

Aalto University
School of Science
Degree Programme in Engineering Physics and Mathematics

Henri Tuovila

Optimised Strategies for Dynamic Asset Allocation

Master's Thesis
Helsinki, May 1, 2016

Supervisor: Professor Ahti Salo
Advisor: Jari Mikkonen M.Sc. (Math.)

The document can be stored and made available to the public on the open internet pages of Aalto University. All other rights are reserved.

Author:	Henri Tuovila	
Title:	Optimised Strategies for Dynamic Asset Allocation	
Date:	May 1, 2016	Pages: vi + 80
Major:	Systems and operations research	Code: Mat-2
Supervisor:	Professor Ahti Salo	
Advisor:	Jari Mikkonen M.Sc. (Math.)	
<p>Modern portfolio theory is a widely used framework in the financial industry. It has a solid theoretical background, and has been successfully employed by the practitioners for decades. Traditional models based on Harry Markowitz's portfolio theory, and its further improved versions, have one significant shortcoming: they are single-period models by definition, and are not able to accommodate multi-period considerations.</p> <p>In this thesis, instead of modern portfolio theory and mean-variance optimisation, we use stochastic programming. To employ stochastic programming as a technique to find the optimal allocations, we need to develop scenarios, or scenario trees that describe the stochastic variables and their distributions.</p> <p>To generate the scenarios, we employ a methodology called moment matching, where the relevant properties of stochastic variables in our generated scenarios are fitted to counterparts estimated by means of time series analysis and econometric modelling. These stochastic factors are also called market invariants in this context. Market invariants are then translated into asset returns, which make it possible to find optimal asset allocations in each stage of the scenario tree.</p> <p>An illustrative asset allocation example is presented in this thesis to demonstrate how the dynamic allocation strategy performs compared to a fixed allocation decision. The results are rather intuitive, and as expected, the dynamic allocation strategy outperforms the fixed strategy in the scenarios generated. A comparison to traditional mean-variance framework is conducted, and it is seen that the resulting allocations for both dynamic and fixed strategy are close to being mean-variance efficient.</p> <p>Further research topics include changing the scenario generation methodology, and more sophisticated modelling of interest bearing instruments. An interesting direction for further development would be constructing the entire term structure of a yield curve, which would allow flexible valuation of assets and liabilities based on their present values.</p>		
Keywords:	stochastic programming, scenario tree, scenario generation, asset allocation, portfolio optimisation, short-rate model	
Language:	English	

Tekijä:	Henri Tuovila		
Työn nimi:	Optimoidut strategiat varojen dynaamiseen kohdentamiseen		
Päiväys:	1. toukokuuta 2016	Sivumäärä:	vi + 80
Pääaine:	Systeemi- ja operaatiotutkimus	Koodi:	Mat-2
Valvoja:	Professori Ahti Salo		
Ohjaaja:	Jari Mikkonen, FM		
<p>Moderni portfolioteoria on rahoitusallalla yleisesti käytetty. Sillä on vahva teoreettinen pohja, ja sen sovelluksia on käytetty onnistuneesti vuosikymmenien ajan. Harry Markowitzin kehittämän portfolioteorian, ja siitä kehitettyjen parannettujen versioiden yksi ilmeinen heikkous on kuitenkin se, että ne ovat rakenteeltaan yksiperiodisia malleja. Ne eivät näin ollen sovellu moniperiodiseen tarkasteluun.</p> <p>Tässä diplomityössä portfolioteorian perinteisten mallien sijaan sovelletaan stokastista ohjelmointia optimaalisten omaisuuslajiallokaatioiden löytämiseen. Jotta stokastista ohjelmointia voisi hyödyntää, on ensin kehitettävä skenaariot, jotka kuvaavat satunnaismuuttujat ja niiden jakaumat, joiden perusteella optimointi voidaan tehdä.</p> <p>Skenaarioiden luomiseksi käytämme momenttien sovittamiseksi kutsuttua menetelmää, jossa ongelman kannalta relevantit satunnaismuuttujien ominaisuudet sovitetaan skenaarioissa aikasarja-analyysin ja muiden ekonometristen menetelmien avulla estimoituihin vastineisiin. Stokastisia muuttujia kutsutaan tässä yhteydessä markkinainvarianteiksi, ja ne voidaan muuntaa omaisuuslajien tuotoiksi, joiden perusteella voidaan laskea optimaalinen omaisuuslajiallokaatio skenaariopuun jokaisessa haarassa.</p> <p>Työssä esitellään havainnollistava esimerkki dynaamisen allokaatiostrategian ja kiinteän allokaatiostrategian vertailua varten. Tulokset ovat intuitiivisia ja kuten odotettua, dynaaminen strategia pärjää kiinteää paremmin. Tehty vertailu perinteiseen Markowitzin mallin mukaiseen optimointiin osoitti, että sekä dynaaminen että kiinteä stokastisen optimoinnin strategia ovat lähellä Markowitzin mallin mukaista tehokasta rintamaa.</p> <p>Jatkotutkimuskohteita ovat skenaarioiden generointiin käytetyt menetelmät ja korkoperustaisten sijoituslajien tarkempi mallintaminen. Kiinnostava tutkimussuunta olisi koko korkokäyrän mallintaminen, joka mahdollistaisi mielivaltaisten tase-erien markkina-arvostamisen nykyarvoonsa.</p>			
Asiasanat:	stokastinen ohjelmointi, skenaariopuu, skenaarioiden generointi, omaisuuslajiallokaatio, portfolion optimointi, korkomalli		
Kieli:	Englanti		

Acknowledgements

This master's thesis was written as a by-product of development of a multi-stage financial planning tool for OP Financial Group's Wealth Management department. The work was instructed by a colleague Jari Mikkonen, to whom I am very grateful for sharing his expertise, and sacrificing his time and effort. Regular follow-ups and coaching sessions with Jari were the key to push the sometimes so sluggishly advancing project until the very end.

I am thankful for being supervised by professor Ahti Salo, who provided valuable feedback along the way, even at short notice, and with flexible means.

Lastly I wish to thank everyone who participated in the process of proofreading, and providing encouragement. In addition, many thanks to Unisport for providing meaningful activities while procrastinating.

Helsinki, May 1, 2016

Henri Tuovila

Contents

1	Introduction	1
1.1	Background and motivation	1
1.2	Context and research objectives	2
1.3	Structure of the thesis	4
2	Background and literature review	5
2.1	Modern portfolio theory	5
2.2	Multi-period investing	7
2.3	Asset allocation strategies	7
2.4	Utility functions	10
2.5	Scenarios for stochastic programming	14
2.5.1	Scenario generation	14
2.5.2	Clustering	15
2.5.3	Sequential simulation with k -means clustering	16
2.5.4	Moment matching	16
2.5.5	Property matching	17
2.5.6	Integration quadratures	17
3	Theoretical model description	18
3.1	Stochastic programming models	18
3.1.1	Portfolio optimisation by stochastic programming	19
3.2	Scenarios for stochastic programming	19
3.2.1	Scenario trees	20
3.2.2	Moment matching	24
3.3	Interest rate models	26
3.4	Objective function	28
4	Implementation and numerical results	31
4.1	Data for model calibration	31
4.1.1	Market invariants	33
4.1.2	Money markets	33

4.1.3	Government bonds and bond rate	35
4.1.4	Investment grade bonds and credit spread	37
4.1.5	Property index	38
4.1.6	Matching moments	39
4.1.7	Annualising moments	39
4.1.8	Parameters for short-rate model	40
4.2	Generated scenario trees	43
4.3	Optimised allocations	44
4.3.1	Dynamic second stage decisions	50
4.3.2	Comparison with fixed-mix strategy	52
4.3.3	Comparison with mean-variance optimisation	53
5	Evaluation and discussion	56
5.1	Stability of the optimal solution	56
5.2	Potential issues in model specification	58
5.3	Solver selection and computational issues	60
5.4	Further research topics	61
6	Conclusions	63
A	Data characteristics	70
B	Estimated short-rate model parameters	74
C	Optimal allocations	76
D	Parameters for mean-variance optimisation	80

Chapter 1

Introduction

1.1 Background and motivation

Defining the asset allocation between major asset classes in an investment portfolio is the biggest decision investors make. At the highest level this means choosing the proportions in which to invest in interest bearing investments, and in equities. This high level allocation decision is the strategic asset allocation, and it determines for most part the expected rate of return and risk in the portfolio. (Zenios and Ziemba, 2007)

Portfolio analysis and portfolio optimisation are well known and widely applied topics. The models in this context are often referred to as the modern portfolio theory, and they are an important tool in today's financial industry. However, these models have a shortcoming: they are by definition single-period models, and are optimal in a multi-period setting only under certain assumptions. These single-period models do not explicitly allow adjusting decisions along the planning horizon when new information becomes available.

Anticipating the development of the state of the economy and an investment portfolio in a longer horizon is in the interest of all investors, and especially such as insurance companies, pension funds and other institutions facing the asset and liability management problematics.

This thesis explores the models for analysing and optimising an investment portfolio in a multi-period setting exploiting scenario based stochastic optimisation methods.

In order to develop a reasonable model for stochastic optimisation problem to be solved, various modelling issues have to be encountered. These topics include detailed return modelling of fixed income investments, whose distributions are thought to differ significantly from the usual normality as-

sumption.

Modelling the market invariants properly finally leads to scenario generation, which is a broad field with various different models to choose from. Scenario generation leads to scenario trees, which can finally be fed to an optimiser. The results and resulting decisions will be only as good as are the scenarios used, hence the importance of setting them carefully.

One of the key elements of the decision making process is defining the investor's preferences, which is not at all trivial. Every investor has their individual attitude towards risk, and many investors measure risk in different ways. For institutional investors such as insurance companies the risk might be measured relative to a regulatory solvency capital requirement. On the other hand, private individuals may be willing to avoid big changes in the value of their portfolio.

Similar scenario based optimisation models are in use in other than financial industries as well, because they provide a rigid way to account for uncertainty in a multi-period setting. These models are not included here, but one example of an application from the field of chemical engineering can be found in Calfa et al. (2014).

1.2 Context and research objectives

In this thesis we study the problem from a viewpoint of an institution with assets and a return target. This could be seen as an arbitrary target set by managers of the company, or an external target set by the liabilities the institution has. Institutions like this include pension funds and life insurers. The approach will not be more specific, however this thesis is written having a pension insurance company in mind. The needs such a company may have are taken as a motivational starting point.

The approach we intend to use involves setting an annual return target, which could be seen set by the liabilities the institution has. The liabilities can have an estimated annual growth rate, which needs to be covered by the asset side of the balance sheet. A typical growth rate for liabilities would be the inflation, which the central banks typically try to maintain at about 2%. (European Central Bank, 2016) In a real asset and liability management problem (shorthand: ALM) the liabilities should be modelled separately, which would result in interesting objectives for the optimisation, like remaining solvent, and holding a buffer to ensure the solvency of the institution.

Some investors face an optimisation problem, where the absolute returns of an asset portfolio matter only little, but the main point of interest is

Table 1.1: An example of the contents of a simplified balance sheet.

Assets	Liabilities
Cash	Cash flow stream
Government bonds	
Corporate bonds	
Equity	
Real estate	

actually the net wealth. In this case the liabilities have to be taken into account as well.

Liabilities are typically modelled as cash flows in time. They are typically independent of the assets, except from the interest rates. Measuring the market values of a cash flow stream is relatively straight forward given the yield curve for discounting, which enables us to deduce the market values of liabilities in each stage of the scenario tree. An illustrative example of a very simple balance sheet is presented in table 1.1.

Having the simplified balance sheet defined we can further define the optimised variables in an asset and liability management terms.

Depending on the modelled company, liabilities can be market valued, which would make it compliant with the regulation implemented in Europe. Alternatively in a more simple case, fixed forecasts could be given by an actuary used as such, and the goal would be simply to remain solvent.

For the model to be entirely compliant with the legislative requirements, it needs to take into account the additional solvency capital requirement, namely the amount by which the assets need to exceed the liabilities in value.

Extending the model to comply with more complex regulatory requirements, but also for more complex estimates for the liabilities would provide reasonable areas for further research. Similar models have been employed earlier for both pension funds and life insurance companies. These earlier models would need, however, updating to be still valid today due to a very different economic environment than in the past when the models were built. One such important factor is the low level of interest rates, which makes it hard to come up with reasonable parameter estimates for the models, but also makes it hard to employ models that assume implicitly that interest rates never go below zero.

For now, we omit, however, the complex regulatory requirements set by the authorities, and adopt a simplified approach with no explicitly modelled liabilities. We do so to avoid delving too deep into the regulatory pecu-

liarities, and to concentrate on the basics of multi-period asset allocation modelling.

The research objectives for this thesis are consequently:

- Asset allocation model in a multi-period setting covering major asset classes of pension fund's portfolio
- Estimate up-to-date parameters and develop scenarios for multi-period optimisation
- Optimise a dynamic strategy in a two-stage problem and compare dynamic strategy to fixed strategy
- Study the stability of the optimal solution and discuss the parameter choices.

1.3 Structure of the thesis

The structure of this thesis is as follows: we start by introducing the subject and mentioning some of the most important literature in chapter 1. In the next chapter 2 we will have a detailed literature review and a look on the background of the problem. The next chapter 3 will be the theoretical part of this thesis covering the models we are interested in. In the last part (chapter 4) we demonstrate the theoretical framework in an empirical asset allocation example mimicking a real-life decision making situation. In chapter 5 we evaluate the performance of the model, and present a broader discussion on the topic. Further research topics are also presenter in chapter 5. In the last chapter 6 we conclude the thesis.

Chapter 2

Background and literature review

2.1 Modern portfolio theory

Modern Portfolio Theory (MPT) began with the work of the Nobel prize winner economist Harry Markowitz who developed the fundamental framework of portfolio theory and portfolio optimisation.

The Markowitz's model in its most basic formulation is

$$\begin{aligned} \min \quad & \sum_{i,j}^n w_i w_j \sigma_{i,j}, \\ \text{s.t.} \quad & \sum w_i \bar{r}_i = \bar{r}, \\ & \sum_i^n w_i = 1. \end{aligned} \tag{2.1}$$

In (2.1) w_i is the amount invested in asset i , $\sigma_{i,j}$ is the covariance of assets i and j , \bar{r}_i is the expected return of asset i and \bar{r} is the expected return of the portfolio. In (2.1) the minimisation is done so that the expected return of the portfolio is the weighted average of the constituents of the portfolio, and all the available resources are invested. (Markowitz, 1952; Luenberger, 1998)

Markowitz's work was further extended by Tobin by inclusion of the risk-free asset, in which the investor can invest. The work was further advanced by Sharpe by assuming that the same risk-free rate can be used for borrowing as well. The findings of Sharpe are more generally known as the capital-asset pricing model (CAPM). In the CAPM it is assumed that the expected return of an investable security is related only to its beta, a factor formalised by Sharpe and the CAPM. (Markowitz, 1999)

In the heart of the CAPM lies the capital market line given by

$$\bar{r}_i - r_f = \beta_i(\bar{r}_M - r_f), \quad (2.2)$$

where \bar{r}_i is the expected return of asset i , r_f is the risk-free rate, β_i is the beta coefficient of asset i , and \bar{r}_M is the expected market return. Equation (2.2) holds if the market portfolio M is efficient. The factor β_i is defined

$$\beta_i = \frac{\sigma_{iM}}{\sigma_M^2}, \quad (2.3)$$

where σ_{iM} is the covariance of excess asset returns and excess market returns, and σ_M^2 is the variance of excess market returns. In other words the factor β_i is the sensitivity of excess asset returns on the excess market returns. This factor is thought to capture the degree of non-diversifiable risk the asset is exposed to.

Another direction in which the models were developed was that of combining investors' market views on the supposedly neutral parameters typically estimated from the historical time series of market returns. The first model to succeed in doing this was called the Black-Litterman -model, named after the inventors. The BL-model managed to blend subjective market views on the neutral estimated market parameters in a coherent way. (Black and Litterman, 1992) More advances towards this direction were taken to eliminate the restrictive assumption of normality, Attilio Meucci proposed a way to use copula functions to blend investor views on the non-normal market distributions. (Meucci, 2005)

The mentioned developments above sought to broaden the model to a certain direction: making the estimates more accurate and more credible. However, these models are all single period models, hence they do not explicitly allow dynamic decision making and updating. It is easy to see some of the limitations these one-period models can have. By definition, these models do not account for the possibility of having to adjust the portfolio during the planning period. In addition, it is not clear what would be an optimal planning horizon for a one-period model.

For these reasons we resort to stochastic optimisation, which is a theoretical framework that allows multi-period considerations in portfolio optimisation problems. In the literature, it is even stated that applying a single-period model repeatedly for more than one period result to sub-optimal dynamic decisions. (Dupačová et al., 2003) The need for another approach, which allows adjustments to portfolios during the planning period is the starting point and motivation for this thesis.

2.2 Multi-period investing

The earliest dynamic models were developed soon after Markowitz's ground-breaking research. Merton, another Nobel laureate, published in the late 60's with his doctoral advisor Samuelson articles treating the multi-periodicity of the underlying problem.

The developments started by the Merton and Samuelson are not among their most notable findings, but the two are more famous for the Merton model, and Samuelson on the other hand has been titled the Father of Modern Economics.

The rationale for studying dynamic multi-period models more is that according to experts in the field they lead to superior results compared to myopic or static techniques. (Zenios and Ziemba, 2007) Instead of making long term fixed decisions an investment portfolio for example may be updated daily, weekly and so forth depending on the investment management problem in hand. In this kind of setting it is interesting to study models that allow a type of an updating scheme. The results tend to predict better performance for models that allow and take into account the possibility of updating the portfolio structure along the way.

A notable finding by Samuelson (1969) and Merton (1969) was that if the investor's utility function is iso-elastic, then the portfolio-selection is an independent decision on the consumption decision. Iso-elastic utility functions are also referred to as power utility functions.

Another popular belief Samuelson (1969) debunked was that an investor should invest differently according to their age and prospective future income streams. What Samuelson (1969) proved was that this is not the case, assuming the risk-tolerance and utility function remaining intact the investor's entire life. However, if the utility function changes, the previous reasoning does not hold.

2.3 Asset allocation strategies

In this section we make a quick overlook on different allocation strategies studied in the literature.

Buy and hold

First and simplest possible allocation strategy is the "buy-and-hold" strategy, where the decision maker makes a decision in the beginning of time, and holds the decision until the end of investment horizon. Such a decision could be

made based on single-period optimisation models described in the previous section. Best known and most widely used allocation model being probably the model by Harry Markowitz. (Markowitz, 1952)

Fixed mix

Slightly more sophisticated strategy would be so called "fixed-mix" strategy, where the investment decision is made in the beginning, and held constant through the investment period until the end of horizon. This strategy already assumes a multi-period setting, where asset allocations in the portfolio need to be updated every now and then.

Fixed-mix strategies play an important role in the asset allocation industry, because portfolio managers are often compared to performance of such strategies.

