

Mean-variance optimization for life-cycle pension portfolios

by

J. M. Peeters Weem

to obtain the degree of Master of Science
in Applied Mathematics
at the Delft University of Technology,
Faculty of Electrical Engineering, Mathematics and Computer Science,
to be defended publicly on Wednesday May 24, 2017 at 04:00 PM.

Student number: 4152581
Project duration: September 5, 2016 – May 24, 2017
Thesis committee: Prof. dr. ir. C. W. Oosterlee, TU Delft, supervisor
Dr. M. van der Schans, Ortec Finance, supervisor
Dr. P. Cirillo, TU Delft

An electronic version of this thesis is available at <http://repository.tudelft.nl/>.

Abstract

In this thesis we discuss a framework for life-cycle construction. For the construction of life-cycles we use mean-variance optimization. Mean-variance optimization is a portfolio selection method used to find a combination of asset classes that has an optimal risk-return trade-off. We choose the replacement ratio, the pension income as fraction of labour income, as the quantity to be optimized. We find that using mean-variance optimization for the construction of deterministic life-cycles yields results that contradict conventional investment wisdom. It is mean-variance optimal to increase risk-taking as time passes, whereas conventional investment wisdom states that risk should decrease as time goes by. We introduce dynamic mean-variance optimization, where the asset allocation can adapt to changing circumstances, as an alternative to deterministic mean-variance optimization. We introduce an algorithm for dynamic mean-variance optimization of the replacement ratio, an extension of the dynamic mean-variance algorithm by Cong and Oosterlee. We show that dynamic mean-variance optimization can be used for life-cycle construction and that dynamic life-cycles outperform deterministic ones.

Preface

This thesis has been submitted for the degree Master of Science in Applied Mathematics at Delft University of Technology. The responsible professor is Kees Oosterlee, professor at the Numerical Analysis group of Delft Institute of Applied Mathematics. The research for this thesis project was carried out at Ortec Finance. The daily supervisor at Ortec Finance is Martin van der Schans, researcher at Ortec Finance. Ortec Finance aims to improve investment decision making by providing consistent solutions for advice and risk management through a combination of market knowledge, mathematical models and information technology.

I would like to thank Kees Oosterlee and Martin van der Schans for their input, guidance and feedback during this thesis project. Also, I would like to thank Pasquale Cirillo for being part of the thesis committee. Furthermore, I am grateful the members of the Research and EFIS departments at Ortec Finance for providing a pleasant working environment and interesting lunch discussions. Finally, I would like to thank my family, friends, and girlfriend for providing support through the duration of my studies.

*J. M. Peeters Weem
Rotterdam, May 2017*

Contents

1	Introduction	1
2	Mean-variance optimization	3
2.1	One-period mean-variance optimization	3
2.2	Utility functions.	4
2.3	Towards multi-period: cumulative returns	7
2.4	Multi-period mean-variance optimization	9
2.4.1	Definition of multi-period mean-variance optimization	9
2.4.2	Numerical examples	9
2.4.3	Two-period portfolio problem	10
2.5	Conclusion	12
3	Life-cycle investing	13
3.1	Definition of life-cycle investing	13
3.2	Literature overview	14
3.3	Life-cycle optimization	14
3.3.1	Replacement ratio, interest rates and bond returns	16
3.4	Example: a realistic base case	17
3.4.1	Bogle's rule.	17
3.4.2	Additional assumptions	18
3.4.3	Optimization results	19
3.5	Example: normal returns	21
3.5.1	Linear asset allocation	21
3.5.2	Individual weight optimization	23
3.6	Conclusion	25
3.A	Appendix to Chapter 3	26
4	Dynamic mean-variance optimization	29
4.1	Dynamic versus deterministic asset allocation	29
4.2	Definition of the dynamic mean-variance problem	29
4.3	Analytic solutions to the dynamic mean-variance and pre-commitment problem	31
4.3.1	Continuous-time solutions	31
4.3.2	Discrete-time solutions	31
4.4	Numerical solutions to the dynamic-mean variance problem.	32
4.4.1	Discretizing partial differential equations	32
4.4.2	Simulation-based methods	33
4.5	A simulation-based algorithm for mean-variance optimization of the replacement ratio	35
4.5.1	Generating an initial guess: the forward algorithm.	36
4.5.2	Updating towards the optimum: the backward algorithm	37
4.5.3	Analysis of the forward algorithm	38
4.5.4	Convergence of the backward algorithm	41
4.6	Conclusion	41
5	Regression and conditional expectation	43
5.1	Definition of conditional expectation through regression	43
5.2	Regression techniques	43
5.2.1	Regression applied to scenarios	44

