Scenario 3b optimal allocation solutions had a lower expected return than the Scenario 1 solutions.

From the relevant area graphs in Figure A.9 it can be seen that Scenario 3b has 100% allocation to bonds at its highest risk level, while Scenario 1 has 100% allocation to equity at its highest risk level. The resulting narrower risk range for Scenario 3b is then plausible if one considers that, in general, bonds are considered to be the less risky asset class relative to equity.

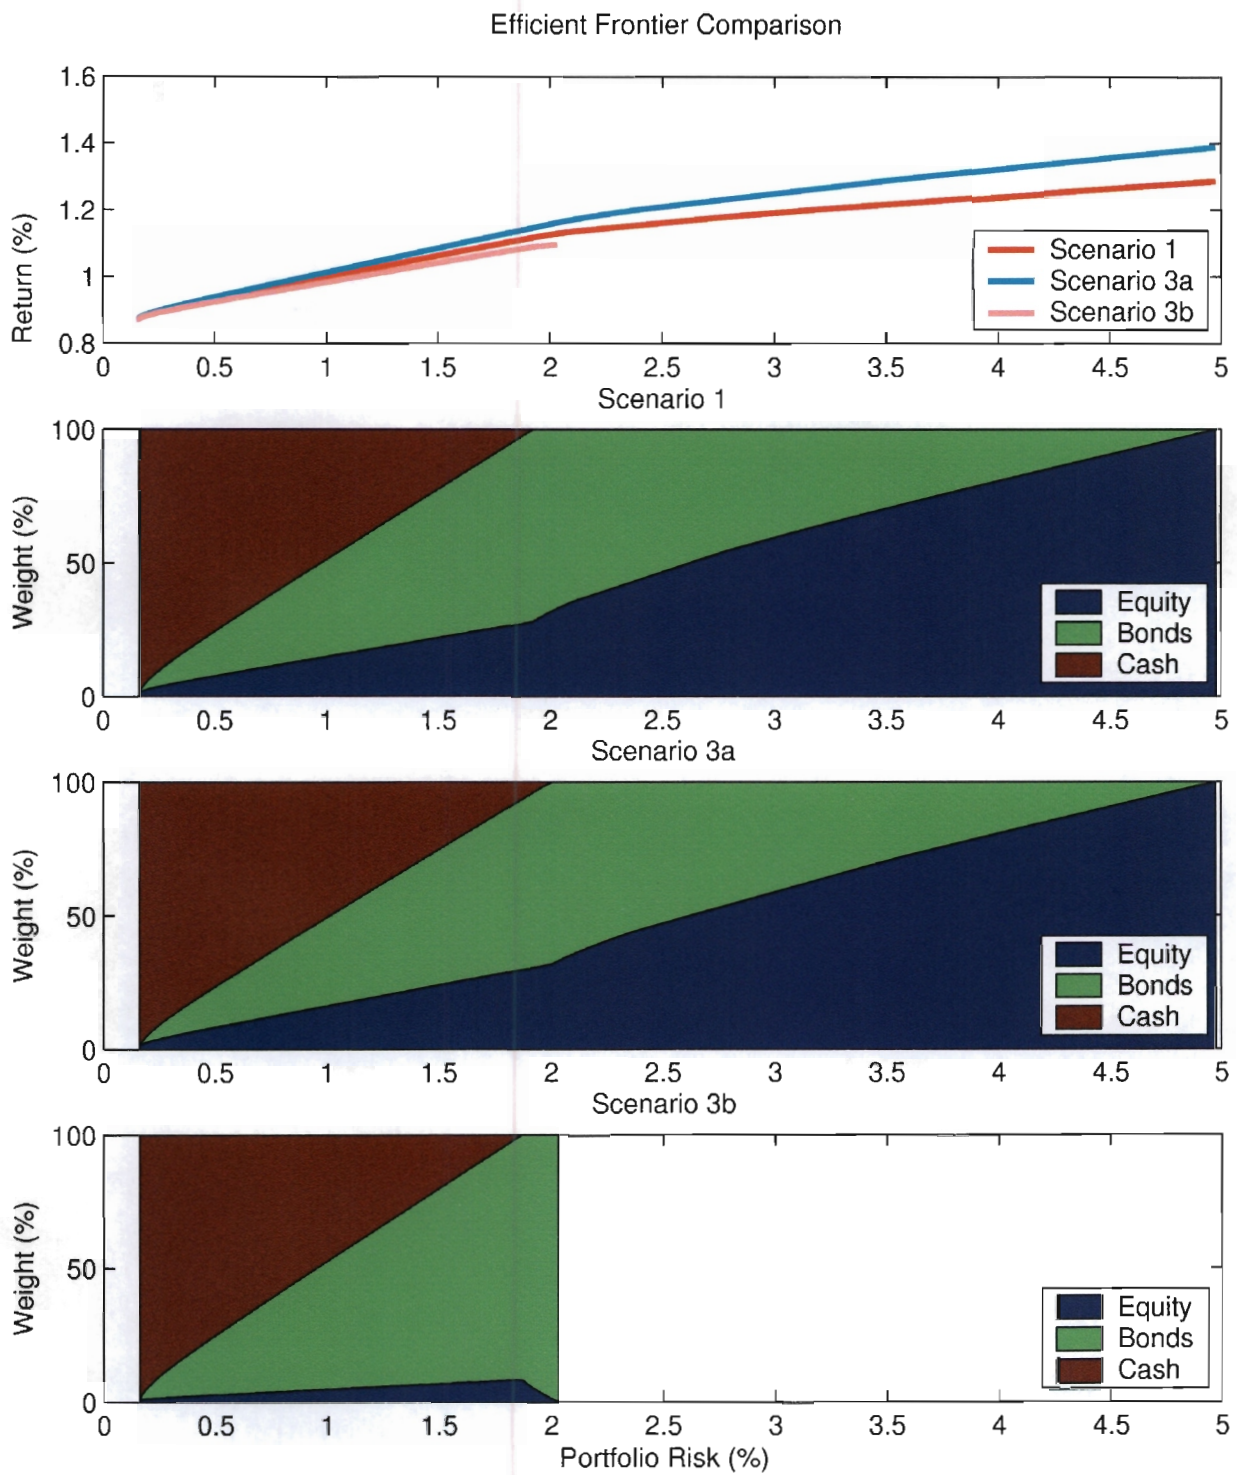


Figure A.9: Comparison of the Efficient Frontiers and Optimal Asset Allocations Solutions for Scenario 3

The area graphs also show that for every risk value common to both Scenario 3b and Scenario 1, Scenario 3b has a lower equity allocation and thus a higher combined bond and cash allocation relative to Scenario 1. Considering the risk-reward tradeoff principle, and the fact that bonds and cash are considered to be the less risky asset classes, it stands to reason that their expected return estimates should be lower than those of the higher risk equity asset class. Given Scenario 3b's higher combined allocation to bonds and cash at every risk level it is understandable that this Scenario 3b would have lower expected return estimates than Scenario 1.

This analysis shows that the Black-Litterman adjustment is effective in changing the expected return estimates as a result of a subjective overlay, and in so doing affecting the optimization results. The results obtained are thus skewed in accordance with the subjective overlay provided.

Bibliography

- [1] Averaging and exponential smoothing models. <http://www.duke.edu/~rmau/411avg.htm>. visited 7 June 2005.
- [2] Modern portfolio theory. http://en.wikipedia.org/wiki/Modern_portfolio_theory. visited 22 July 2005.
- [3] PCA-principal component analysis. <http://www.eng.man.ac.uk/mech/merg/research/datafusion.org.uk/pca.html>. visited 26 May 2005.
- [4] Carol Alexander, editor. *Risk Management and Analysis, Volume 1. Measuring and Modelling Financial Risk*. John Wiley & Sons, 1998.
- [5] C Bengtsson and J Holst. On portfolio selection: Improved covariance matrix estimation for Swedish asset returns. Master's thesis. Centre for Mathematical Sciences, Lund Institute of Technology, Lund University, 2002.
- [6] Z Bodie, A Kane, and Alan J Markus. *Essentials of Investments*. Irwin, 2nd edition edition, 1995.