The theoretical justification for fixed-mix strategy exists as well, at least as long as asset returns are assumed to be i.i.d., the utility function is CRRA (constant relative risk aversion), only investment income is taken into account, and transaction costs are not included.

However, these assumptions listed above are some of the reasons why a more sophisticated multi-period framework is of interest at the first place. Transaction costs are in some cases significant, so they can not be ignored. In addition, asset returns may not be i.i.d. Moreover, incorporating future cash-flows due to a more complex balance sheet is one of the reasons to consider a more sophisticated multi-period model.

Constant-proportion strategies

Constant-proportion strategies are defined by Perold and Sharpe (1988) as a strategy, where the investor sets a floor for size of the portfolio below which it should not fall at any point. The floor is set to grow according to a riskless rate.

Formally the constant-proportion strategy is given by

$$\text{Invested amount in stocks} = m(\text{Assets} - \text{Floor}). \quad (2.4)$$

In Perold and Sharpe (1988) the authors explain the difference in parenthesis to form a "cushion", and according to the CPPI decision rule the investor should simply keep their exposure to equities a prefixed constant multiple of this "cushion".

Option-based portfolio insurance

In an option-based portfolio insurance strategy the definition begins by fixing the investment horizon, and a floor level in for value of assets in that given planning horizon. Although it is not said explicitly, the floor set to the end of the horizon defines implicitly floors for each time step until the end of planning horizon. This is done by discounting the floor from the end of planning horizon using the riskless interest rate. (Perold and Sharpe, 1988)

An important drawback of the OPBI strategy is that it is calendar dependent. Typically long-term investors have in real life an investment horizon much longer than their planning horizon, and it causes problems when the current OPBI strategy expires. A new set of investing rules is needed, when the previous one expires. (Perold and Sharpe, 1988)

Stochastic programming

One method to solve problem of generating dynamic allocation strategies is called stochastic programming. In stochastic programming, the optimisation problems are probabilistic. This means that instead of optimising a certain utility function, we are optimising the expected value of such a function.

Stochastic programming has proved to be an method to deal with the most general models, and there taking the transaction costs into account is easy. Moreover, the serial dependency of the return distributions of asset classes can be modelled. (Zenios and Ziemba, 2007)

Another advantage the stochastic programming (SP) has is that it can be directly expanded from an asset only analysis into an asset and liability management (ALM) problem. Including the liability side of the balance sheet is an essential requirement for a financial planning model by most institutional investors: pension funds and insurance companies. (Zenios and Ziemba, 2007; Ziemba, 2003)

In stochastic programming the uncertainty involved in investment decisions is modelled by scenario trees. Different techniques for estimating a scenario tree are discussed later in this thesis.

When a stochastic program in a multi-period setting is solved, the solution typically contains the optimal values of decision variables in each stage of the scenarios. Still, even though the solution contains the values for latter decisions in the program, the most interesting part of the solution is of course the first period decision, since that decision is done first, and the others can be considered later.

Stochastic dynamic programming

In the book *Handbook of Asset and Liability Management* (Zenios and Ziemba, 2007) stochastic dynamic programming is provided as another method to come up with decision policies for dynamic asset allocation strategies.

A major drawback of stochastic dynamic programming is the number of state variables it can cope with, the state variables being the decisions made in each time step. This is a widely recognised issue and often called as the "curse of dimensionality".

The main difference between stochastic programming models, and stochastic dynamic programming is the solution concept. In stochastic programming most of the emphasis is put to the first period decision, whereas in stochastic dynamic programming the aim is to establish decision rules, which could be applied by the system in which ever state. (Dupačová et al., 2003)

2.4 Utility functions

In order to optimise, we need to define what is the property to either maximise or minimise depending on the problem setting. This is an important issue at the heart of the problematics of portfolio selection, and in this section we show different types of typical utility functions available to describe the preferences of the decision maker.

Utility function is an individual trait of each investor. They are typically noted as $U(x)$, and in most cases we are interested in the expectation of the utility function: $E[U(x)]$. The argument x inside the utility function denotes the wealth level at which the utility is considered. In many cases wealth is also denoted as w . In some other contexts the argument could be measured in other units as well, like as the number of desirable outcomes, or an amount of a good or commodity in broader terms.

The most obvious and simple utility function is the linear risk neutral utility function: $U(x) = x$. For this utility function the investor ranks the obtained wealth levels linearly and does not account the risk, hence it is called the risk neutral utility function. (Luenberger, 1998)

Utility functions are increasing and continuous functions. If $x > y$ then $U(x) > U(y)$. The utility functions are used to rank different outcomes, so their cardinal value has little or no interpretation at all.

Table 2.1 lists well-known utility functions, see e.g. Luenberger (1998). A more profound look at the utility functions in portfolio selection problems can be found in an article by Kallberg and Ziemba (1983).

Now that we have the investor's utility function defined, we can study

Table 2.1: Widely used utility functions. (Luenberger, 1998)

Function name	Expression
Linear	$U(x) = x$
Exponential	$U(x) = -e^{-ax}$
Logarithmic	$U(x) = \ln x$
Power	$U(x) = bx^b$
Quadratic	$U(x) = x - bx^2$

the investor's attitude towards risk via the risk tolerance. The Arrow-Pratt absolute risk aversion coefficient (ARA) is defined (Luenberger, 1998)

$$ARA(x) = -\frac{U''(x)}{U'(x)}. \quad (2.5)$$

Respective relative risk aversion coefficient (RRA) is

$$RRA(x) = -x \frac{U''(x)}{U'(x)}. \quad (2.6)$$

Risk tolerance is then defined as the reciprocal of risk aversion. From previous equations (2.5) and (2.6) we can define the following risk tolerance coefficients

$$ART(x) = \frac{1}{ARA(x)}, \quad (2.7)$$

$$RRT(x) = \frac{1}{RRA(x)}. \quad (2.8)$$

In a typical case, risk aversion is thought to decrease when wealth increases. This feature is incorporated in the class of utility functions called hyperbolic absolute risk aversion (shorthand: HARA). In general case HARA functions can be written as

$$U(C) = \frac{1-\gamma}{\gamma} \left(\frac{\beta C}{1-\gamma} + \eta \right)^\gamma. \quad (2.9)$$

The class of so called HARA utility functions can obtain numerous forms by varying the parameters in the definition (2.9). Luenberger (1998) gives an exercise to show how the general form HARA yields all the different utility functions seen in the table 2.1 by just altering the parameters.

The utility maximisation problem leads to the formulation of von Neumann and Morgenstern, and the utility maximisation, so called von Neumann-Morgenstern function is

$$\max \sum_s p_s \cdot U(W_{s,t}), \quad (2.10)$$

where p_s is the probability of a scenario, $W_{s,t}$ the wealth in period t in scenario s , and U the utility function of an individual investor.

There exists a whole array of problem specific utility functions derived for different purposes and different types of asset allocation problems. A very general formulation is seen in the publication by the CFA institute *The stochastic programming approach to asset, liability, and wealth management* written by Ziemba (2003).

The general formulation of Ziemba is shown in equation (2.11). In this objective function the model maximises the expected value of wealth in the final period, penalised with the shortfall cost taking into account the risk tolerance the investor has

$$\max E \left[W_T - \frac{\text{Accumulated penalised shortfalls}}{\text{Risk tolerance}} \right]. \quad (2.11)$$

Here the risk tolerance is the reciprocal of the Arrow-Pratt or absolute risk aversion (ARA) coefficient seen earlier in the equation (2.7). (Ziemba, 2003) In previous equation (2.11) W_T stands for the wealth at time T .

In Ziemba's definition of objective function the accumulated penalised shortfalls are defined via an annual return goal, and falling short of this goal is then something penalised for. In a problem specific case, this shortfall variable could be defined as for example different levels of solvency capital.

Institutional investors like pension or life insurers face various types of legally imposed solvency constraints, which aim at keeping the companies away from going bankrupt. These legal restrictions are typically defined as the minimum amount for the net wealth, the positive difference between the value of assets and liabilities. In the European Solvency 2 regulation the solvency capital to cover the market risk is calculated based on the risk factor exposures of the asset portfolio and liabilities. Different asset classes face different capital requirements. For example, equities face higher capital requirement, and if the asset portfolio is considered very risky, the authority requires higher amount of net assets. The local regulation the Finnish pension funds face is qualitatively similar, and its properties are explained in the article by Hilli et al. (2007).

Similarly as the Ziemba's objective function in (2.11) is a construction of a risk neutral, but also a risk accounting part. The shortfall part accounting for the risk in the objective function could be defined as falling short of the solvency capital requirement.

Høyland and Wallace (2001a) construct the utility function for their insurance fund as the value in the end of planning horizon penalised by the discounted accumulated shortfall of certain regulative objective. The max-

imisation problem proposed by Høyland and Wallace (2001a) for a Norwegian life insurance company is

$$\max \sum_{s \in S_{scen}} P_s \left[\sum_{i \in S_{assets}} m_{val,i,T,s} - \sum_{j \in S_{sf}} \sum_{t \in S_{per}} \rho_{s,t} \cdot C_j(sf_{j,t,s}) \right]. \quad (2.12)$$

Here P_s stands for the probability of scenario s , where as scenario refers to an entire path from the beginning until the end of planning horizon. m_{val} is the market value of the investments in asset class i at time T in scenario s . The risk is incorporated in the objective function in the form of accumulated compounded shortfall costs, which are summated through time periods S_{per} . The compounding factor in scenario s at time t is $\rho_{s,t}$, and it is used to scale shortfalls occurring at different times to the same future value at the end of the planning horizon. The shortfall cost function is noted as $C_j(\cdot)$ and $sf_{j,t,s}$ is the shortfall of type j in scenario s at time t . The shortfalls of types j are different kinds of shortfalls, first of which is defined as *shortfall relative to the target capital adequacy*, and another *shortfall relative to the target solvency*. The definition of the model is fully described in Høyland and Wallace (2001a).

Another definition for the objective function is found in an article by Hilli et al. (2007), where a similar procedure is exploited. Here the authors have constructed an overall objective function

$$\max E \left[\sum_{t=1}^{T-1} (d_t u(C_t, B_t, H_t, L_t)) + d_T u_T(C_T, L_T) \right], \quad (2.13)$$

where d_t is a time-variant discount factor, C_t capital, B_t solvency border, H_t transfer to bonus reserve, L_t the technical reserves (liabilities), and u and u_T respective utility functions constructed by penalising the violation of solvency constraints.

Overall, the selection of the objective function in the optimisation is a process that requires many considerations to capture the of the investor preferences. We will construct our own objective function later in this thesis following the same principles discussed in this section.

Later in their article Høyland and Wallace (2001a) propose two alternatives for utility functions to test the derived results in the regulatory framework.

First of their proposals is the power utility function given by

$$U(\Delta W) = \frac{\Delta W^\gamma - 1}{\gamma}, \quad (2.14)$$

where ΔW refers to change in wealth. Here in the power utility function the wealth argument is defined to be the return of the portfolio to the end of the planning horizon.

And the other alternative for the utility is given by

$$U(\Delta W) = \Delta W - B \cdot e^{-C \cdot \Delta W}, \quad (2.15)$$

where ΔW is the change in wealth, B and C parameters that determine the level of risk aversion. In this latter formulation the terminal wealth is taken as the difference of the terminal portfolio value and the current portfolio value.

2.5 Scenarios for stochastic programming

Stochastic programming models need rigorous scenarios as an input to perform the optimisation and to come up with decision policies. The optimisation results depend entirely on the scenarios they are optimised on.

2.5.1 Scenario generation

In order to employ stochastic programming methods one has to develop scenarios under which the optimisation is carried out. Scenario generation is an important and integral part of the modelling process, and if done poorly the results will have little relevance to the real life. Hence, the process for generating scenarios should not be overlooked.

Methods for generating scenarios are numerous. Some different scenario tree generation methods are listed below:

- Clustering or path-based methods (Dupačová et al., 2000; Kaut and Wallace, 2003)
- Conditional sampling method (Kaut and Wallace, 2003)
- Sequential importance sampling (Dupačová et al., 2000)
- Moment matching (Høyland and Wallace, 2001b)
- Optimal discretisation (Pflug, 2001; Kaut and Wallace, 2003)
- Integration quadratures (Hilli et al., 2007)

Many of these methods require additional assumptions about the tree: the number of time steps, and the branching structure. Some of the methods mentioned are parametric, like the sampling methods, where we assume distributions, for which we have estimated parameters. Others are non-parametric, like moment matching, where the practitioner does not need to provide explicit assumptions for the distribution.

One important issue that is raised in the case of each of these scenario generation methods is the dimensionality. Each asset class can be seen to consist of multiple stochastic factors that need to be modelled independent from one another. On the other hand some assets may share some of the factors, which may allow dimension reduction by the means of factor analysis. (Dupačová et al., 2000)

During the process of building the model for this thesis we tried and applied clustering method with no success, hence the results from those studies are not included here. The scenarios in this thesis are based on the moment matching methodology as suggested and described by Høyland and Wallace (2001b). The method of integration quadratures as applied in an article by Hilli et al. (2007) seemed an interesting possibility, but it is not used due to the complexity of the model. Some aspect of the modelling in the article Hilli et al. (2007) are taken to supplement the proposed method by Høyland and Wallace (2001b).

Next we describe briefly each of the mentioned scenario generation methods found in the literature.

2.5.2 Clustering

In the clustering approach we take a simulated sample of data paths as a starting point for the discretisation process. In sequential clustering we cluster first according to ω_1 , after which we continue the clustering in ω sequentially to the subvector ω_2 and so forth. In a book by Dupačová et al. (2003) the authors refer to this method as a multi-level clustering scheme.

The simulated data path is noted as $(\omega_1, \dots, \omega_T)$. The clustering is done based on a dissimilarity measure. In (Dupačová et al., 2003) this measure is defined

$$d(\omega^s, \omega^{s'}) = \sum_{t=1}^T w_t \|\omega_t^s - \omega_t^{s'}\|. \quad (2.16)$$

In (2.16) the dissimilarity measure can be chosen accordingly, and weights w_t can be set so that earlier observations are given more importance in the estimation procedure. The original data path is noted by ω_t^s , and the clustered path is $\omega_t^{s'}$.

In multi-level clustering the first clustering results in a number of clusters, noted as $C_1^2, \dots, C_1^{K_2}$. The clustering scheme continues for each defined cluster C_1^k separately. For this second round of clustering the simulations ω_2 will be used.

Xu et al. (2012) describe a k -means clustering algorithm accounting for interstage dependencies. The model they describe does not, however, take into account the problem of having multiple variables, each of which having their own properties and respective scenario trees.

Clustering as such may not be a sufficient method for scenario generation, and ways to improve its performance are discussed in Pflug (2001). The technique they have developed is called "optimal discretisation".

Drawbacks of clustering methods include the generated tree may not be statistically close to the original distribution, and running into problems in clustering when the sequentially clustered group becomes very small in a prescribed scenario tree generation, and the clustering can no longer be done according to the specification.

2.5.3 Sequential simulation with k -means clustering

One way to overcome the evident drawbacks of clustering methods is to use sequential simulation combined with k -means clustering algorithm. This method is described in an article by Xu et al. (2012), and is called by the authors "The Hybrid Sequential Generation Method for Multi-stage Scenario Tree".

Essentially what happens in this algorithm is that the data path is simulated using a time-series model for each time step. In Xu et al. (2012) it is a combined VAR-MGARCH model. For this simulated path k -means clustering is utilised to select a representative arch from each node. It is noteworthy that the number of possible outcomes grows exponentially with the number of variables.

2.5.4 Moment matching

Because the clustering approach may be seen as a crude ad-hoc solution for generating scenario trees, we studied some more suited methods as well. A method called moment matching (or moment fitting) means estimating the scenario tree so that the properties of the discretised distribution match as well as possible the properties of the theoretical, or observed distribution from node to node in the scenario tree.

The notation used in this thesis for moment matching method follows the one used in Høyland and Wallace (2001b). Their model is described in more

detail in the next chapter 3.

2.5.5 Property matching

Moment matching method is relatively straightforward and intuitive, but so is its slightly more sophisticated counterpart which is called property matching, or distribution matching. In property matching all the available information concerning the underlying distribution is used, and no information is wasted, as in moment matching, when only certain moments are fitted.

The methodology is described more thoroughly in Calfa et al. (2014). In property matching the cumulative distribution function of the prescribed scenario tree is compared and matched to the actual cumulative distribution function, which can be empirical, or fitted to observed data.

2.5.6 Integration quadratures

The most sophisticated model mentioned here is the integration quadratures method applied by Hilli et al. (2007) and explained in Pennanen and Koivu (2002). The method is actually a modification and improvement to the conditional sampling methods used to generate scenario trees.

The approach in Hilli et al. (2007) is based on so called vector equilibrium correction model, which is a type of an econometric model taking into account the long term equilibria of certain econometric factors. The short term dynamics of these economic factors is taken into account as well.

The discretisation scheme built atop of the vector equilibrium model is based on earlier work by Pennanen and Koivu (2002). These together form an integrated model, which employs sophisticated econometric modelling, and a reasonable discretisation, which ensures fast convergence of the optimal solution when the number of scenarios grows.

Chapter 3

Theoretical model description

3.1 Stochastic programming models

Stochastic programming models described here are extensively covered by for example Jitka Dupacova in her article. (Dupačová, 1995)

The notation used in this chapter follows Dupacova's notation. In the general formulation of a T -stage stochastic program following we need to define a stochastic data process

$$\omega = \{\omega_1, \dots, \omega_T\}. \quad (3.1)$$

The realisations of ω are trajectories or paths in a probability space (Ω, \mathcal{F}, P) . We define a decision process $x = \{x_1, \dots, x_T\}$, which is a measurable function of ω . We will have a sequence of decisions and observations

$$x_1, \omega_1, x_2(x_1, \omega_1), \omega_2, \dots, x_T(x_1, \dots, x_{T-1}, \omega_1, \dots, \omega_{T-1}), \quad (3.2)$$

where every decision made depends on the data and decision path so far. In a sophisticated model the function would account also for the overall costs related to the problem. In addition an important requirement is that the decision process is nonanticipative, which means the current decision does not depend on future realisations. The T -stage stochastic program can then be written in general case

$$\begin{aligned} \min_x \quad & \mathbb{E}[f_0(x, \omega)], \\ \text{subject to} \quad & \mathbb{E}[f_i(x, \omega)] \leq 0, \quad i = 1, \dots, k, \\ & \mathbb{E}[f_i(x, \omega)] = 0, \quad i = k + 1, \dots, k + r, \\ & x \in \chi. \end{aligned} \quad (3.3)$$

In this notation by Dupačová (1995) ω is a random parameter in a probability space (Ω, \mathcal{F}, P) , χ is a nonempty closed set. The objective function to be minimised is $f_0 : R^n \times \Omega \rightarrow R^1 \cup \{+\infty\}$, and the constraint functions $f_i : R^n \times \Omega \rightarrow R^1$. f_0 and f_i are given functions. The decision variable in this stochastic program is x and ω is a random parameter as shown in (3.2). Indices $i = 1, \dots, k$ note the inequality constraints if the program, and the remaining $i = k + 1, \dots, k + r$ stand for the equality constraints. Different decision stages $t = 1, \dots, T$ are embedded in the decision variable x and in the stochastic variable ω as seen in (3.2).