5.3	Stratified state aggregation	44
5.4	Numerical tests	45
5.4.1	Test procedure	45
5.4.2	Geometric Brownian Motion.	46
5.4.3	Autoregressive returns	49
5.4.4	MGARCH returns	52
5.5	Conclusion	53
5.A	Regression techniques	54
5.A.1	Ordinary Least Squares regression	54
5.A.2	Lasso regression	55
5.A.3	Support Vector Regression	56
5.B	Appendix to Chapter 5	58
6	Numerical experiments	59
6.1	Numerical issues in implementation	59
6.1.1	Estimation of conditional expectation of asset returns	59
6.1.2	Convex optimization algorithm	59
6.1.3	Estimation of the value function	60
6.2	Normal, independent asset returns and constant risk-free rate	60
6.2.1	Regress-later	62
6.3	Ortec Finance Scenario returns	64
6.4	Conclusion	66
6.A	Appendix to Chapter 6	68
7	Conclusion	71
7.1	Conclusions.	71
7.2	Recommendations for further research	71
7.3	Recommendations for Ortec Finance	72
	Bibliography	73

Introduction

In this thesis, we introduce a framework for the construction of life-cycle portfolios for retirement savings. A life-cycle portfolio adjusts the asset allocation to the investor's age or years to retirement. The general belief is that a young investor can invest more risky, because he has many years left to retirement. Also, investors close to retirement, want to lock in what they have build up so far and prefer a less risky allocation [46]. This allocation problem is often solved by a combination of equity and bonds. Young investors may have a high allocation towards the risky equity and older investors may have a high allocation towards safer bonds. This relation between age and asset allocation is characterized by the 'glide path', named after its downward sloping shape. The glide path represents the fraction of the portfolio invested in equity, and hence is decreasing in time [20].

Retirement income is supposed to replace labour income. Therefore, we will evaluate a life-cycle by the density forecast of its replacement ratio. The replacement ratio is the yearly retirement income as a fraction of the average yearly labour income. In a mean-variance framework [32], the investor optimizes the expected value versus the variance of the replacement ratio. In this thesis, we introduce a mean-variance framework for the replacement ratio, with the goal of obtaining a life-cycle asset allocation for a pension portfolio. We have the following research objectives:

- investigate and evaluate methods to find the mean-variance optimal asset allocation;
- compare mean-variance optimal allocations to current standards in life-cycle investing.

Our research is organized as follows. In Chapter 2, we introduce mean-variance optimization for the buy-and-hold investor and for the investor that can choose (a priori) to rebalance or adjust his asset allocation periodically. We show why mean-variance optimization on wealth does not yield classic life-cycle allocations, i.e., decreasing equity allocation over time.

In Chapter 3, we introduce the deterministic life-cycle problem. Under simplifying assumptions, we construct mean-variance optimal life-cycles. We find that mean-variance optimization does not adhere to conventional life-cycle wisdom.

In Chapter 4, we introduce a dynamic approach to life-cycle investing. A dynamic approach implies that an investor can rebalance his asset allocation depending on the course of events. We introduce dynamic mean-variance optimization, in which the optimal asset allocation is dependent on the past. We extend the dynamic mean-variance optimization algorithm by Cong and Oosterlee [12] to optimize the replacement ratio. Furthermore, we extend the algorithm so that it can be used for asset returns of which the underlying process is not known.