- [7] Dave Bradfield and Jason Swartz. Demonstration of the Black-Litterman portfolio design procedure in the south african setting. Quantitative Research, Cadiz, October 2002.
- [8] G Connor. The three types of factor models: A comparison of their explanatory power. *Financial Analysts Journal*, 51:42–45, 1995.
- [9] Richard C Grinhold and Ronald N Kahn. *Active Portfolio Management - A Quantitative Approach for Providing Superior Returns and Controlling Risk*. McGraw - Hill, 2nd edition edition, 1999.
- [10] Guangliang He and Robert Litterman. The intuition behind the Black-Litterman model portfolios. Investment Management Research, Goldman, Sachs & Company, December 1999.
- [11] Thomas Idzorek. A step-by-step guide to the Black-Litterman model. http://faculty.fuqua.duke.edu/~charvey/Teaching/BA453.2003/How_to_do_Black_Litterman.doc. visited 21 September 2005.
- [12] Olivier Ledoit. Improved covariance matrix estimation. http://www.ledoit.net/ole2_abstract.htm. visited 26 May 2005.
- [13] Olivier Ledoit and Michael Wolf. Improved estimation of the covariance matrix of stock returns, with an application to portfolio selection. *Journal of Empirical Finance*, 10(5):603–621, 2001. http://www.ledoit.net/ole2_abstract.htm.
- [14] Olivier Ledoit and Michael Wolf. Honey, I shrunk the sample covariance matrix. <http://www.econ.upf.es/~wolf/honey.pdf>, 2003. visited 26 May 2005.

- [15] Robert Litterman and Quantitative Resources Group Goldman Sachs Asset Management. *Modern Investment Management - An Equilibrium Approach*. John Wiley & Sons Inc. Hoboken, New Jersey, 2003.
- [16] A Madhavan and J Yang. Practical risk analysis for portfolio managers and traders. *ITG whitepapers*, 2003. <http://www.itginc.com/research/whitepapers/madhavan/RISKJPM2.pdf>.
- [17] H Markowitz. Portfolio selection. *Journal of Finance*, 7:77-91, 1952.
- [18] Robert D Mason and Douglas A Lind. *Statistical Techniques in Business and Economics*. Irwin McGraw-Hill, 9th edition edition, 1990.
- [19] Harry H Panjer, editor. *Financial Economics - With Applications to Investments, Insurance and Pensions*. The Actuarial Foundation, Schaumburg, Illinois, 1998.
- [20] L Rahl. *Risk Budgeting - an Approach to Investing*. Risk Books, 2000.
- [21] R Tibben-Lembke. Forecasting introduction. <http://equinox.scs.wmr.edu/homepage/rtl/ForecastingIntro17.pdf>, 2004. visited 7 June 2005.
- [22] B Waring, D Whitney, J Pirone, and C Castillo. Optimizing manager structure and budgeting manager risk. *Journal of Portfolio Management*, 26:90-104, 2000.