Solving the stochastic programming problems has been extensively covered in an article by Dupačová (1995). They have concentrated mostly on solving stochastic linear programming models, but these are no longer relevant to us because the objective functions we intend to use are nonlinear. This results in more complicated solving, but fortunately the capacity modelling such problems has developed much since the earliest stochastic linear programs.

One important element in defining stochastic programs is to ensure the nonanticipativity property holds. This means simply that decisions are made before events occur. In the scenario tree formulation it implies that in each node independent of the actual scenario the decisions have to be equal.

3.1.1 Portfolio optimisation by stochastic programming

So far the model we have explained has been the general case, but now we will show more detailed version of the model intended for asset allocation.

Now the problem can be written as a deterministic program

$$\begin{aligned} \max_x \quad & \sum_s p_s u(x_s), \\ \text{subject to} \quad & \\ & A_s x_s = b_s, \\ & s = 1, \dots, S. \end{aligned} \tag{3.4}$$

In (3.4), function u is the utility function. Decision variables are x and A and b are the constraint matrices and vectors. Probability of each scenario s is p_s . Formulation is the so called Von Neumann-Morgenstern utility function.

3.2 Scenarios for stochastic programming

In this section we first introduce the basic requirements for scenarios for stochastic programs. We then cover some of the models used for generating




Figure 3.1: Illustrative graph represented by its adjacency matrix in (3.5)

scenarios and scenario trees. Some of those methods are based on expert opinion, cutting and pasting, and manual combining of intuitive scenarios, but others have more solid theoretical background.

In this thesis we use moment matching as the scenario generation algorithm as mentioned in chapter 2.5.4 and explained in depth in chapter 3.2.2.

3.2.1 Scenario trees

Many calculations in the tree are carried out using the graph theoretical representation of a directed graph, namely the adjacency matrix. The adjacency matrix A of a directed graph of five elements $N = 5$ reads

$$A = \begin{pmatrix} 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 \end{pmatrix}, \quad (3.5)$$

and the respective graph is shown in 3.1.

The graph in figure 3.1 is an example to illustrate the use of adjacency matrices. In the matrix (3.5) the rows and columns represent a node in the graph. Non-zero elements in a row represent an arc pointing at the node, and the columns with non-zero elements the arcs leaving the node. Each element in (3.5) represents an arc of the graph. Elements that are one exist, and zeros do not exist. For example, the fifth node has two arcs pointing at it, but none leaving. Hence, the fifth column has two non-zero elements, but the fifth row is all zeros.

The graph theoretical representation of a scenario tree has some advantages, for example calculating the probabilities for each state reduces to a matrix multiplication

$$C = B \cdot A^t, \quad (3.6)$$

where A is the adjacency matrix, B is a matrix with all elements zeros except the one representing the initial node, typically $\{i, j\} = \{1, 1\}$. The resulting

matrix C contains on the first row the final nodes that were accessible though the graph by t time steps in the graph from the initial state.

Other useful properties are that if ones in the adjacency matrix are replaced by the probabilities of its arcs, the resulting matrix C in multiplication (3.6) contains on the first row the probabilities of being in the respective state after t transitions from the initial state represented by matrix B . This modified adjacency matrix is usually called transition matrix.

Moreover, instead of probabilities, the adjacency matrix can be populated with arch values, and in this thesis these values represent the asset returns occurring in transition. In this case the multiplication (3.6) allows us to calculate the final value of an asset portfolio.

The above formulation is practical and easy to use, but it has one very severe shortcoming: matrix operations and especially matrix power operations are computationally very costly. Hence, this method is not suited for optimisation, where the final values need to be calculated iteratively many times.

Another way to express a scenario tree is matrix notation

$$M = \begin{pmatrix} 1 & 2 & 5 & 11 \\ 1 & 2 & 5 & 12 \\ 1 & 2 & 6 & 13 \\ 1 & 2 & 6 & 14 \\ 1 & 3 & 7 & 15 \\ 1 & 3 & 7 & 16 \\ 1 & 3 & 8 & 17 \\ 1 & 3 & 8 & 18 \\ 1 & 4 & 9 & 19 \\ 1 & 4 & 9 & 20 \\ 1 & 4 & 10 & 21 \\ 1 & 4 & 10 & 22 \end{pmatrix}, \quad (3.7)$$

where each node is indexed similarly as in figure 3.2. This matrix has one row assigned for each individual scenario. This type of a matrix can accomodate any tree shape, be it regular or irregular. In a book by Dupačová et al. (2003) the matrix in (3.7) is called a scenario tree nodal partition matrix.

This compact matrix formulation allows simple and effective scenariowise calculations in a tool like MATLAB, which can then be aggregated by the scenario probabilities.

In this thesis for simplicity we assume that the scenario trees have regular shapes. This means that for each node in a specific time step the branching is identical. In different times branching does not need to be identical, for

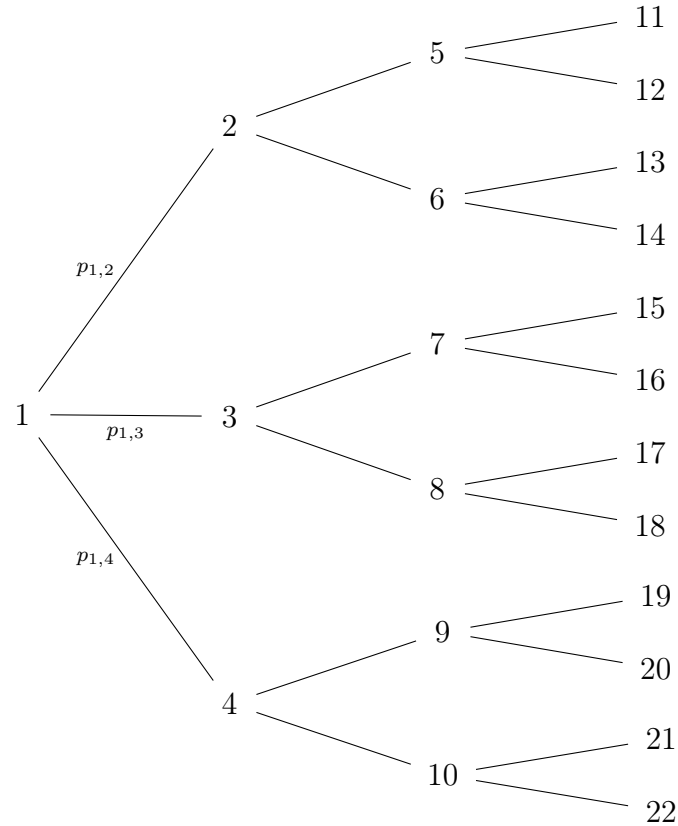


Figure 3.2: Larger scenario tree of branching structure $\{3, 2, 2\}$, and $T = 3$. First transition probabilities from root node $1 \rightarrow j$ are noted as $p_{1,j}$.

example at the earlier stages more branching can occur, whereas in the later stages tree may be sparse.

As uncertainty grows in time it could be a good idea to allow more branching in the early stages and less in the later ones. An illustrative example is seen in figure 3.2. This tree has $N = 1 + 3 + 3 * 2 + 3 * 2^2 = 22$ nodes, and $N - 1$ arcs. The tree is a directed graph from left to right, so it is unnecessary to draw the arrows to indicate the direction. After $T = 3$ transitions we end up to the final node, which are found on the right edge of the graph.

The transition matrix for the scenario tree in figure 3.2 is written in (3.8). This matrix A_p is of the size $n \times n$, where n is the number of nodes in the tree. In this case $n = 22$. Only the first 10 rows of A_T contain non-zero elements.

$$A_p = \begin{pmatrix} 0 & p_{1,2} & p_{1,3} & p_{1,4} & 0 & \cdots & 0 \\ 0 & 0 & 0 & 0 & p_{2,5} & \cdots & 0 \\ \vdots & & & \ddots & & & \vdots \\ 0 & 0 & 0 & 0 & 0 & \cdots & p_{10,22} \\ 0 & & & \cdots & & & 0 \\ \vdots & & & \ddots & & & \vdots \\ 0 & & & \cdots & & & 0 \end{pmatrix} \quad (3.8)$$

An important property of transition matrix is that the sums of its rows are ones for all the parent nodes (10 first nodes), and zeros to all final nodes (nodes 11-22).

Probability of each node in the final stage can then be calculated by

$$C_p = B \cdot A_p^T, \quad (3.9)$$

where C_p is the matrix of final probabilities, B a matrix indicating the first node, A_p the transition matrix as in (3.8) and T the length. Probabilities of the final stages can be read from the first row of C_p in columns 11-22.

Dimensions of scenario trees

By the depth of the tree we mean the number of time steps modelled. There are numerous factors to consider when choosing the planning horizon. In this thesis we settle for planning period length $T = 2$. We chose this length, because it is short enough to keep the problems of manageable size, but long enough to demonstrate the effects of multi-periodicity. In our example one time step will be one year, and the planning horizon two years. Any other period length would be equally suited for the model, and the steps are not required to be of equal length. As an example, in the model by Hilli et al. (2007), the model consists of three time steps, each of which is of different length.

Another dimension in which the tree structure needs to be considered is the width, in other words, the number of branches from each node in each state. This relates back to level of accuracy required of the model. For example, if the number of branches from a certain node is large enough, one will be able to match the moments of the distribution perfectly. But when the number of branches is small, there is no perfect fit. In this case it is important to give relative weights for each moment so that the matching algorithm will not form a tree based on the less important higher moments instead of the most important ones, such as mean and standard deviation. Moreover, if the branching reduces to only two branches leaving the node, higher moments like

skewness and kurtosis are no longer defined, and matching those moments will not yield anything.

3.2.2 Moment matching

To measure the similarity of the moments, we use the square norm as the measure as suggested by Høyland and Wallace (2001b)

$$\begin{aligned} & \min \sum_{i \in S} w_i \cdot (f_i(x, p) - V_{VALi})^2, \\ & \text{subject to} \\ & \sum p \cdot M = 1, \\ & p \geq 0. \end{aligned} \tag{3.10}$$

The optimised variables are vectors x and p , where x is the value of the random variable, and p its probability. Function $f(x, p)$ stands for the statistical property of index i in scenario S . M is a matrix of zeros and ones, and it has equally many rows as p has elements, and equally many columns as the scenario tree has nodes. Ones in a column of M indicate the conditional distribution in the node. Lastly w_i is the weight, hence the importance, of the statistical property i in S . In this thesis the highest weights were given to the lower order moments, and weight decrease when the moment order grows. In the model used in this thesis we treat p as a variable that is optimised in the previous equation 3.10. It could also be treated as a parameter

In the tree generation process we follow Høyland and Wallace (2001b) in the methodology. The authors suggest keeping certain moments constant through the scenario tree and the time horizon, where as some are state dependent. Kurtosis and skewness are held constant the entire period, but standard deviation and expected value are state dependent. The authors Høyland and Wallace (2001b) provide formulas for calculating the dependent moment value.

For standard deviation they propose

$$SD(x_{i,t}) = VC_i \cdot |x_{i,t-1} - E(x_{i,t-1})| + (1 - VC_i) \cdot SD_{AV}(x_{i,t}), \tag{3.11}$$

where $VC_i \in [0, 1]$ is the volatility clumping parameter, and SD_{AV} the average standard deviation for the specific asset class i .

For interest rate dependent asset classes the mean reversion effect needs to be modelled. The mean reversion effect is

$$E(x_{i,t}) = MRF_i \cdot MRL_i + (1 - MRF_i) \cdot x_{i,t-1}. \tag{3.12}$$

Table 3.1: Parameters for the moment matching model equations.

Symbol	Value	Explanation
VC	0.3	Volatility clumping
MRF	0.2	Mean reversion factor
MRL_c	4.0%	Mean reversion level for cash
MRL_b	5.8%	Mean reversion level for bonds
RP	0.3	Risk premium

In (3.12) $MRF_i \in [0, 1]$ is the mean reversion factor, and MRL_i the mean reversion level.

For equity asset class similar mean reversion effect does not occur, but there is assumed a risk premium. This means, that for additional risk taken a higher return has to be in sight. In Høyland and Wallace (2001b) the expected return for equity asset class is given by

$$E(x_{i,t}) = r_{t-1} + RP_i \cdot SD(x_{i,t}). \quad (3.13)$$

In (3.13) r_{t-1} is the previous risk-free interest rate, $SD(x_{i,t})$ the standard deviation for the respective equity asset class. RP_t is the risk premium constant.

The fixed parameters and their given values in the model in (Høyland and Wallace, 2001b) are listed in the following table 4.3.

Clearly, one can argue that the levels for mean reversion seen in table 4.3 are no longer relevant due to current level of interest rates. Hence, at least the parameters for mean reversion need to be re-calibrated in order to use the model today.

Calculating the individual moments has been explicitly described in the literature, and the moment matching problem is stated as follows. (Calfa et al., 2014)

Following the notation in Calfa et al. (2014) we write the L^2 moment matching problem (shorthand: MMP)

$$\min_{x,p} z_{MMP}^{L^2} = \sum_{i \in I} \sum_{k \in K} w_{i,k} (m_{i,k} - M_{i,k})^2 + \sum_{\substack{i \leq i' \\ (i,i') \in I}} w_{i,i'} (c_{i,i'} - C_{i,i'})^2,$$

subject to

$$\begin{aligned} \sum_{j=1}^N p_j &= 1, \\ m_{i,1} &= \sum_{j=1}^N x_{i,j} p_j, \quad \forall i \in I, \\ m_{i,k} &= \sum_{j=1}^N (x_{i,j} - m_{i,j})^k p_j, \quad \forall i \in I, k \geq 1, \\ c_{i,i'} &= \sum_{j=1}^N (x_{i,j} - m_{i,j})(x_{i',j} - m_{i',1}) p_j, \quad \forall (i,i') \in I, i \leq i', \\ x_{i,j} &\in [x_{i,j}^{LB}, x_{i,j}^{UB}], \quad \forall i \in I, j = 1, \dots, N, \\ p_j &\in [0, 1], \quad \forall j = 1, \dots, N. \end{aligned} \tag{3.14}$$

In (3.14) the minimised measure of goodness of fit is the L^2 -norm $z_{MMP}^{L^2}$. Weights of each moment are $w_{i,k}$, where i is the index corresponding to each asset class, and index k corresponds to each moment order. $M_{i,k}$ is the moment value targeted, and $m_{i,k}$ the current value in the scenario tree. Similarly $C_{i,i'}$ is the ideal correlation, and $c_{i,i'}$ is the correlation in the tree. Formulas for calculating $m_{i,k}$ and $c_{i,i'}$ are given below, and upper and lower bound for optimised variable are given by $x_{i,j}^{LB}$ and $x_{i,j}^{UB}$.

A clear shortcoming of the model described by Calfa et al. (2014) is that they have not incorporated a time dependent part to their model. This is, however, taken into account by Høyland and Wallace (2001b) and Xu et al. (2012).

3.3 Interest rate models

In (3.12) interest rate is modelled to be mean reverting. This is a typical assumption for interest rate models, and the model in (3.12) and proposed in an article by Høyland and Wallace (2001b) is compatible with the oldest so called short rate model, The Vasicek model. (Vasicek, 1977) The Vasicek model as its most common formulation is

$$dr_t = a(b - r_t)dt + \sigma dW_t, \tag{3.15}$$

where dr_t is the rate change at time t , a is a mean reversion factor, b is a long-term equilibrium rate, r_t the rate at time t , σ is standard deviation and dW_t a Wiener process.

An important issue arising from the Vasicek model is that it allows rates to go negative without any restrictions. During the historical periods of relatively high interest rates this issue does not rise, or at least it does not cause significant issues for modelling. However, during current extremely low yield levels it is necessary to take into account the anomalies this may cause.

This is a problem in the traditional economic sense, because it is usually thought that negative interest rates do not occur. This is caused by the assumption that the banks are not able to impose negative rates on their clients. During current turmoil and economic environment in the developed countries this assumption needs to be reconsidered.

One of the early contributions to tackle the issue with negative rates was a model by Cox, Ingersoll and Ross, which is nowadays known as the CIR-model. (Cox et al., 1985) The CIR-model is similar to Vasicek's model, but in the CIR-model it is assumed that the standard deviation of the underlying process depends on the current yield level.

One formulation of the CIR-model is

$$dr_t = a(b - r_t)dt + \sigma\sqrt{r_t}dW_t. \quad (3.16)$$

A more general formulation of short-rate models can be written as

$$dr_t = \alpha(\beta - r_t)dt + \sigma \cdot r_t^\gamma dW_t, \quad (3.17)$$

where γ is the model parameter determining the dampening of volatility. The formulation shown in (3.17) as seen in an article by Chan et al. (1992). Different models with slightly differing parameter specifications exist, but in this thesis we use only the simplest models, the model by the Vasicek and the CIR-model.

Such one factor short-rate models can be estimated by various means. One of the most popular ones is generalised method of moments (GMM), which is first used in this context in an article by Chan et al. (1992), and was first described by Hansen (1982).

Following the econometric formulation in the article by Chan et al. (1992) we define the discrete version of a short-rate model as follows in (3.18) and (3.19)

$$r_{t+1} - r_t = a + b \cdot r_t + \epsilon_{t+1}, \quad (3.18)$$

where $a = \alpha$, and $b = 1 + \beta$, and

$$E(\epsilon_{t+1}) = 0, \quad E(\epsilon_{t+1}^2) = \sigma^2 r_t^{2\gamma}. \quad (3.19)$$

Further, the error term in (3.18) can be written as

$$\epsilon_{t+1} = r_{t+1} - r_t - a - b \cdot r_t. \quad (3.20)$$

The modification taking into account the dampening of volatility at lower rate levels is easy to take into account in the scenario tree generation process. In this thesis we replace the original formulation of standard deviation seen in (3.11) for interest rate trees by

$$SD(r_{t+1} - r_t) = \sqrt{r_t} \sigma, \quad (3.21)$$

where $SD(r_{t+1})$ is the standard deviation of the next branching, r_t the current rate, and σ the estimated standard deviation parameter as in the CIR-model (3.16).

An obvious issue arises when the yield level drops below zero. For rates $r_t < 0$ the square root in (3.16) can become non-real number containing an imaginary part. This does not have any sensible interpretation, hence we will modify the standard deviation for interest rates as

$$SD(r_{t+1} - r_t) = \sqrt{|r_t|} \sigma, \quad (3.22)$$

where we take the absolute value of the interest rate. This keeps the model properly behaved even in the event of below zero rates without the need to impose an artificial floor for interest rate values at zero level.

The inclusion of taking the absolute value of the argument in the square root in the CIR-model is a rather ad-hoc solution proposed by practitioners. We were unable to find any scientific articles proposing similar methodology. Theoretically this problem does not arise, because the model is a continuous-time model by definition. In practice, it is used as a discrete model, and the original model assumptions no longer hold.