In Chapter 5, we investigate methods to estimate conditional expectations in a simulation setting. These methods improve the estimation of conditional expectation in the algorithm developed in Chapter 4. We employ different methods for estimation of the conditional expectation by testing them on various models of asset returns.

We bring Chapter 4 and 5 together in Chapter 6, where we test the developed algorithm. We test the algorithm on two models: a basic model for asset returns and an involved model of which we do not know the underlying driver. We compare the results of the dynamic optimization to deterministic asset allocations

and one of the current standards in life-cycle investing: Bogle's rule, which states that an investor should invest $(100 - \text{age})\%$ of his wealth in equity [7].

In Chapter 7, we present our conclusions and recommendations for further research.

2

Mean-variance optimization

In 1952, Markowitz [32] introduced mean-variance optimization: one of the first methods for portfolio selection. Portfolio selection methods find a combination of asset classes that, according to the investor's belief, have an optimal risk-return trade-off [32].

This chapter introduces the mean-variance framework, which is still a widely accepted standard in portfolio selection, and its derivation from investor preferences, that is, utility functions. Also, we extend the framework to a multi-period approach where it is possible to change asset allocation over the investment horizon.

2.1. One-period mean-variance optimization

This section introduces the mean-variance framework as posed by Markowitz [32]. Suppose an investor can divide and invests a fraction w_i in asset i with stochastic return R_i . The mean and variance of the return $R = \sum_{i=1}^N w_i R_i$ of the portfolio are now given by:

$$\mathbb{E}[R] = \sum_{i=1}^N w_i \mu_i = \vec{w}^\top \vec{\mu}, \quad (2.1)$$

$$\text{Var}(R) = \sum_{i=1}^N \sum_{j=1}^N w_i w_j \sigma_{ij} = \vec{w}^\top \Sigma \vec{w}, \quad (2.2)$$

where:

$\vec{w} := [w_1, w_2, \dots, w_N]^\top$ are the asset weights,

$\vec{\mu} := [\mu_1, \mu_2, \dots, \mu_N]^\top$, $\Sigma := [\sigma_{i,j}] \in \mathbb{R}^{N \times N}$,

$\mu_i = \mathbb{E}[R_i]$, $\sigma_{i,j} = \mathbb{E}[R_i R_j] - \mu_i \mu_j$.

The mean-variance optimal allocation balances risk and return with a risk-aversion parameter λ in the following optimization problem:

$$\max_{\vec{w}} \quad \mathbb{E}[R] - \lambda \text{Var}[R] \quad (2.3a)$$

$$\text{s.t.} \quad \sum_{i=1}^N w_i = 1, \quad (2.3b)$$

$$w_i \geq 0, \quad \forall i \in \{1, \dots, N\}. \quad (2.3c)$$

Constraint (2.3b), called the no-leverage constraint, enforces that the total value of assets is equal to the total capital available for investment and constraint (2.3c), called the no-shorting constraint, enforces positive weights.

In his work, Markowitz [32] posed the problem slightly differently: "The $E - V$ rule states that the investor would (or should) want to select one of those portfolios which gives rise to the (E, V) combinations indicated as efficient ..., those with minimum V for given E or more and maximum E for given V or less." [32]. Where

E and V are the expectation and variance of the total return, respectively. These efficient combinations are given by the solution to either

$$\min_{\vec{w}} \quad \text{Var}[R] \quad (2.4a)$$

$$\text{s.t.} \quad \mathbb{E}[R] \geq \mu_0 \quad (2.4b)$$

$$\sum_{i=1}^N w_i = 1, \quad (2.4c)$$

$$w_i \geq 0, \forall i \in \{1, \dots, N\}, \quad (2.4d)$$

or

$$\max_{\vec{w}} \quad \mathbb{E}[R] \quad (2.5a)$$

$$\text{s.t.} \quad \text{Var}(R) \leq \sigma_0^2 \quad (2.5b)$$

$$\sum_{i=1}^N w_i = 1, \quad (2.5c)$$

$$w_i \geq 0, \forall i \in \{1, \dots, N\}. \quad (2.5d)$$

With the restrictions on asset allocation as above, all three representations of the mean-variance problem presented are mathematically equivalent, see [27]. Formulation (2.4) and (2.5) are intuitive for the general investor, picking a minimum level of risk or a maximum level of return and optimizing the portfolio according to that criterion is more natural than picking a risk-aversion parameter.