For the credit spreads, whose properties will be described in more detail later in this thesis, we impose a lower bound at zero level. Here credit spread means the difference between a risky corporate interest rate, and a riskless or low-risk government interest rate. A negative spread would imply the respective corporate bonds to be considered less risky than government debt. Now it is in principle possible that even the interest rate for corporate debt could go below zero, but this is very unlikely, because the positive spread usually keeps the corporate rate higher than the potentially sub-zero government rate.

3.4 Objective function

In this thesis we will not make a full balance sheet analysis, but our approach concentrates on the asset side. There are numerous reasons for this choice,

and one of the most important ones is to keep the model of manageable complexity.

The objective function for our optimisation problem will be constructed by taking inspiration from articles mentioned in the literature review section.

We will have an objective function consisting of two parts, and it is very close to the one described by Ziemba (2003). The objective function of the optimisation problem is

$$U(x) \equiv U(x, w(x)) = w_T(x) - \sum_{t=1}^T \rho_t \cdot sf_t(x). \quad (3.23)$$

Here x is the decision variable, namely the allocation in different asset classes. The wealth in the end of the planning horizon is w_T , ρ_t is a compounding factor at time t to project the cost to the planning horizon. sf is the shortfall cost function, formally given by

$$sf_t(x) \equiv sf(r(x), t) = \begin{cases} a \cdot e^{-b(r_t(x) - r_g(t))} - 1, & \text{if } r_t < r_g(t), \\ 0, & \text{if } r_t \geq r_g(t). \end{cases} \quad (3.24)$$

In (3.24) $r_g(t)$ is the cumulative return goal. If the goal is set to 4% annually then the cumulative return goal at the end of year two would be given by $(1 + 0.04)^2 - 1$. Wealth w_t translates similarly into return by means of linear return: $r_t = w_t/w_0 - 1$.




Figure 3.3: The utility function in equation (3.23) in a one-period example and its risk aversions as in equations (2.5) and (2.6).

The parameters of the shortfall function determine the level of risk aversion. Parameter b in the exponent controls the aversiveness towards extreme events, whereas the multiplier a controls a more general type of risk aversion. The utility function as specified in (3.23) is shown in figure 3.3 with three

different parameter combinations. Here the utility function is plotted in a single-period case for simplicity, and no compounding of target returns or shortfalls occur.

The utility function is linear above wealths corresponding to exceeding the target return, which is in 3.3 set to 2%. Because linear utility function is the risk neutral utility function, the risk aversion coefficients drop to zero when the threshold is reached. Below the wealth level indicated by the target return, we can see how the utility function starts to penalise for shortfalls. Higher b means higher penalty for extreme events. (Høyland and Wallace, 2001a) This behaviour can be seen in the plots of ARA and RRA in figure 3.3, where the risk aversions cross at certain point, and risk aversions for utilities with high b are higher, even though the order was the other way around for low shortfalls.

In the deterministic equivalent form the objective function for the optimisation problem is

$$U(x) = \sum_{s \in S} p_s w_{s,T}(x) - \sum_{s \in S} p_s \sum_{t=1}^T \rho_{s,t} \cdot sf_t(x). \quad (3.25)$$

Here, p_s is the probability of scenario s , and $\rho_{s,t}$ is the compounding factor defined by the risk-free money market rate, which we calculate by linear interpolation from the two interest rates the model specifies. The discounting for our model will be in this case for period lengths 3 and less, hence using either of the two estimated stochastic rate parameters of terms 0.5 years and 5 years would be flawed.

Only two low risk interest rates are estimated in each scenario, and it caused that constructing a yield curve is not a solution without making additional assumptions. Therefore linear interpolation seems the most justifiable way to have some sort of an estimate for the compounding factors.

Chapter 4

Implementation and numerical results

In this chapter we describe the implementation results of our work. First the results for estimating the model parameters are explained, and later these are used to generate scenarios. Finally we end up with optimised allocation decisions based on the scenarios.

4.1 Data for model calibration

In this thesis the implementation has been written in MATLAB. (MATLAB, 2015) The model has been built to employ four different risk factors, one of which is an equity factor, two others interest rate linked factors, and the fourth represents the investments for real estates.

To further justify our choice of modelled asset classes, we look as an illustrative example the investments of Finnish pension funds. The Finnish Pension Alliance (TELA) reports high level summaries of an average allocation of the industry. In this review investments are divided into six asset classes. With our simplified model of only four distinct asset classes we can maybe get an idea of the allocation problem pension funds are facing. The selected asset classes we selected can be seen to cover approximately 93% of the portfolio of an average pension fund in Finland. (The Finnish Pension Alliance TELA, 2016)

These asset classes and their relative weight in the average portfolio of a pension fund is shown in figure 4.1. Same information is presented in the table 4.1. The data is given by The Finnish Pension Alliance TELA publicly on their website. (The Finnish Pension Alliance TELA, 2016)

Finnish pension funds have their assets split roughly to these six classes,




Figure 4.1: Finnish pension fund asset allocation, percentages.

Table 4.1: Pension fund asset allocation in percentage.

Asset class	2007	2008	2009	2010	2011	2012	2013	2014
Money market	5.5	6.9	5.9	4.9	8.1	6.4	6	4.4
Loans	2.7	3.9	3.4	2.8	2.7	2.4	1.9	1.9
Loans to policyholders	0.6	2.9	3.9	3.7	2.8	2.1	1.4	0.9
Bonds	35.7	41.1	37.4	31.9	35.2	35.6	33.8	34.2
Real estate	9.1	12.5	11.1	10.3	11.2	10.8	10.5	9.9
Equity	46.4	32.8	38.4	46.4	40	42.6	46.4	48.9

some of which require a short explanation. Loans to policyholders are a historically important, nowadays less important part of the investment portfolio. These loans are issued to policyholders, and there the customers of a pension fund can borrow back according to a fixed set of rules. The other class of loans, on the other hand, are investment loans, where the company and the borrower negotiate the terms, and no fixed set of rules exist for issuing them.

These two categories are omitted in this thesis due to their shrinking importance, and a rather peculiar nature. Similar arrangements are not very common elsewhere, and their weight in an average portfolio of a pension fund is less than 3%.

In our model we include five asset classes: equities, money market investments, government bonds, investment grade bonds and real estate investments.

The asset class of bonds makes a major stake of the asset portfolio of the Finnish pension funds. This asset class divides further into government bonds, and corporate bonds. Corporate bonds can be categorised even further to investment grade, and high yield bonds according to the credit rating of the bond. As the government bond yields are historically low at this time, and pension funds are looking for higher returns to cover the expenses in future, and investing in close to zero rates is unappealing. Furthermore, the shift from government debt to corporate debt is due to shift in risk sentiment: government debt is perhaps no longer regarded as a perfectly riskless asset class, although very low risk depending on the government in question.

Real estate investments make roughly 10% of the asset portfolio of an average pension fund in Finland. Their nature is rather illiquid, and modelling their return distributions based on historical data is difficult.

4.1.1 Market invariants

In order to model returns of asset classes listed earlier, we need to define their respective market invariants.

For equity type of investments the invariant is usually defined to be the compound return, see e.g. Meucci (2009). Market invariant for equity and real estate investments, the compound returns are

$$X_{t,\tau} = \ln(P_t) - \ln(P_{t-\tau}). \quad (4.1)$$

For fixed income investments the respective market invariant is not the compound return, but the difference of yields. This difference is

$$X_{t,\tau} = Y_t - Y_{t-\tau}. \quad (4.2)$$

Selected asset classes and the respective indices used for modelling them are shown in table 4.2. In table 4.2 and some other occasions later in this thesis shorthand IG is used for investment grade bonds. Government bonds are shortened as gov. bonds. Money markets are referred to by the shorthand notation MM.

Development of equity index and calculated real estate index are shown in figure 4.2. Development of fixed income investments, the yield of government bonds and credit spread of investment grade bonds is shown in figure 4.3.

4.1.2 Money markets

Short term interest, the money market rate, is modelled with a short rate model as shown earlier in section 3.3. The return generated by money market

Table 4.2: Selected asset classes for the model, and the indices used to calibrate the model. Real estate data is based on Official Statistics of Finland (OSF) (2016), and other indices downloaded from Bloomberg Terminal (Bloomberg Finance L.P., 2016).

Asset class	Data since	Index name
Equities	29/03/88	Deutsche Boerse AG German Stock Index DAX
Money markets	29/03/99	The BofA Merrill Lynch <1 Year Euro Government Index
Gov. bonds	29/03/88	The BofA Merrill Lynch 1-10 Year Euro Government Index
IG bonds	27/12/96	The BofA Merrill Lynch 1-10 Year Euro Corporate Index
Real estate	29/03/88	Real estate prices in Finland




Figure 4.2: Historical development of studied indices. 100 points in March 1999, but longer history included for equity and real estate indices.

investments will be modelled using the approximation given in Hilli et al. (2007), which is

$$R_t = ((1 + r_{t-1})(1 + r_t))^{\frac{1}{2}}. \quad (4.3)$$




Figure 4.3: Historical development of studied money market yield government bond yield, and credit spread of investment grade bonds.

In other words, money market return is calculated as the geometric average of preceding interest rates.


4.1.3 Government bonds and bond rate

The bond rate will be modelled according to the interest rate models explained in section 3.3. The model we assumed is the CIR-model, which has the property of dampening volatility when rates fall near zero.


The total return of the government bonds and money market investments will be calculated using the approximation proposed in the book *The Econometrics of Financial Markets*, Campbell et al. (1997) and used in practice in Pennanen and Hilli in their article (Hilli et al., 2007). The approximation is given by

$$R_t = \left(\frac{1 + Y_{t-1}}{1 + Y_t} \right)^D + \frac{1}{2} \cdot (Y_{t-1} + Y_t) \cdot \tau. \quad (4.4)$$

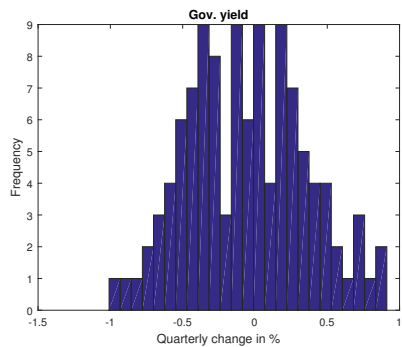
The previous equation (4.6) consists of the yield Y_t and bond duration D . Parameter τ is the length of time step in years. The first part of the summation accounts for the change in the bond value, and the latter part is the




(a) Quarterly compound returns of equities



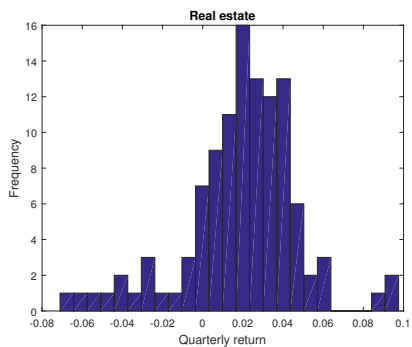
(b) Quarterly changes of money market rate.



(c) Quarterly changes of government yield.



(d) Quarterly changes of spread.



(e) Quarterly compound returns of real estate

Figure 4.4: Historical distributions of studied market invariants.

average of the two most recent yield levels accounting for the cash component of return. These components together make the total return of the bond.

According to our quick test, the approximation given by Pennanen and Hilli works well and the correlation between the observed index returns and the returns given by the approximation is high.

4.1.4 Investment grade bonds and credit spread

Riskier corporate bonds are modelled using their credit spread to the riskless government bond rate. The spread is calculated as shown in equation (4.5), where Y_S is the credit spread, Y_C the yield of a corporate bond, and Y_G yield of the government bond

$$Y_S = Y_C - Y_G. \quad (4.5)$$

Like the government bond yield, the credit spread is modelled according to the interest rate models shown in subsection 3.3. Another alternative could be treating the changes in corporate bond yield as the market invariant similarly to government yields. In this thesis, however, we treat the spread separately, and assume that in the end we can solve the corporate bond yield according to the equation (4.5). Assuming a mean reverting process for credit spread seems reasonable regarding the historical development seen in figure 4.3.

Price of the credit bond will be calculated using the same approximation as in an article by Pennanen and Hilli shown in equation (4.6). (Hilli et al., 2007) However, we make a small adjustment for the returns of investment grade returns: the model assumes the return consisting of only change in value, and cash component generating cash flow. It neglects the fact that sometimes corporates do default, which generates losses. Omitting the loss rate could make the scenarios for corporate bonds look too optimistic, and they dominate optimised allocations in the cases of high risk aversion. Hence, we adjust the return induced by investment grade bonds through

$$R_t = R_{bond}(t) \cdot (1 - LR(R_{eq}(t), t)), \quad (4.6)$$

where $R_{bond}(t)$ is the bond return as in (4.6), and the loss rate will be defined through equity returns

$$LR(t) = \begin{cases} 0, & R_{eq}(t) \geq 0, \\ \frac{dr_{max}}{-0.3} \cdot R_{eq}(t) \cdot LGD, & 0 > R_{eq}(t) \geq -0.3, \\ dr_{max} \cdot LGD, & R_{eq}(t) < -0.3. \end{cases} \quad (4.7)$$

In (4.7) value -0.3 is the worst case estimate for an annual equity return after which even lower returns do not affect the loss rates. In this thesis we assumed

a fixed parameter for loss given default ($LGD = 0.57$), which is based on long term averages of recovery rates of corporate bonds given in *The Oxford Handbook of Quantitative Asset Management*. (Scherer and Kenneth, 2012) Maximum default rate is defined as $dr_{max} = 0.0158$, which is the highest annual default rate occurred for investment grade bonds. These factors together define the loss rate, which is in our model at highest $\max LR(t) = 0.009$. Even at the highest levels the loss rate remains low, and for investment grade bonds it could be ignored. For higher risk speculative grade bonds these figures are much higher, and if such an asset class was incorporated in the model, loss rate modelling would become a much more important topic.

Even though loss rates are a matter of fact for the safer investment grade bonds, they are perhaps not the most important property to look at, because the loss rate has been modest even at its highest historical levels. If we were to incorporate the even more risky debt as an investment alternative, modelling the loss rate would become a more important issue, because the default rates for speculative grade bonds are significantly higher, as expected.

4.1.5 Property index

In this thesis we calculate the total return index for the real estate index applying the approximation of Koivu et al. (2005). Possibilities for modelling the real estate returns are numerous. In a model developed specifically for Finnish pension funds the real estate returns are modelled by two factors: property price and rental yield. (Koivu et al., 2005)

The total return of property is given by

$$R_{re,t} = \frac{P_t}{P_{t-1}} + \frac{1}{2} \left(\frac{R_{t-1}}{P_{t-1}} + \frac{R_t}{P_t} \right) - 0.03. \quad (4.8)$$

In the equation (4.8) the total return of real estate index at time t is $R_{re,t}$. P_t is the property price at time t , R_t is the rental return at time t , and the maintenance costs are assumed to be 3%. The total return of real estate investments consists of two stochastic variables, and the data for them is available from statistics of Finland (Official Statistics of Finland (OSF), 2016).

The approximation in (4.8) gives us the total return index, and the distribution of quarterly compound returns. The modelling of total return index for real estate investments follows the approach of equity investments as seen in equation (4.1), and as suggested in Meucci (2009).

Young and Graff (1995) suggest that real estate return distributions are non-normal. This makes our model assumption of moment matching seem

more suited than for example simulating from conditional (fitted) distributions. (Meucci, 2010)

As seen in figure 4.4, the real estate index has numerous sub-zero observations. Most of these negative real estate returns are from the recession of 1990s in Finland. Estimating the moments of distributions from shorter time frame would lead omitting this downturn, and potentially underestimate the risk related to investing in property in Finland.

4.1.6 Matching moments

The moments taken into account when using the moment matching algorithm are mean, standard deviation, skewness and kurtosis. Additionally to these variable specific moments, correlation is taken account as a co-moment of the variables.

The theoretical background for moment matching was described in section 3.2.2. Now we need to establish our own estimates for the model parameters. This is important especially for the mean reversion levels of interest rate and investment grade spread. We keep the other parameters unchanged from the original article by Høyland and Wallace (2001b), but redefine the mean reversion levels. For government bonds we take the current level as the mean reversion level. For investment grade spread we define the mean reversion level as the long term average of the spread. Also the mean reversion level for spread is close to current level.

Arguably other choices for mean reversion levels could be justified. An obvious way to define the mean reversion levels would be basing it on an expert opinion. An expert opinion would allow exploiting more sophisticated view on the state of the economy. An expert opinion could be used for the expected returns, replacing the estimated means of distributions.

4.1.7 Annualising moments

When we are estimating the model from quarterly data, but the natural timeframe for calculations is annual, an obvious issue is annualising the parameters. Best practice would be to use annual observations, but this is not possible due to short histories of studied indices. Estimating distribution moments from less than 20 data points would definitely be not sufficient. This is the reason why we resort to quarterly data, and then annualising the moments.

For mean and variance, or standard deviation, this is a relatively straightforward process assuming the variables are independent and identically distributed. For the higher moments, however, the issue is more complex. We

Table 4.3: Parameters used for moment matching algorithm in this thesis. CL is the current level on 29/12/15, RP the risk premium, VC the volatility clumping parameter. α is the mean reversion level, β is the mean reversion factor, σ is the volatility and γ volatility dampening.

Symbol	Equity	Money market	Gov. yield	Spread	Real estate
CL	-	-0.18	0.31	1.30	-
RP	0.30	-	-	-	0.30
VC	0.30	-	-	-	0.30
α	-	0.06	0.31	0.59	-
β	-	-0.64	-0.42	-0.51	-
σ	-	0.45	0.48	0.70	-
γ	-	0.5	0.50	0.50	-

Table 4.4: Annualised first six moments of distributions of studied market invariants. Quarterly moments before projection can be seen in table A.1.

Moment	Equity	Money market	Gov. yield	Spread	Real estate
Mean	0.08	-0.21	-0.20	0.07	0.08
Std	0.24	0.89	0.82	0.74	0.06
Skew	-0.53	-0.54	0.07	-0.06	-0.19
Kurt	3.47	3.85	2.88	4.45	3.39
5th mom	-5.32	-6.70	0.62	-1.34	-1.69
6th mom	23.94	31.63	13.25	37.23	20.56

do not delve deep into theory behind scaling these standardised moments, but use the method of Meucci (2010). This methodology is in line with the specifications we have made concerning the studied market invariants.

Moreover, Høyland and Wallace (2001b) specify the distributions for interest rates directly for the given interest rate level. This specification would not be compatible with our need to project the moments from quarterly to yearly time frame.

For correlations projecting for a longer horizon does not cause similar difficulties. However, in practice, especially in the shorter horizon, some issues may arise.

4.1.8 Parameters for short-rate model

The parameters for the short rate model are estimated for the three interest rate factors we employ in this thesis. Estimated parameters and the estimation method are described in this section.