Optimization problems (2.3), (2.4) and (2.5) have no analytically tractable solution due to the inequality constraint on w_i , however, the problems are quadratic programs for which efficient numerical methods exist.

Mean-variance optimization can also be formulated in terms of the investors wealth $W = W_0(1 + R)$, instead of the portfolio return R . In formulation (2.3) this leads to a different optimal risk-return trade-off, the optimal risk-return trade-off depends on wealth. This dependence on wealth has a consequence: when wealth is larger, investments are less risky. This is caused by a difference in scaling between expectation and variance. The difference in scaling can be seen when plugging in cW instead of R into formulation (2.3): $\mathbb{E}[cW] - \lambda \text{Var}[cW] = c\mathbb{E}[R] - c^2 \lambda \text{Var}[R]$. This shows that the expectation operator is linear: the expectation of cW is equal to c times the expectation of W . On the other hand, the variance operator is non-linear; scaling wealth by c means the variance scales by c^2 . Hence for higher wealth, variance becomes relatively bigger compared to expectation and, for a fixed level of λ , investments are less risky. This is caused by the underlying assumptions on investor preferences, the topic of the next section.

2.2. Utility functions

Despite of its practical use, mean-variance optimization has no intuitive links with investor's preferences and thereby lacks a proper theoretical foundation [2]. An alternative is provided by expected utility theory where investor's preferences are modelled with a utility function.

The expected utility hypothesis states that the rational investor tries to maximize his expected utility [21]. This theorem is supported by the Von Neumann–Morgenstern utility theorem [34], which states the conditions under which the investor is rational. In this section, we summarize the aspects of expected utility theory which support mean-variance optimization.

Definition 2.1 (Utility function). A function $u : X \rightarrow \mathbb{R}$ is called a utility function if it is non-decreasing, concave and continuous on X .

We can summarize the important properties of a utility functions in the following lemma:

Lemma 2.1. A utility function u , as defined by Definition 2.1, has the following properties:

1. $\frac{du(x)}{dx} \geq 0$
2. The utility function corresponds to a risk-averse individual. A risk-averse individual prefers a certain payment A over an uncertain payment with expectation B , where the expectation of the uncertain payment is higher than the certain payment: $B > A$.

Proof. For a proof of property 1 we use the fact that $u(x)$ is non-decreasing. From an investment standpoint, property 1 implies that an investor never finds more wealth unattractive.

Property 2 follows from the concavity of the utility function. When $u(x)$ is concave and X is stochastic: $\mathbb{E}[u(X)] \leq u(\mathbb{E}[X])$ by Jensen's inequality. This shows that the expectation of the *uncertain* utility is smaller than the utility of the *certain* expectation. Certainty is preferred over uncertainty, hence concave utility functions imply risk-aversion. $\square$

We define three types of utility functions from which mean-variance optimization can be derived:

Definition 2.2 (Quadratic utility). A function u of the form $u(x) = ax - bx^2$ with $a, b > 0$ and $x < \frac{a}{2b}$ is called quadratic utility.

Definition 2.3 (Exponential utility). A function u of the form $u(x) = 1 - \exp(-cx)$ with $c > 0$ is called exponential utility.

Definition 2.4 (Power utility). A function u of the form $u(x) = \begin{cases} \frac{x^{1-\gamma}-1}{1-\gamma}, & \gamma > 0, \gamma \neq 1 \\ \log(x), & \gamma = 1 \end{cases}$ is called log or power utility.

The following theorem links utility functions to mean-variance optimization:

Theorem 2.1. *Mean-variance optimization can be derived from expected utility optimization in the following ways:*

1. *Optimization of expected quadratic utility, as defined in Definition 2.2, is equivalent to mean-variance optimization.*
2. *Under the assumption that portfolio returns are normal, optimization of expected exponential utility, as defined in Definition 2.3, is equivalent to mean-variance optimization.*
3. *Under the assumption that asset returns are log-normal, optimization of power utility, as defined in Definition 2.4, is equivalent to mean-variance optimization.*

Proof. Proof of statement 1; for quadratic utility we have:

$$\begin{aligned} \max \quad \mathbb{E}[aW - bW^2] &= \max \quad a\mathbb{E}[W] - b(\text{Var}(W) + \mathbb{E}[W]^2), \\ &= \max \quad a\mathbb{E}[W] - b\mathbb{E}[W]^2 - b\text{Var}(W), \end{aligned} \quad (2.6)$$

when $W \leq \frac{a}{2b}$ the term $a\mathbb{E}[W] - b\mathbb{E}[W]^2$ is increasing in $\mathbb{E}[W]$. Hence expected quadratic utility optimization is equivalent to mean-variance optimization.

Proof of statement 2; normality of the returns R implies $1 - \exp(-cW)$ is log-normally distributed and hence:

$$\max \quad \mathbb{E}[1 - \exp(-cW)] = \max \quad 1 - \exp\left\{-c\mathbb{E}[W] + \frac{c^2}{2}\text{Var}(W)\right\}, \quad (2.7)$$

which is equivalent to:

$$\max \quad \mathbb{E}[W] - \frac{c}{2}\text{Var}(W). \quad (2.8)$$

So under the assumption that portfolio returns are normal, expected exponential utility optimization is equivalent to mean-variance optimization.

Proof of statement 3; maximizing expected power utility of wealth yields:

$$\max \quad \mathbb{E}[W^{1-\gamma}] = \max \quad \mathbb{E}[\exp((1-\gamma)X)], \quad (2.9)$$

where $X = \log(W)$ is normally distributed. We see the same as for exponential utility:

$$\max \quad \mathbb{E}[\exp((1-\gamma)X)] = \max \quad \exp\left\{(1-\gamma)\mathbb{E}[X] + \frac{(1-\gamma)^2}{2}\text{Var}(X)\right\}, \quad (2.10)$$

which is equivalent to:

$$\max \quad \mathbb{E}[X] + \frac{1-\gamma}{2}\text{Var}(X). \quad (2.11)$$

$\square$

In the proof of statement 3 from Theorem 2.1, we optimize $X = \log(W)$. When we define wealth W as:

$$W = W_0(1 + R),$$

where W_0 is wealth at the start and R is the portfolio return, then

$$X = \log(W) = \log(W_0(1 + R)) = \log(W_0) + \log(1 + R).$$

So, in maximizing X , the $\log(W_0)$ term will drop out and hence we mean-variance optimize the portfolio log-return, the log of $1 + R$. The removal of W_0 in the optimization eliminates the dependence on wealth; the optimization is no longer dependent on size of wealth, only on the returns.

Dependence on wealth can be explained by the measures of absolute risk-aversion (ARA) and relative risk-aversion (RRA):

Definition 2.5 (Absolute risk-aversion). Let $u(x)$ be a utility function. Then $-\frac{u''(x)}{u'(x)}$ is called the coefficient of absolute risk-aversion (ARA).

Definition 2.6 (Relative risk-aversion). Let $u(x)$ be a utility function. Then $-\frac{xu''(x)}{u'(x)}$ is called the coefficient of relative risk-aversion (RRA).

We summarize the risk-aversion properties of utility functions in the following lemma.

Lemma 2.2. For exponential utility, absolute risk-aversion is given by

$$ARA = c,$$

which has constant absolute risk aversion (CARA). Furthermore, relative risk-aversion is given by

$$RRA = cW,$$

which is increasing in wealth. For log utility, absolute risk-aversion is given by

$$ARA = \frac{\gamma}{W},$$

which has decreasing absolute risk aversion (DARA). Furthermore, relative risk-aversion is given by

$$RRA = \gamma,$$

which has constant relative risk-aversion (CRRA). For quadratic utility, absolute risk-aversion is given by

$$ARA = \frac{2b}{a - 2bW}.$$

Furthermore, relative risk-aversion is given by

$$RRA = \frac{1}{\frac{a}{2bW} - 1},$$

and both ARA and RRA for quadratic utility are increasing in wealth.