Table 4.5: Correlations of market invariants. Based on quarterly data since 29/03/99.

	Equity	Money market	Gov. yield	Spread	Real estate
Equity	1.00	0.16	0.22	-0.78	0.40
Money market	0.16	1.00	0.78	-0.03	0.39
Gov. yield	0.22	0.78	1.00	-0.17	0.41
Spread	-0.78	-0.03	-0.17	1.00	-0.57
Real estate	0.40	0.39	0.41	-0.57	1.00

The GMM

We use the generalised method of moments (GMM) for estimating the parameters for our interest rate models. The GMM is not the only method available for short-rate model parameter estimation, but it is widely used in this context and has a solid theoretical background. (Hansen, 1982; Chan et al., 1992)

The estimation process has, however, some complications: it fails to give reasonable parameter estimates for the government bond yield, and for the money market yield. One, and probably the most important reason for this is that the historical time series for which we are trying to fit the model, nor the government bond yield or money market yield has not been stationary, but it has drifted towards zero the entire period. Of course such stationary periods exist in the history, but selecting such a distant period for estimating the model parameters today would seem unreasonable.

The reason the models are unable to be calibrated is that the short-rate models are by construct based on an average equilibrium level. In recent history such equilibrium levels have not existed, and the model calibration would lead to heteroscedastic error terms. For a downwards drifting time series the model would have underestimated the variable until crossing the equilibrium level, after which it would have overestimated it. Hence, the model calibration is not possible for European money market yields or government bond yields.

In contrast, for the spread of European investment grade bonds the model does work and produce reasonable estimates. Even by looking at the time series of the credit spread it seems justifiable to assume a model with a constant equilibrium parameter. The model runs for estimating the parameters for the spread model can be seen in table B.1. The selected parameters for the spread model are shown in table 4.3.

Comparison of different estimation trials can be seen in table B.1. The model is selected due to its high p-value, which means it is very close to

the unrestricted model counterpart, where the parameter γ is not prefixed. Moreover, the model parameters seem reasonable, and other models have rather low p-values, which suggests they are not properly set. The chosen time interval changes the estimation outcome drastically.

Caused by the downward trend in the European government bond yields during recent history we were unable to fit a short-rate model, and need to consider other approaches.

The most obvious way would be to employ expert opinion for setting the model parameters. Practitioners in the industry tend to have their own views on future levels of interest rates, and such could be used directly instead of estimating the model parameters.

For this thesis, however, such expert opinions are not available, hence we use another approach. Also, setting the volatility level and the speed of correction towards the long term equilibrium level would be very difficult to base on experts judgement.

We consider now the historical interest rates from Japan as an alternative for estimating the parameters for our model. Reasoning for this choice goes that as the Japanese economy has already experienced a rather long period of extremely low interest rates, maybe this could be used as a proxy for the future development in Europe. Comparisons of historical yields of European and Japanese bonds and money markets can be seen in table A.5.

Unfortunately, even the Japanese rates have been extremely low past decade, they have not been below zero until very recently. This makes estimating the models even more difficult than before, because the most common assumption of interest rate modelling, namely flooring rates at zero level, is no longer a valid assumption.

Comparison of the historical time series of European and Japanese yields are shown in figures A.2b and A.2a. The distributions of rate changes of respective time series are shown in figures A.3 and A.3.

The parameters for the estimated models are seen in table 4.3. The model for the money market rate is set according to the model shown in table B.2. The model for money market is estimated from semiannual data, and it has a relatively high p-value, and parameter reasonable parameter values. For money market we use the α given in B.2 as the mean reversion level. Corresponding runs for the government bond yield are shown in B.1. For the 5 year bond yield we take again a model with high p-value, but we replace the parameter α with the current level of European rate.

For spread the model estimated from semiannual data was selected, because it had very high p-value, and parameters close to the unrestricted model. The mean reversion level for spread was set according to the α value indicated by the semiannual model. Different spread model parameters are

shown in table B.1.


4.2 Generated scenario trees

Scenario trees generated according to the specification in the previous section are shown in figure 4.5. The planning horizon is now two time steps, and the branching structure of the tree is set to $\{6,6\}$. The scenario tree definition is shown in table form in C.2, where we have included the asset returns, probabilities, and node indices.

Qualitatively the generated trees seem reasonable. They all include scenarios of high and low returns, and their shapes resemble the original parameters set in the estimation process. Fixed income asset classes are much less risky, than equities or real estates. These scenario trees are generated by the moment matching algorithm, using MATLAB *fmincon* optimisation routine. The optimisation algorithm used was the interior-point algorithm.

Qualitatively the worst case scenarios in equities and real estates seem justified. For equities this means losing approximately 60 % of the initial value until the end of the planning horizon. Such drops have occurred in the recent history, as we can see from figure 4.2. Similarly, the real estate prices lose almost 15 % of their initial value in the horizon of three planning periods.

The scenarios for money market investments are the least risky ones of all asset classes. This is no surprise, and their historical time series in figure 4.2 does not include any drops. Indeed, in the recent history there exists only one quarterly period, when the return generated by money market investments has fallen below zero. Now as the current yield level is below zero, the return formula (4.3) used for money market investments generates sub-zero returns even from the beginning of the planning horizon.

Figure 4.5: Scenario trees for all asset classes with branching structure $\{6,6\}$.

4.3 Optimised allocations

Next we show the optimisation results on the generates scenarios discussed in the previous section and shown in figure 4.5.

In the optimisation runs we have varied the model parameters defining

the level of risk aversion, namely parameters a and b in equation (3.24). This gives us slightly different objective functions (3.25), which are then fed to the optimiser.

Optimisation routine is carried out in MATLAB using function *fmincon*, which is an optimiser for non-linear constrained optimisation problems. The algorithm the solver uses is the interior-point algorithm, which is the same algorithms many other optimisers suited for similar problems use.

Each optimisation run returns a slightly different set of optimal decisions based on the risk aversiveness of the objective function. We compare the results by fixing the risk measure to a shortfall cost function with parameters: $\{a, b\} = \{1, 10\}$. Fixing neutral parameters for the risk function allows us to plot efficient frontiers according to these parameter choices.

Utility is measured by the expected return at the end of planning horizon, and it is penalised by the shortfall of a return target. Three different target returns were utilised. The target returns were in annual terms: 4%, 2%, 0%. These translate to cumulative target returns according to compounding $(1 + r)^t - 1$. The efficient frontiers for each target returns are shown in figure 4.6.

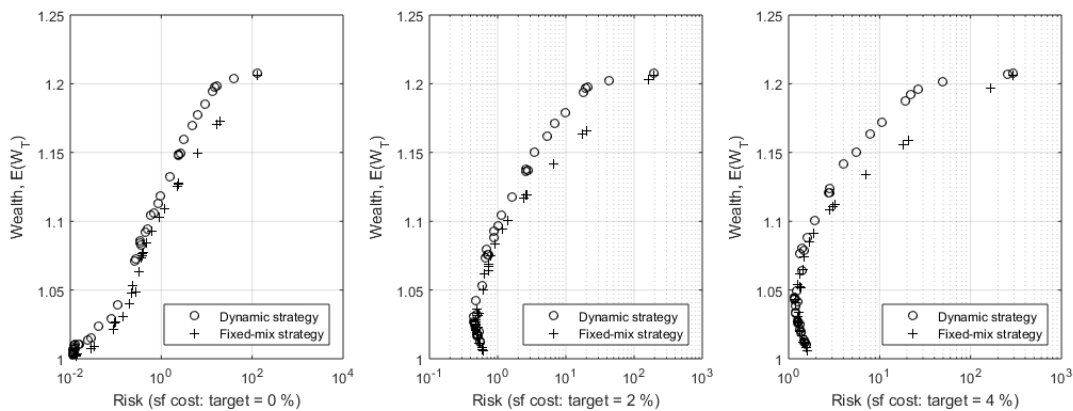


Figure 4.6: Efficient frontiers for all target returns in dynamic strategy and fixed-mix strategy in scenario tree of branching structure $\{6, 6\}$. Risk on a log-scale.

In figure 4.6 we have plotted the risk in logarithmic scale, because the shortfall cost rises sharply when we move to riskier portfolios and higher expected returns. All frontiers for different target return sets are concave, even though the leftmost panel for target return 0% annually does seem convex. This illusion is due to the logarithmic horizontal axis, which is distorted.

We can observe that when target return is rising, the low-risk end of the efficient frontier moves right, while the highest return points remain still. This is an obvious effect, because when the target return rises, it becomes increasingly difficult to obtain it, especially in the low-risk portfolios. The reason why very low risk portfolios are still obtained in the higher target return lies in the risk aversion: if the risk aversion is extreme, especially towards extreme events, the penalty for missing the target slightly matters much less than being very far off.

The resulting asset allocation in the first period of dynamic strategy is shown in figure 4.7. In figure 4.7 we have split the allocations according to the risk aversion parameter b , even though in the efficient frontiers in 4.6 moving along the frontier results from altering both parameters, a and b . In figure 4.7 only the parameter a changes along the expected wealth at the end of planning horizon. Having the effect of a and b plotted in the same figure would lead to slightly unintuitive results, and we would have sudden shifts in the optimal portfolio composition, because very different portfolios may generate expected returns that are very close to each other.

The resulting allocations given by the optimisation procedure for the first period allocation decision are rather intuitive, and in line what one would expect from such an experiment. For higher risk aversion we obtain low-risk portfolios with high weights in money market investments. When the risk aversion changes, we shift from lower risk to slightly higher risk portfolios, and finally the portfolio consists solely of equities, which is the asset class inducing highest expected return, but is also the riskiest one.

In portfolios aiming at higher expected returns the problem simplifies mostly to choosing the allocation between real estate and equities. When the target return is increased, we see smaller and smaller allocations to the lowest risk asset classes, especially money market investments. This is an expected result: when we penalise for missing target return annually 4%, it is obvious that high weights to a low-risk asset that generates almost certainly returns lower than the target, are very low. Reasons why the lowest risk asset is still included with small weights even in the case of high target return, is the benefit of diversification, but also the aversion to extreme events. Money market investments provide a safe investment opportunity, if we want to avoid very extreme shortfalls at all costs.

One slightly surprising observation we make is the absence of allocation to government bonds, one of the asset classes in the low-risk end of the spectrum. This is simply due to the scenarios we have generated for government bonds, and their very low expected return. From table D.1 we can see that in the two-period model the expected return for government bonds our generated scenarios is actually below zero, which makes it easy to understand why the

first period allocations do not suggest government bonds in almost any case. It is only included in the portfolios with low target return, high risk aversion, but relatively low risk aversion towards extreme shortfalls.

Other three asset classes are present in all runs. Allocation to real estates seems to be shrinking when the parameter b controlling risk aversion towards extreme events is increasing. The allocation to real estate investments is surprisingly large, and due to their illiquid nature it could be worth considering to imposing a constraint on their maximum weight, or making their transaction costs very high. Another way to adjust the very high proportion of real estate investments would be adjusting the annual maintenance costs in (4.8).

In figure 4.8 we have the expected wealth at the end of the planning horizon as a function of risk aversion parameter a . Each line represents a combination of one of the target returns, and a risk aversion parameter b .

The interesting shifts in optimal allocation, and therefore in the expected wealths, seem to happen in the region of $a = [0, 1]$, after which increasing risk aversion factor a does not change the expected wealth at the end of planning horizon. Lower values in factor b result to faster shifts from very high risk portfolios to low risk portfolios. This is due to the fact that when b is higher, we want to avoid the extreme allocation to equities for longer, where a low value of b allows the optimiser to shift from low risk portfolios to extreme allocation to equities with smaller changes in parameter a .

One interesting observation we make from figure 4.8 is that for target returns of 0% all the lines converge to the same wealth level, when the risk aversion coefficient a is increased. For target returns of 0% the factor b does not change the final value where lines converge to.

For higher annual target returns the results are different. Expected final wealth levels where the lines converge to when a increases seem to differ according to the risk aversion coefficient b . For lower b the converging level of expected wealth is significantly higher when a increases, than when b is very high. This result can be observed from the allocations shown in figure 4.7. Even the low-risk portfolios for high target returns contain much higher proportions of riskier assets, than in the case of lower target returns.

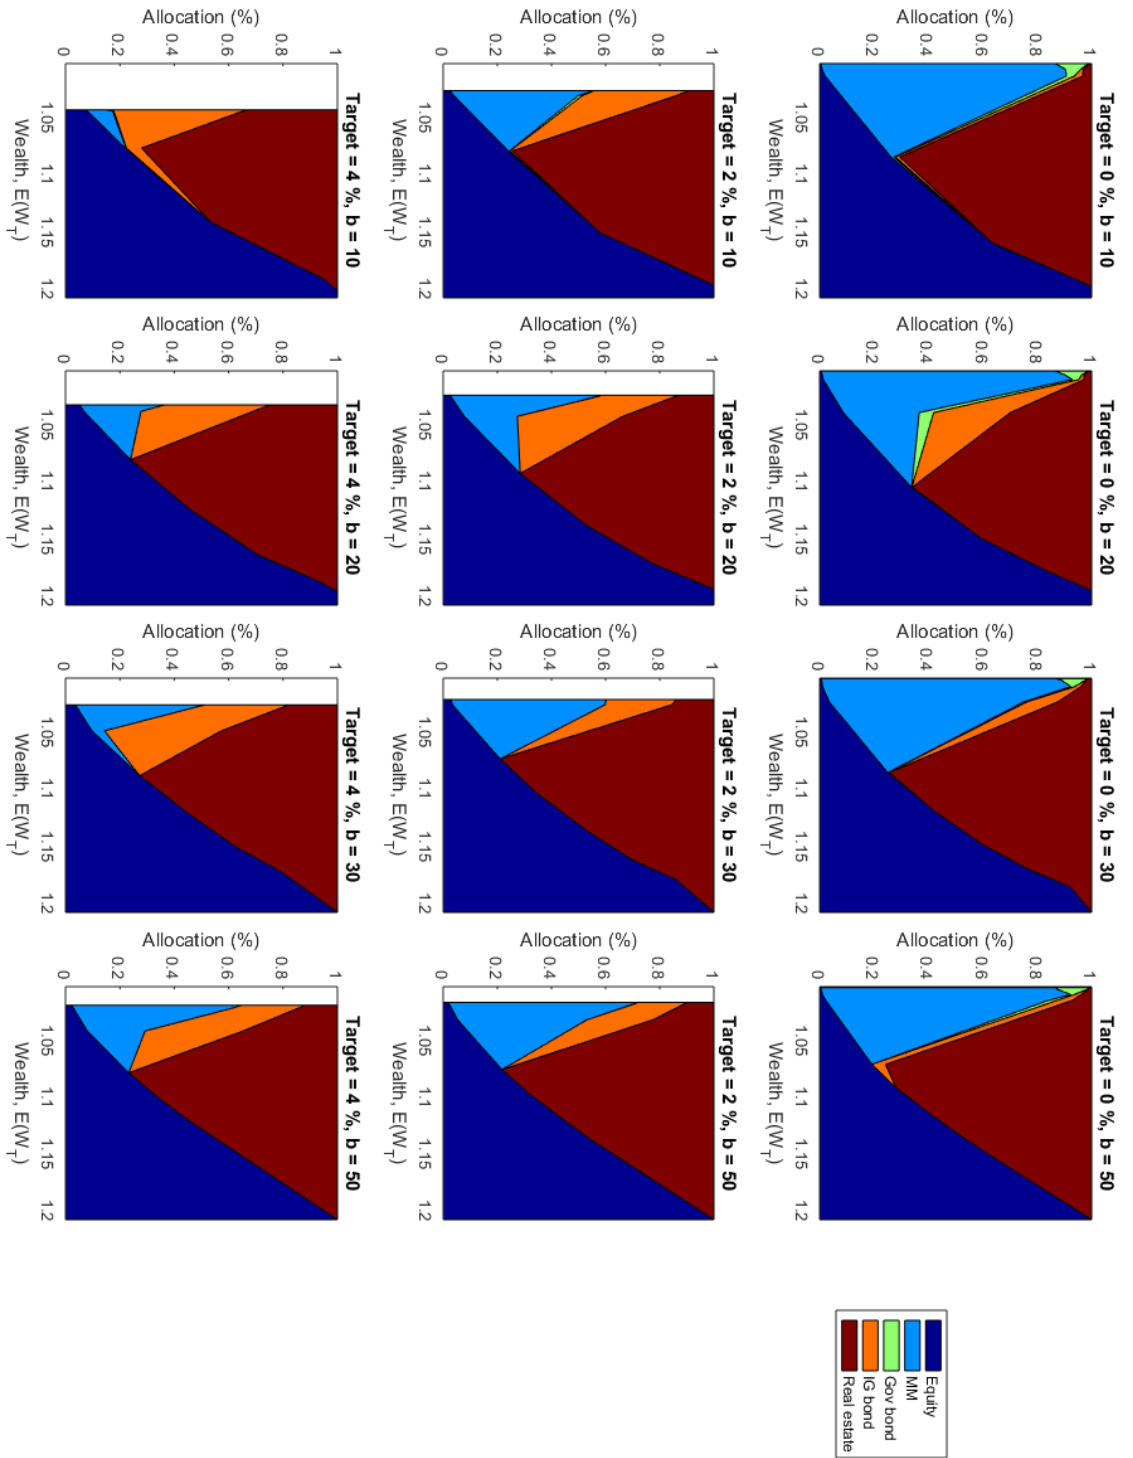


Figure 4.7: Different allocations among the efficient frontier for different parameter values for the dynamic strategy. For each panel the target return and risk aversion b are fixed, a changes and results in different values in expected wealth.




Figure 4.8: Wealth as a function of risk aversion coefficient a. Functions are shown for different target returns and different values of risk aversion coefficient b.

4.3.1 Dynamic second stage decisions

To get an idea of the dynamic behaviour of the model we chose two risk aversion levels, namely two combinations of risk aversion parameters: $\{a_1, b_1\} = \{10^{-2}, 10\}$ and $\{a_2, b_2\} = \{10^{-2}, 50\}$. First combination leads to very high risk portfolios, and the second somewhat lower risk portfolios. The parameter a , which controls the general risk aversion level, is kept constant, while b , controlling the aversion towards extreme events, is altered.

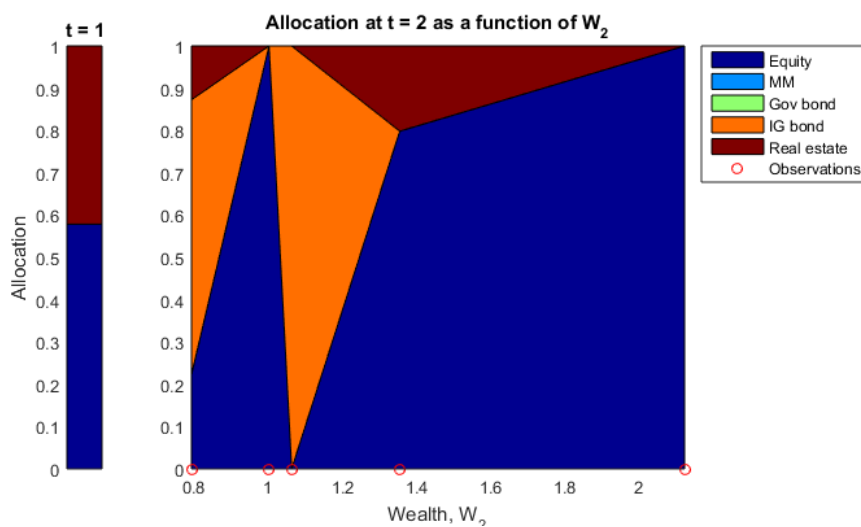
Corresponding scenario tree for the risk aversion parameter values in the first figure 4.9a are shown in figure C.1. Data in figure 4.9a can be seen in numerical form in table C.1.

The observed decisions in the nodes of the tree given by the optimisation are indicated with red circles along the horizontal axis in figure 4.9. Trees are sparse, and some of the observed wealth levels are very far apart. On the other hand, we can observe very sudden shifts in the optimal allocation when wealth changes. In some cases for similar level of wealth the optimal decision given by optimisation are very different. This is due to differences in the scenarios thereafter.


To draw conclusions about the dynamic nature of the decision making, it would be beneficial to implement larger scenario trees. Moreover, the allocation given in between the observations is only suggestive, and should not be taken as a result. With larger scenario trees we could be able to shorten the gaps between observations. This comes, however, with a cost, because the complexity and calculation times increase.

Even though the trees are sparse, we may be able to draw one conclusion from resulting second stage allocations: when the obtained wealth level is very high, it is beneficial to adjust the allocation to a riskier one. This may be due to the fact that with higher wealth we are further away from the penalty threshold, under which the penalising part of the objective function would return non-zero costs. In other words, when the wealth is high, we have more means for risk taking, and seeking even higher returns. This is clearly visible in figure 4.9b, where the allocation is riskier among higher wealth.

The behaviour of shifting to a riskier allocation after favourable outcome from the first period is in line with the constant-proportion strategy mentioned earlier in section 2.3. In constant-proportion strategy the investor adjust the amount of a risky asset class in the portfolio according to the value of a "cushion", which is the difference between the value of the portfolio, and a floor under which investor does not wish to fall. Moreover, as demonstrated earlier in figure 3.3, the risk aversion decreases to zero when the target return is obtained. This decreasing risk aversion may explain the



(a) Risk aversion parameters: $\{a_1, b_1\} = \{10^{-2}, 10\}$



(b) Risk aversion parameters: $\{a_2, b_2\} = \{1, 10\}$

Figure 4.9: Second stage allocation decisions for two different risk aversion levels in a two-stage tree with a branching structure $\{6,6\}$. Observations marked with red circles along the horizontal axis. First stage decision shown in far left of the figure.

shifts in high wealth end of figure 4.9b.

Still, as noted before, strong conclusions should be avoided from these results, because the scenario trees are sparse, and the resulting second stage allocations seem to differ significantly according to the risk aversion coefficients.

4.3.2 Comparison with fixed-mix strategy

The dynamic strategy yields seemingly reasonable allocation results. Next we will discuss what if we constraint the strategy to a fixed mix strategy, where the proportions of each asset class is rebalanced to the original weight in each period.

The optimisation is run similarly as in the previous section for the dynamic strategies. The only difference is now that the allocation is kept constant in each period. This limits the decision maker, and does not allow for corrective actions during the planning period.

The resulting efficient frontiers for fixed-mix strategies in the stochastic scenario optimisation are shown in figure 4.6 along with the dynamic frontiers. The result seems clear: dynamic strategy performs better than the fixed-mix strategy. Performing better means that the efficient frontier of dynamic strategy is above the frontier of fixed-mix strategy. This means, in other words that in the dynamic strategy same expected wealth can be obtained by lower shortfall cost.

Difference is clear in the midway along the efficient frontier, but in both ends of the frontiers they seem to converge. This is due to the fact that for very high risk aversion and low risk portfolio there is not that much space for adjustment: the allocation is almost entirely to the lowest risk asset, namely money market. Similar logic holds for the other end of the frontiers: when risk aversion is low and investor is seeking for very high expected returns, the allocation is solely on equities, which is the riskiest asset class.

Allocations given by the optimisation are shown in figure C.2. As in the corresponding figure of dynamic first period allocations, the frontiers have been split here according to the target return, and the risk aversion coefficient b .

The first thing we notice in figure C.2 is that the optimisation suggests very high weights for real estate investments in the midway of the frontier. This is a consistent result for all target returns, and for all risk aversion coefficients b .

Allocation to government bonds is even smaller than in the dynamic first stage decision, and it is not visible to naked eye in figure C.2. When comparing figures 4.7 and C.2 it is very difficult to say any definitive differences between the two strategies. It seems that in the dynamic strategy it is possible to adjust the portfolio to slightly riskier composition, because the possibility for adjustment exists.

Allocations to equities and on the other hand money markets remain relatively intact compared to the first period allocation in the dynamic strategy as seen in figure 4.7. This is in line what we observed already from the effi-

cient frontiers: they converge at both ends of the spectrum. When the goal is very high return, the only option is to allocate all to the riskiest asset, and when risk aversion is very high, the allocation will be mainly to money market investments, in both strategies.

Although the dynamic strategy seems superior when compared to the fixed-mix strategy, one important consideration is missing from the model used: the costs that occur when allocation is changed. These can be due to transaction costs, or the asset class may be highly illiquid, which would make adjusting its allocation nearly impossible. It is probable that the difference in performance between dynamic and fixed-mix strategies would be smaller if we accounted properly for the costs occurring from adjusting the allocation.

On the other hand, if we split the planning period to shorter intervals, in theory we should see even better performance from the dynamic strategy, because more opportunities for adjusting the portfolio exist.

4.3.3 Comparison with mean-variance optimisation

So far we have conducted optimisations and comparisons in the framework of stochastic optimisation and generated scenario trees. Now we change the framework to the classical mean-variance optimisation, and study how efficient the generated portfolios from stochastic programming framework are in this more traditional model.


To compare the results of both dynamic strategy and the fixed-mix strategy we derived expected returns and standard deviations for the two-period planning horizon, which are required in order to run a meaningful optimisation. Means and standard deviations are derived by calculating these values from the scenario trees shown in figure 4.5 at the end of the planning horizon.

As a correlation matrix for this purpose we use the correlation matrix shown in table A.4, which is calculated from the quarterly linear returns of observed historical total returns of studied asset classes. Estimating parameters from historical data is a typical approach to get the parameters for mean-variance optimisation.

For expected returns and standard deviations we can not use historical data, because that would require us to make additional assumptions on how to make projections of quarterly returns to annual, and to two-year scale. For correlations this should not be a problem, because they should remain constant for different time scales. If we projected a covariance matrix to another horizon according to formulas given in Meucci (2009), the covariance indeed changes for different horizons. But the normalised measure of dependence, the correlation coefficients, remain constant.

Mean-variance optimisation is based on the theory developed by Harry

Markowitz, and the precise formulation of the optimisation problem is given in equation (2.1). In practice, the optimisation is carried out by using portfolio optimisation tools in MATLAB. Results of this comparison can be seen in figure 4.10. The efficient frontier given by the mean variance optimisation is the blue line on the left panel. Individual asset classes, of which each portfolio consists of, are plotted by blue stars according to their means and standard deviations. Corresponding data is in table form in D.1.



(a) Mean variance efficient frontier. (b) Mean variance optimised allocation

Figure 4.10: Mean variance efficient frontier and optimised allocation.

Results from earlier experiments, first period allocations in the dynamic strategy, and the fixed allocations in the fixed-mix strategy, are included in the figure 4.10. It seems that the fixed-mix strategy gives very similar results that the mean-variance does, and the fixed-mix portfolios are very close to the efficient frontier.

Also the first period allocations of the dynamic strategy are very close to the efficient frontier. This is of course no surprise, because the allocations consist of mainly assets found along the efficient frontier.

For individual asset classes it seems clear from figure 4.10 why certain assets make a high proportion of the portfolio with certain risk levels. Government bonds are almost always omitted, which is not surprising given their sub-zero expected return, but higher standard deviation than for money markets.

From the allocations given by the mean-variance optimisation seen in figure 4.10 we can observe that real estate investments take smaller stake of the portfolio in the midway along the efficient frontier. The explanation is that in the sense of mean-variance optimisation investment grade bonds are a more efficient asset class, as they are closer to the frontier.

In this comparison the results from fixed-mix strategy and mean-variance

optimisation are quite similar. The key differences in the fixed-mix strategy, and mean-variance strategy, are related to the risk measure. In stochastic programming our risk measure is asymmetric: we penalise only the shortfalls, but not high returns. Standard deviation as a risk measure is, however, symmetric. In theory it penalises equally for big shifts both up and down, even though outcomes with extremely good outcomes would be preferred.

Other factor explaining the differences seen in these two methods lie in the differences seen in correlations of asset class returns. In mean-variance framework we used the correlations calculated from historical returns, where as the tree has its own correlation structure for the asset returns. These differences can be observed from tables 4.5 and A.2 for the moments. The correlations of historical asset returns can be seen in A.4 and their counterparts calculated from the tree D.2.

Most significantly the correlations seem to differ for money market returns. The reason is that their return is calculated according to the geometric average of current yield as seen in equation (4.3). The return of money market investments goes to opposite direction than government bonds, because the modelling of money market investments ignores the effect of bond value changing when rates change. Modelling money market returns this way the correlations are more or less the opposite sign than historical correlations, and this is consistent across the correlations between money markets and all other asset classes.

Apart from money market correlations, the correlations among other asset class returns in the scenario tree seem to be in line with the ones calculated from quarterly history. They are not exactly the same, which is not surprising, because our model has made numerous assumptions on how to calculate the asset returns, some of which may need revision.

Chapter 5

Evaluation and discussion

5.1 Stability of the optimal solution

To study the stability of our model we performed two experiments. In the first experiment we ran the scenario tree estimation procedure for four different branching structures. The highest branching tried was 17 branches per node, which makes a total of 289 scenarios in a two-period tree. No larger trees were tried, because solving trees larger than this became a very slow procedure.

We ran the test for two sets of risk aversion parameters that were from the midway of the efficient frontier. The result can be seen in figure 5.1. Stabilising of the final value of the objective function and utility measures does not seem evident, but the decision variables are relatively stable for all runs.

The two distinctive parts of the objective function are plotted separately, because their relative importance is in a very different scale. Changes in values of objective functions between the two larger branching structures are much smaller than for the first ones. This gives an idea that maybe the problem is stabilising.

To establish a good estimate on the model stability and what would be a minimum size for a tree, larger trees should be tried. In Kaut (2003) and Hilli et al. (2007) the authors suggest that stabilising occurs for trees larger than 1000 scenarios in the final node. Due to the model performance we did not run such tests.

Secondly, we studied the stability of the decision variables, namely the allocations, in small scenario trees. Results for decision variable stability can be seen in tables 5.1 and 5.2.

It is seen that for high risk tolerance the stability of the decision variable becomes an issue. For low risk portfolios the deviations in optimal values of




Figure 5.1: Stability of the utility measures and first stage decision variables. Low risk aversion, high risk portfolio. $\{a, b, r_{target}\} = \{10^{-2}, 30, 2\%\}$ Compare to table 5.2.

Table 5.1: Stability of the first stage decision variables, high risk aversion, low risk portfolio. $\{a, b\} = \{1, 30\}$, 25 runs.

Asset Class	Equity	MM	Gov bond	IG bond	Real estate
Average	3.3 %	63.5 %	0.1 %	22.8 %	10.3 %
Standard deviation	1.0 %	6.3 %	0.2 %	7.2 %	1.9 %
Smallest value	1.3 %	49.4 %	0.0 %	6.5 %	6.5 %
Largest value	5.2 %	75.8 %	0.9 %	37.9 %	13.9 %

decision variables remain stable.

Finally we plotted the moments of the equity returns in the scenario tree used earlier in the main studies, and compared them to the moments the moment matching algorithm is matching them to. Moments and target moments can be seen in figure 5.2. Here we can see all moments obtaining a relatively good fit, but not perfect. It is difficult to say the reason for these matches being only approximate, but it may be an accuracy in the solver, which we could overcome by adjusting the tolerance level.

Table 5.2: Stability of the first stage decision variables, low risk aversion, high risk portfolio. $\{a, b\} = \{10^{-2}, 30\}$, 25 runs.

Asset Class	Equity	MM	Gov bond	IG bond	Real estate
Average	14.3 %	13.0 %	0.1 %	30.1 %	42.5 %
Standard deviation	7.1 %	17.6 %	0.2 %	25.3 %	24.9 %
Smallest value	6.1 %	0.0 %	0.0 %	0.0 %	11.4 %
Largest value	30.1 %	58.1 %	0.7 %	69.5 %	79.1 %




Figure 5.2: Moment values of equity return trees for two a tree with branching $\{6, 6\}$.

Still, in figure 5.2 even the higher moments obtain a very close match to their target values. Here it could be worth considering a more sophisticated relative weighting in the matching procedure, to primarily match the lower two moments and then the higher ones.

5.2 Potential issues in model specification

In this section we discuss some of the choices concerning the model specification, and what other choices could have been possible. There are arguably numerous aspects that could have been made otherwise.

First of all, the selection of scenario generating methodology. We chose moment matching as the scenario tree generating algorithm, because it had been previously used successfully in a similar problem. Moreover, the moment matching algorithm had the property that it may be able to accommodate more information in smaller trees, than for example simulation based procedures. It was indeed suggested in Høyland and Wallace (2001b) that for

a five dimensional problem like the problem we studied in this thesis, the smallest number of branches would be at around six.

On the other hand, the model suggested by Hilli et al. (2007) would have been utilised in a similar problem. However, the model Hilli et al. (2007) proposed was much more complex. Their model would have required the use of more sophisticated optimisation tools than MATLAB, and it remains unclear how big scenario trees would have been possible to solve with the current solution.

Of the individual asset class returns money market investments seem to be the one with most issues concerning the modelling. Due to the fact that the return correlations differ significantly from the ones observed in the historical data, it is easy to draw a conclusion that other methods could have been applied.

Modelling of the investment grade bonds is also topic that contains many aspects that could be criticised. First of all, including the loss rate as we did, was not suggested anywhere in the literature we were able to find. Therefore it is fair to say that it complicates the model unnecessarily. But if we were to include a riskier debt asset class, namely the speculative grade bonds to this model, it would be much more important to include the loss rates, because they are significantly higher for speculative grade bonds.

Real estate investments were modelled according to their total returns, and the moments of total return index were based on an index we composed ourselves from historical data of the Finnish real estate market. This makes it a country specific, and if we were to apply it on another geographical area, the local real estate market should be modelled separately. In Meucci (2009) the real estate investments are suggested to be modelled according to distributions of their compounded total returns, which is the strongest argument for our model specification. In Hilli et al. (2007) the real estate investments are modelled, however, with two distinct factors, which we have blended together in our model. Hence it is fair to say that our modelling of real estates has room for improvement.

In more general terms, the modelling decisions made concerning the interest bearing asset classes could be reconsidered. Asset classes are assumed to be of fixed duration, about five years, and their return modelling is based on solely one fixed point in the interest rate curve. This omits the changes in value caused by the shape of the yield curve, and changes in the shape of the curve. Having only two interest rates of different maturities modelled, we are not able to construct an entire yield curve. Constructing an approximation for the entire yield curve across all maturities would enable us to model arbitrary cash flows according to their present values. As an example, liability modelling could be based on estimated cash flows, which could

then be market valued in the scenarios generated. This would complicate the model further, and the number of required scenarios to accommodate the problem would increase.

For modelling the interest rates we employed the CIR model, which is a short-rate model based on one random factor. Other alternatives for selected model existed, but because the CIR model has properties we valued, we chose it. As shown before, estimating parameters for a short-rate model was challenging, because the recent history of European interest rates is not consistent with the core assumption of the short rate models: the mean reversion. We estimated the model parameters from Japanese data, making an assumption that the past history of low rates seen in Japan would occur in Europe too. If this assumption was altered, it is very likely that the results we obtained would not hold.

Dependencies between asset class returns were modelled according to linear correlations. This is a common assumption in financial models. More sophisticated possibilities exist, and copula functions could prove to be an alternative solution for modelling the dependencies. The problem that is embedded in linear correlations is that they do not adjust in time, nor along the probability of the event. In the events of extreme events in the financial markets it is typical that the correlations tend to rise. This means that when the equity markets sink unexpectedly, the corresponding bond market sinks as well, even though under normal market conditions their correlation would be low. Copula functions would complicate the model even further, hence it would be a trade-off to consider the easier tractability, to better modelling.

5.3 Solver selection and computational issues

The model is implemented entirely in MATLAB programming environment without additional solver capacity available. To improve the model performance using other solvers could be worth considering.

In the literature similar models have been solved using more sophisticated optimisation tools. In Hilli et al. (2007) the authors suggest using the AMPL modelling language and MOSEK solver. These modelling tools employ same interior-point algorithms as MATLAB, the tool we used in our work, so the algorithm does not provide an obvious reasons to change the modelling language.

Høyland and Wallace (2001a,b) do not mention the solver used. They do, however, mention the scenario generation being run on a Sun Ultra Sparc 1 machine. Hence, a modern PC should be able to cope with a similar problem if specified efficiently.

If we were to develop the model further, considering the performance issues could be one of the key issues to continue working with. One of the decisions made in the model building phase, that proved to be inefficient, was the use of matrix formulation in (3.6) to calculate the function values at the end of planning horizons. The formulation may be useful in calculating model values once, but it is not suited for iterative evaluation, which is present in all optimisation algorithms.

5.4 Further research topics

During the process of building the model and analysing results we have found several directions in which the development could be continued.

From the software and performance point of view, exploiting more powerful optimisation tools than MATLAB could be considered, as mentioned before. In the algorithm there remains parts that could be made more efficient, and then enable modelling larger problems.

Considering other possibilities for interest rate modelling would be one of the key aspects to consider when developing the modelling further. Assessing the problem of extremely low interest rates would be an interesting issue, and it would be beneficial to seek for further justification for selected model parameters.

Like the models by Hilli et al. (2007) and Høyland and Wallace (2001a), it would be interesting to implement regulatory solvency requirements in the model. The shortfall cost could be adjusted to be defined according to these regulatory limits set by the financial supervisory authorities. Conducting these studies with revised and updated model parameters would be of interest as well, because the market environment has been very different during the times these models were developed and proposed.

Extending the model to consider the regulatory solvency requirements would require us to consider modelling of market risk of liabilities. This would lead us back to considering better modelling of the interest rates, even the entire yield curve. Estimating the yield curve would enable modelling arbitrary future cashflows based on present valuation. A potentially suitable model for such parsimonious yield curve modelling could be the Nelson-Siegel model, where the interest rate term structure is modelled based on a small number of factors. (Nelson and Siegel, 1987)

One of important issues not addressed in this thesis is the transaction costs investors face when adjusting their portfolios. This remains as a further research topic. Especially for the illiquid assets, such as real estates, considering a costs for adjusting the portfolio would probably change the

results. Moreover, every investor has their current portfolio - taking into account the costs occurring from adjusting from the initial portfolio should change results as well. These considerations could narrow down the dominant allocations to real estates, which are perhaps a bit unrealistic for institutional investors.

Chapter 6

Conclusions

The starting point for this thesis was to develop a sound multi-period asset allocation model, potentially including liability modelling. While liability modelling proved to be too challenging at this point, we simplified the approach, and instead of liabilities, the risk was modelled according to an annual target return.

The multi-period model we developed is a stochastic programming model, which is based on scenario trees. In stochastic programming we do not have deterministic objective functions for the optimisers, but the optimisation is conducted in a probabilistic setting, which means that instead of optimising an outcome, we optimise the expected value of an outcome.

Scenario trees were generated to accommodate the properties of five distinct asset classes, each of which was modelled according to a market invariant, or a combination of multiple market invariants. Scenario tree generation procedure was based on a moment matching algorithm, where the properties of the scenario trees were fitted to pre-specified properties modelled by means of time series and econometrics. Econometric modelling of the interest rates had some complications, because the mean reverting models could not be fitted to time series with a clear trend. To overcome this issue we used the corresponding historical data from the Japanese interest rates, because the current extremely low interest rates have been the prevalent market condition for more than a decade in Japan.

This is, however, a significant factor reflected to the results. Because the interest rates were modelled according to very low rates, their return potential is also very low. This makes the other than interest bearing asset classes more attractive, because they provide better income for the investor.

One of the key hypotheses we had concerning the outcome of our model was that the dynamic asset allocation strategy would be dominant when compared to a fixed-mix strategy. In our study the scenario trees were two-

stage trees. The only difference between the fixed-mix strategy, and the dynamic strategy, was that in the dynamic strategy the investor is allowed to adjust their portfolio once during the planning period. In the fixed-mix strategy the allocation remains intact until the end of the horizon.

The results of this comparison were as expected: allowing the decision maker to adjust their decision should in principle benefit the outcome. However, when we studied more closely the second period decisions the investor should make, a bigger picture was left somewhat unclear. Because the scenario trees generated were sparse, we were unable to develop comprehensive decision rules for updating the portfolio based on the observed outcome.

The result we obtained from the second period allocation suggests that a behaviour like in the constant proportion portfolio strategy, may be justifiable. This means that adjusting the portfolio to a riskier composition after a favourable outcome seems reasonable, because the investor has a large safety margin to the penalty threshold, namely the return target. Strong conclusions of any sort should however be avoided, because the model specification had multiple attributes that could be specified differently.

A comparative study with the traditional mean-variance optimisation developed by Harry Markowitz, was conducted. The results of this study showed that the portfolios given by the stochastic programming, are indeed close to their counterparts generated by mean-variance optimisation. Even the portfolios for the first period in the dynamic model were close to the efficient portfolio. This result may suggest that if the investor is used to conducting similar optimisation in the mean-variance framework, the decision maker could understand the problem better through mean-variance parameters, even when the actual optimisation is conducted in the stochastic programming framework.

Intuitive results for constructing extremely high and low risk tolerance portfolios were confirmed: if the risk tolerance is high, the investor should invest most their funds in equities. In the opposite end of the risk spectrum the investor should invest in money market investments. Under the assumptions of low interest rates, government bonds seemed the worst asset class to invest in, because according to the model developed, they yield negative expected returns in the two-period planning horizon.

The stochastic programming model seemed to favour three asset classes: equities in the high risk end, money markets in the low risk end, and real estates in between. In lower end of the risk and return spectrum small allocations to investment grade bonds occur, and even minuscule allocations to government bonds emerge.

Differences compared to other strategies, namely the fixed-mix strategy, and the mean-variance optimal strategy seemed surprisingly small. This may

be, however, because only one adjustment to the portfolio in the planning period is allowed. More adjustments should mean better performance for the dynamic strategy compared to the fixed strategies.

Model performance was an issue that arose along the way, and even though it improved significantly during the development process, it remained an issue to be addressed. Similar models have previously been built in more powerful programming environments than MATLAB.

The stability of the model developed remained an open question, and it seemed likely that larger scenario trees would benefit the accuracy of the results. Still, the smaller scenario trees shown in this thesis could provide a basis for bracketing the investor preferences. As the computation times increase exponentially among the number of branches per node, it could be an idea worth considering, to establish the risk aversion level by using rough estimates given by small trees, and finding a more accurate solution for prefixed risk aversion parameters in a larger scenario tree.

Even though the results of this thesis may not be applicable as such, it serves as a basis for further developments. Because stochastic programming is a flexible methodology, the model presented could be further extended to an integrated model, where both assets and liabilities were modelled under the same framework.

Bibliography

- Black, F. and Litterman, R. Global portfolio optimization. *Financial Analysts Journal*, 48(5):28–43 (1992).
- Bloomberg Finance L.P. Bloomberg Terminal (2016). Accessed: 16.1.2016.
- Calfa, B. A., Agarwal, A., Grossmann, I. E., and Wassick, J. M. Data-driven multi-stage scenario tree generation via statistical property and distribution matching. *Computers & Chemical Engineering*, 68:7–23 (2014).
- Campbell, J. Y., Lo, A. W.-C., MacKinlay, A. C., et al. *The Econometrics of Financial Markets*, volume 2. Princeton University press, Princeton, NJ (1997).
- Chan, K. C., Karolyi, G. A., Longstaff, F. A., and Sanders, A. B. An empirical comparison of alternative models of the short-term interest rate. *The Journal of Finance*, 47(3):1209–1227 (1992).
- Cox, J. C., Ingersoll Jr, J. E., and Ross, S. A. A theory of the term structure of interest rates. *Econometrica: Journal of the Econometric Society*, 53(2):385–407 (1985).
- Dupačová, J. Multistage stochastic programs: The state-of-the-art and selected bibliography. *Kybernetika*, 31(2):151–174 (1995).
- Dupačová, J., Consigli, G., and Wallace, S. W. Scenarios for multi-stage stochastic programs. *Annals of Operations Research*, 100(1-4):25–53 (2000).
- Dupačová, J., Hurt, J., and Štěpán, J. *Stochastic Modeling in Economics and Finance*, volume 75 of *Applied Optimization*. Kluwer Academic Publishers (2003).
- European Central Bank. The definition of price stability (2016). Available at <https://www.ecb.europa.eu/mopo/strategy/pricestab/html/index.en.html>. Accessed: 30.3.2016.

- Hansen, L. P. Large sample properties of generalized method of moments estimators. *Econometrica: Journal of the Econometric Society*, 50(4):1029–1054 (1982).
- Hilli, P., Koivu, M., Pennanen, T., and Ranne, A. A stochastic programming model for asset liability management of a Finnish pension company. *Annals of Operations Research*, 152(1):115–139 (2007).
- Høyland, K. and Wallace, S. W. Analyzing legal regulations in the norwegian life insurance business using a multistage asset–liability management model. *European Journal of Operational Research*, 134(2):293–308 (2001a).
- Høyland, K. and Wallace, S. W. Generating scenario trees for multistage decision problems. *Management Science*, 47(2):295–307 (2001b).
- Kallberg, J. G. and Ziemba, W. T. Comparison of alternative utility functions in portfolio selection problems. *Management Science*, 29(11):1257–1276 (1983).
- Kaut, M. *Scenario tree generation for stochastic programming: Cases from finance*. Ph.D. thesis, Norges teknisk-naturvitenskapelige universitet (2003).
- Kaut, M. and Wallace, S. W. Evaluation of scenario-generation methods for stochastic programming. Number 14 in Stochastic Programming E-Print Series. Humboldt-Universität zu Berlin, Mathematisch-Naturwissenschaftliche Fakultät II, Institut für Mathematik (2003).
- Koivu, M., Pennanen, T., and Ranne, A. Modeling assets and liabilities of a finnish pension insurance company: a VEqC approach. *Scandinavian Actuarial Journal*, 2005(1):46–76 (2005).
- Luenberger, D. G. *Investment Science*. Oxford University Press, New York (1998).
- Markowitz, H. Portfolio selection. *The Journal of Finance*, 7(1):77–91 (1952).
- Markowitz, H. M. The early history of portfolio theory: 1600-1960. *Financial Analysts Journal*, 55(4):5–16 (1999).
- MATLAB. *version 8.5.0 (R2015a)*. The MathWorks Inc., Natick, Massachusetts (2015).
- Merton, R. C. Lifetime portfolio selection under uncertainty: The continuous-time case. *The Review of Economics and Statistics*, 51(3):247–257 (1969).

- Meucci, A. Beyond Black-Litterman: Views on non-normal markets. *Risk*, 19(2):87–92 (2005).
- Meucci, A. *Risk and Asset Allocation*. Springer Science & Business Media, New York (2009).
- Meucci, A. Quant nugget 4: annualization and general projection of skewness, kurtosis and all summary statistics. *GARP Risk Professional - "The Quant Classroom"*, pages 59–63 (2010).
- Nelson, C. R. and Siegel, A. F. Parsimonious modeling of yield curves. *Journal of Business*, 60(4):473–489 (1987).
- Official Statistics of Finland (OSF). Prices of dwellings in housing companies. issn=2323-8801. Helsinki: Statistics finland. (2016). Available at http://www.stat.fi/til/ashi/index_en.html. Accessed 16.1.2016.
- Pennanen, T. and Koivu, M. Integration quadratures in discretization of stochastic programs. Number 11 in Stochastic Programming E-Print Series. Humboldt-Universität zu Berlin, Mathematisch-Naturwissenschaftliche Fakultät II, Institut für Mathematik (2002).
- Perold, A. F. and Sharpe, W. F. Dynamic strategies for asset allocation. *Financial Analysts Journal*, 44(1):16–27 (1988).
- Pflug, G. C. Scenario tree generation for multiperiod financial optimization by optimal discretization. *Mathematical Programming*, 89(2):251–271 (2001).
- Samuelson, P. A. Lifetime portfolio selection by dynamic stochastic programming. *The Review of Economics and Statistics*, 51(3):239–246 (1969).
- Scherer, B. and Kenneth, W. *The Oxford Handbook of Quantitative Asset Management*. Oxford University Press, Oxford (2012).
- The Finnish Pension Alliance TELA. Investment statistics and analyses, Summary of investment assets. (2016). Available at http://www.tela.fi/en/economy/investment_statistics_and_analyses/summary_of_investment_assets. Accessed 23.1.2016.
- Vasicek, O. An equilibrium characterization of the term structure. *Journal of Financial Economics*, 5(2):177–188 (1977).

- Xu, D., Chen, Z., and Yang, L. Scenario tree generation approaches using k-means and lp moment matching methods. *Journal of Computational and Applied Mathematics*, 236(17):4561–4579 (2012).
- Young, M. S. and Graff, R. A. Real estate is not normal: a fresh look at real estate return distributions. *The Journal of Real Estate Finance and Economics*, 10(3):225–259 (1995).
- Zenios, S. A. and Ziemba, W. T. *Handbook of Asset and Liability Management: Applications and Case Studies*, volume 2. Elsevier (2007).
- Ziemba, W. T. *The Stochastic Programming Approach to Asset, Liability, and Wealth Management*. CFA Institute, AIMR Publisher (2003).

Appendix A

Data characteristics

Table A.1: Quarterly first six moments of distributions of studied market invariants.

Moment	Equity	MM	Gov. yield	Spread	Real estate
Mean	2.1 %	0.00	-0.05	0.02	1.9 %
Standard deviation	11.9 %	0.61	0.41	0.37	2.8 %
Skewness	-1.06	2.12	0.14	-0.12	-0.51
Kurtosis	4.88	17.04	2.50	8.80	4.59
5th mom	-10.86	80.34	0.86	-7.25	-4.06
6th mom	40.23	500.03	9.00	109.11	31.18

Table A.2: Correlations of market invariants from the first node in the tree. Compare to historical correlations shown in table 4.5.

	Equity	MM	Gov bond	IG bond	Real estate
Equity	1.00	0.09	0.29	-0.45	0.19
MM	0.09	1.00	0.61	-0.07	0.21
Gov bond	0.29	0.61	1.00	-0.12	0.28
IG bond	-0.45	-0.07	-0.12	1.00	-0.33
Real estate	0.19	0.21	0.28	-0.33	1.00

Table A.3: Historical means and standard deviation of quarterly compound returns.

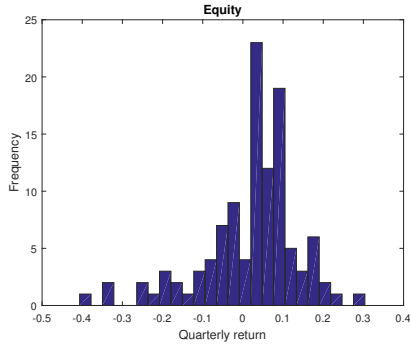
Moment	Equity	MM	Gov. bond	IG bond	Real estate
Mean	1.98 %	0.60 %	1.30 %	1.18 %	1.94 %
Standard deviation	12.00 %	0.41 %	1.61 %	1.78 %	2.82 %

Table A.4: Historical correlations calculated from quarterly returns.


	Equity	MM	Gov bond	IG bond	Real estate
Equity	1.00	-0.29	-0.32	0.22	0.25
MM	-0.29	1.00	0.40	-0.01	-0.12
Gov bond	-0.32	0.40	1.00	0.55	-0.34
IG bond	0.22	-0.01	0.55	1.00	0.13
Real estate	0.25	-0.12	-0.34	0.13	1.00

Table A.5: Historical moments of quarterly changes of both European money market yields and government bond yields and their Japanese counterparts.

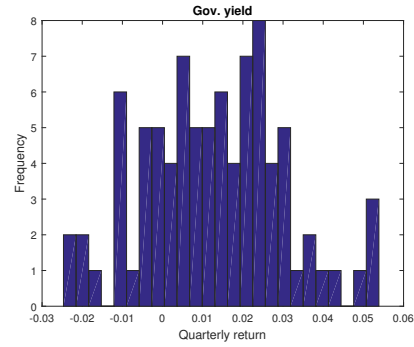
	Europe MM	Europe gov. 5Y	Japan gov. 1Y	Japan gov. 5Y
N	67	111	111	111
Data since	28/06/99	29/03/88	29/03/88	29/03/88
Mean	-0.05	-0.05	-0.03	-0.04
Standard deviation	0.44	0.41	0.40	0.40
Skewness	-1.11	0.14	0.05	0.26
Kurtosis	3.84	-0.47	5.36	3.70
Corr to Europe			0.53	0.44



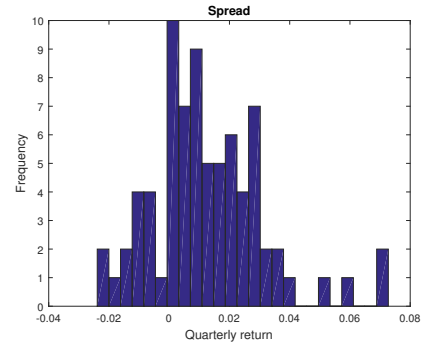
(a) Quarterly compound returns of equities.



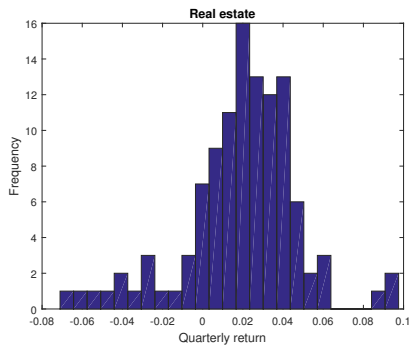
(b) Quarterly compound return of money market investments.



(c) Quarterly compound return of government bonds.




(d) Quarterly compound return of investment grade bonds.



(e) Quarterly compound returns of real estate.

Figure A.1: Historical quarterly compound returns of studied asset classes.



(a) Short rates.

(b) Bond rates.

Figure A.2: Comparison of historical yield time series of European and Japanese yields.




Figure A.3: Observed historical quarterly changes of Japanese and European short rate and bond rate changes.

Appendix B

Estimated short-rate model parameters

Table B.1: Estimated parameters for the CIR-model for Japanese government yield and credit spread with different specifications. T-statistics for each parameter in parenthesis.

Dataset	$1/\Delta t$	α	β	σ^2	γ	χ^2	p-value
Japan	250	0.2818 (1.9836)	-0.5193 (-1.7998)	0.2548 (12.9028)	0.5	3.0745	0.0795
Japan	52	0.2342 (1.6433)	-0.42 (-1.5447)	0.2323 (8.2552)	0.5	1.646	0.1995
Japan	4	0.0952 (1.1529)	-0.2684 (-1.7917)	0.151 (8.7835)	0.5	0.3167	0.5736
Japan	2	0.1029 (0.8554)	-0.3122 (-2.5679)	0.15 (6.746)	0.5	1.1803	0.2773
Japan	1	0.0047 (0.0587)	-0.0845 (-1.0636)	0.0733 (9.0585)	0.5	15.586	0.0001
Spread	250	0.1534 (0.7952)	-0.0851 (-0.3532)	0.0966 (13.281)	0.5	0.0615	0.8041
Spread	52	0.2858 (1.043)	-0.3126 (-0.8411)	0.0834 (6.3232)	0.5	10.5579	0.0012
Spread	4	0.8438 (6.4832)	-0.5359 (-3.9317)	0.6727 (6.674)	0.5	3.4891	0.0618
Spread	2	0.5945 (8.2803)	-0.5137 (-9.3332)	0.4968 (2.6852)	0.5	0.0002	0.9895
Spread	1	0.2427 (5.8535)	-0.2481 (-4.865)	0.1733 (7.2881)	0.5	0.9836	0.3213

Table B.2: Estimated parameters for money market investments in the CIR-model with different specifications. T-statistics for each parameter in parenthesis.

Dataset	$1/\Delta t$	α	β	σ	γ	χ^2	p-value
Europe	250	0,0065 (1,4677)	-1,0738 (-1,2789)	0,0124 (7,8641)	0,5	0,0122	0,9122
Europe	52	0,003 (0,8321)	-0,6764 (-0,9576)	0,0089 (4,6274)	0,5	1,1099	0,2921
Europe	4	-0,0103 (-4,6437)	0,1488 (0,3554)	0,0095 (4,4999)	0,5	8,1207	0,0044
Europe	2	-0,0016 (-0,5878)	-0,6784 (-2,5382)	0,0099 (8,1781)	0,5	7,7389	0,0054
Europe	1	-0,0002 (-0,1234)	-0,0398 (-0,6856)	0,0029 (10,6124)	0,5	13,5926	0,0002
Japan	250	0,0211 (0,6493)	-0,2166 (-0,7791)	0,0946 (10,6698)	0,5	0,0182	0,8928
Japan	52	0,0307 (0,8417)	-0,2759 (-0,931)	0,137 (5,616)	0,5	0,4293	0,5123
Japan	4	0,0342 (1,4755)	-0,4033 (-2,568)	0,1536 (2,992)	0,5	0,3049	0,5808
Japan	2	0,063 (2,576)	-0,6374 (-4,3445)	0,2016 (3,3888)	0,5	0,7814	0,3767
Japan	1	0,0248 (0,9959)	-0,5142 (-4,2571)	0,0852 (3,466)	0,5	9,8775	0,0017

Appendix C

Optimal allocations

Table C.1: Allocations in each node varying the risk aversion.

Annual target = 2.0 %
 a = 1E-04
 b = 10.0
 E(w) fix. = 1.20
 E(w) dyn. = 1.20

Fixed-mix					
Equity	MM	Gov bond	IG bond	Real estate	
98 %	0 %	0 %	0 %	2 %	

Dynamic						
Node	Equity	MM	Gov bond	IG bond	Real estate	
1	100 %	0 %	0 %	0 %	0 %	
2	0 %	0 %	0 %	100 %	0 %	
3	100 %	0 %	0 %	0 %	0 %	
4	100 %	0 %	0 %	0 %	0 %	
5	67 %	0 %	0 %	14 %	18 %	
6	100 %	0 %	0 %	0 %	0 %	
7	100 %	0 %	0 %	0 %	0 %	

Annual target = 2.0 %
 a = 1E-02
 b = 10.0
 E(w) fix. = 1.12
 E(w) dyn. = 1.15

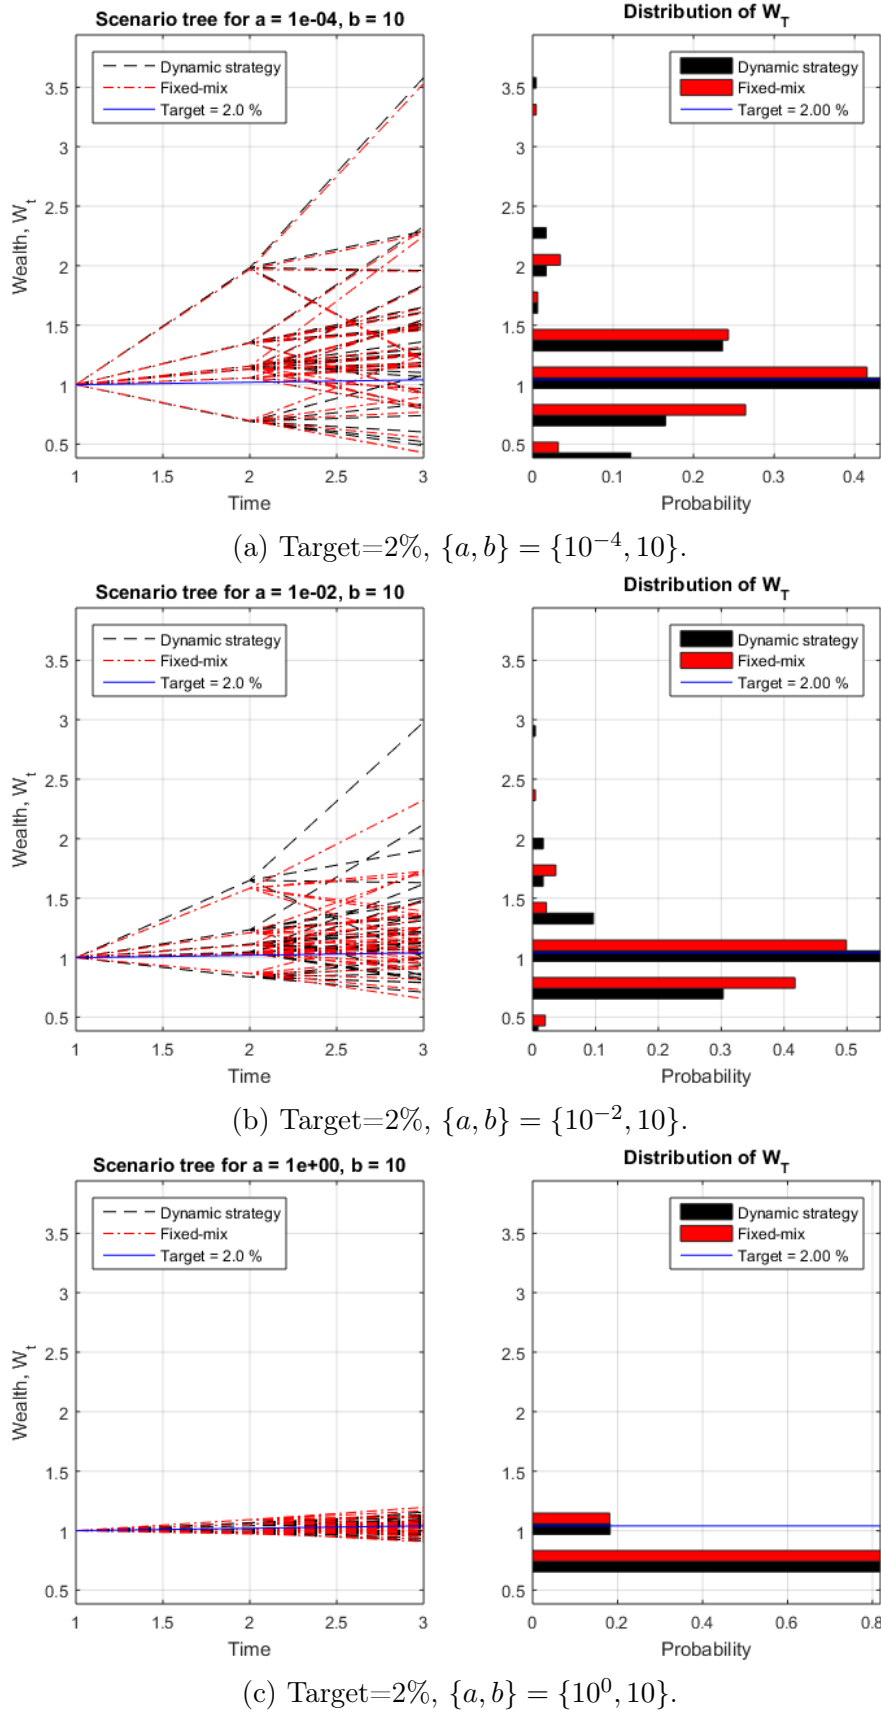
Fixed-mix					
Equity	MM	Gov bond	IG bond	Real estate	
49 %	0 %	0 %	0 %	51 %	

Dynamic						
Node	Equity	MM	Gov bond	IG bond	Real estate	
1	58 %	0 %	0 %	0 %	42 %	
2	0 %	0 %	0 %	100 %	0 %	
3	80 %	0 %	0 %	0 %	20 %	
4	100 %	0 %	0 %	0 %	0 %	
5	22 %	0 %	0 %	65 %	13 %	
6	100 %	0 %	0 %	0 %	0 %	
7	90 %	0 %	0 %	10 %	0 %	

Annual target = 2.0 %
 a = 1E+00
 b = 10
 E(w) fix. = 1.03
 E(w) dyn. = 1.03

Fixed-mix					
Equity	MM	Gov bond	IG bond	Real estate	
5 %	27 %	0 %	47 %	21 %	

Dynamic						
Node	Equity	MM	Gov bond	IG bond	Real estate	
1	4 %	46 %	2 %	35 %	14 %	
2	3 %	8 %	6 %	75 %	9 %	
3	6 %	2 %	2 %	72 %	18 %	
4	5 %	32 %	28 %	25 %	10 %	
5	7 %	5 %	32 %	46 %	10 %	
6	8 %	4 %	54 %	2 %	31 %	
7	11 %	38 %	1 %	2 %	47 %	

Figure C.1: Scenario trees when varying risk aversion coefficient a .

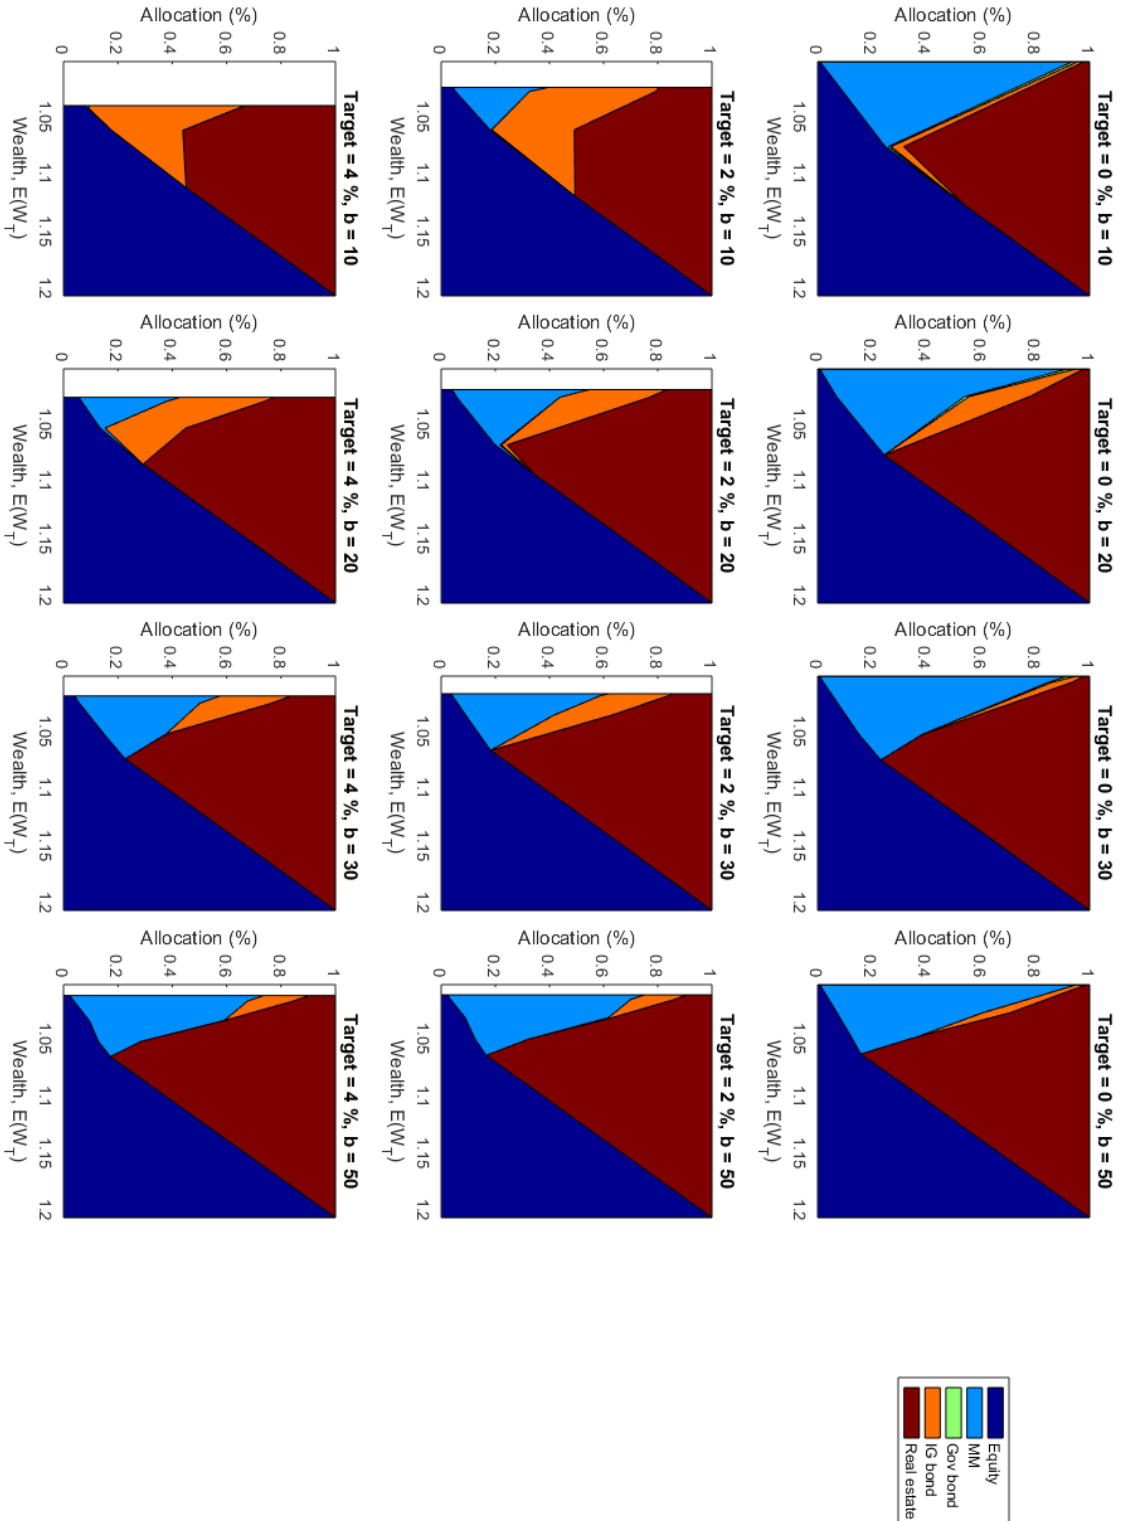


Figure C.2: Different allocations among the efficient frontier for different parameter values for the fixed-mix strategy. For each panel the target return and risk aversion b are fixed, a changes and results in different values in expected wealth.

Table C.2: Definitions of the $\{6, 6\}$ scenario tree in numerical terms corresponding to tree figures in 4.5. Table includes unique index for each node, time step it occurs, the probability of transition to the node, and returns for each asset class in the node.

Time	Node	Parent node	Transition probability	Log-returns for each transition				
				Equity	MM	Gov bond	IG bond	Real estate
1	1	1	1.00	0	0	0	0	0
2	2	1	0.04	14.5 %	-0.1 %	-3.6 %	-10.6 %	5.5 %
2	3	1	0.23	12.5 %	0.0 %	-0.7 %	-0.2 %	-7.1 %
2	4	1	0.03	68.5 %	0.3 %	-1.6 %	0.1 %	17.6 %
2	5	1	0.18	-36.9 %	0.0 %	-0.6 %	-1.4 %	3.7 %
2	6	1	0.22	30.3 %	-0.1 %	-1.7 %	5.1 %	6.6 %
2	7	1	0.31	5.5 %	-0.2 %	1.0 %	1.8 %	1.8 %
3	8	2	0.01	67.5 %	0.2 %	-0.1 %	-4.8 %	17.6 %
3	9	2	0.18	26.5 %	0.1 %	-0.8 %	16.4 %	8.4 %
3	10	2	0.05	-35.5 %	0.1 %	-3.0 %	-8.0 %	4.0 %
3	11	2	0.17	-22.7 %	0.0 %	5.9 %	12.0 %	0.7 %
3	12	2	0.42	9.0 %	0.0 %	2.4 %	9.2 %	2.0 %
3	13	2	0.17	13.7 %	0.0 %	1.8 %	8.0 %	-7.6 %
3	14	3	0.20	29.2 %	0.1 %	-1.2 %	6.6 %	9.5 %
3	15	3	0.13	-32.9 %	0.1 %	2.2 %	2.2 %	5.1 %
3	16	3	0.45	2.5 %	0.0 %	-0.3 %	1.1 %	1.1 %
3	17	3	0.16	10.2 %	0.1 %	2.1 %	2.9 %	-9.9 %
3	18	3	0.04	4.4 %	0.3 %	-3.7 %	-10.4 %	6.0 %
3	19	3	0.02	48.3 %	-0.1 %	3.4 %	3.6 %	21.6 %
3	20	4	0.05	-50.0 %	0.3 %	5.2 %	11.6 %	13.5 %
3	21	4	0.09	-1.2 %	1.0 %	-2.3 %	-2.4 %	17.4 %
3	22	4	0.61	14.4 %	0.5 %	1.7 %	1.9 %	-0.9 %
3	23	4	0.04	-50.0 %	0.5 %	-0.9 %	-1.0 %	-20.2 %
3	24	4	0.04	-50.0 %	0.0 %	0.6 %	-6.8 %	9.6 %
3	25	4	0.16	59.1 %	0.5 %	-1.7 %	4.6 %	12.9 %
3	26	5	0.07	-23.7 %	0.4 %	-2.9 %	-3.1 %	9.3 %
3	27	5	0.20	25.4 %	0.1 %	-0.8 %	7.2 %	3.2 %
3	28	5	0.07	-50.0 %	0.0 %	3.5 %	3.3 %	6.9 %
3	29	5	0.11	59.4 %	0.1 %	-1.9 %	-0.4 %	7.6 %
3	30	5	0.51	10.0 %	0.1 %	1.0 %	1.8 %	-1.0 %
3	31	5	0.05	-50.0 %	0.0 %	-0.8 %	-8.4 %	-10.4 %
3	32	6	0.08	54.2 %	0.0 %	1.9 %	0.7 %	12.2 %
3	33	6	0.05	8.0 %	0.0 %	5.1 %	0.9 %	-12.1 %
3	34	6	0.21	17.4 %	0.1 %	-2.1 %	-4.9 %	5.4 %
3	35	6	0.08	19.9 %	0.2 %	0.1 %	-2.9 %	5.2 %
3	36	6	0.17	-37.6 %	0.0 %	2.6 %	-0.6 %	3.7 %
3	37	6	0.41	9.3 %	0.1 %	0.9 %	-1.4 %	-2.3 %
3	38	7	0.59	9.1 %	-0.1 %	-1.0 %	0.0 %	2.5 %
3	39	7	0.05	-9.2 %	-0.1 %	-2.5 %	-9.6 %	-8.0 %
3	40	7	0.07	11.9 %	0.2 %	-2.5 %	5.0 %	6.2 %
3	41	7	0.16	-28.0 %	-0.2 %	0.5 %	0.3 %	-2.4 %
3	42	7	0.04	8.8 %	-0.4 %	-2.2 %	5.3 %	10.4 %
3	43	7	0.09	38.1 %	-0.2 %	-0.4 %	7.1 %	-4.7 %

Appendix D

Parameters for mean-variance optimisation

Table D.1: Parameters for the mean variance optimisation. Two-period means and standard deviations calculated from the trees in figure 4.5 and in table C.2.

	Equity	MM	Gov bond	IG bond	Real estate
Expected return	20.6 %	-0.1 %	-0.5 %	2.2 %	3.7 %
Standard deviation	39.2 %	0.2 %	1.5 %	3.4 %	8.0 %

Table D.2: Correlations of linear two-period asset returns in the scenario tree with a branching structure {6,6}.

	Equity	MM	Gov bond	IG bond	Real estate
Equity	1.00	0.25	-0.31	0.32	0.39
MM	0.25	1.00	-0.30	-0.15	0.28
Gov bond	-0.31	-0.30	1.00	0.30	-0.33
IG bond	0.32	-0.15	0.30	1.00	0.12
Real estate	0.39	0.28	-0.33	0.12	1.00