Proof. All statements follow from applying Definition 2.5 and 2.6 to the definitions of utility functions, i.e., Definition 2.2, 2.3 and 2.4. $\square$

The combination of CARA and increasing RRA means that an investor with exponential utility who invests an extra sum of capital, will keep his absolute allocation towards risky assets equal, but will decrease his relative allocation, the fraction of the portfolio invested in risky assets. This is because, in absolute terms, size of capital does not influence his risk-aversion. In relative terms, however, his risk-aversion increases as his wealth increases and hence he decreases his relative exposure towards risky assets.

CRRA means that an investor will keep his relative allocation constant when he invests extra capital, the size of capital does therefore not matter, asset allocation will be the same. By many, CRRA is seen as an attractive property of the power utility function. Mean-variance, however, does not have this property (in general).

Finally, when both ARA and RRA are increasing in wealth, the dependence on the size of wealth in mean-variance optimization of wealth is observed. When wealth increases, an investor with quadratic utility will decrease his absolute wealth invested in risky assets, which follows from increasing ARA. Quadratic utility is the only utility function where optimization is equivalent to mean-variance optimization without any assumptions, so mean-variance optimization suffers from the same problem. In the next section we will see how these results influence portfolio optimization with multiple investment periods.

2.3. Towards multi-period: cumulative returns

In this section, an analytically tractable version of the multi-period problem is set out. Multi-period mean-variance optimization differs from the one-period version, as during the investment horizon the allocation can be adjusted.

In multi-period mean-variance optimization we try to optimize the cumulative returns, the total return gathered by investing over multiple periods. Suppose there are N assets of which the log-returns X are multivariate normal distributed: $X \sim \mathcal{N}(\vec{\mu}, \Sigma)$. Then the assets have returns $Y = \exp(X) - 1$, so that $Y + 1$ is log-normally distributed. The return of a single asset Y_i is characterized by:

$$r = \mathbb{E}[Y_i] = \exp\left(\mu_i + \frac{1}{2}\Sigma_{i,i}\right) - 1, \quad (2.12)$$

$$v = \text{Var}(Y_i) = \exp(2\mu_i + \Sigma_{i,i})\left(\exp(\Sigma_{i,i}) - 1\right). \quad (2.13)$$

The n -period log-returns of the assets, the cumulative log-returns, can be represented by the sum of n independent random variables with distribution as X , and we denote this sum by X_n . X_n has expectation $\vec{\mu}_n = n\vec{\mu}$ and (co)variance matrix $\Sigma_n = n\Sigma$. The cumulative log-returns can be translated to cumulative returns; the ratio of the expected cumulative return, and the variance of cumulated return is given in the following lemma.

Lemma 2.3. *The ratio of expectation of cumulative return over variance of cumulative return for n periods is given by:*

$$\frac{r_n}{v_n} = \frac{(1+r)^n - 1}{(1+r)^{2n}(\exp\{n\Sigma_{i,i}\} - 1)}, \quad (2.14)$$

where r_n is the expected cumulative return, v_n the variance of cumulative return, r is specified by equation (2.12) and $\Sigma_{i,i}$ is the variance of the log-return.

Proof. The cumulative return of an asset is characterized by:

$$r_n = \exp(n\vec{\mu}_i + n\Sigma_{i,i}) - 1, \quad (2.15)$$

$$v_n = (1+r_n)^2(\exp\{n\Sigma_{i,i}\} - 1), \quad (2.16)$$

$$= (1+r)^{2n}(\exp\{n\Sigma_{i,i}\} - 1), \quad (2.17)$$

taking the ratio $\frac{r_n}{v_n}$ yields the required result. $\square$

Lemma 2.4. *Taking the limit of the ratio specified in Lemma 2.3 yields:*

$$\lim_{n \rightarrow \infty} \frac{v_n}{r_n} = \lim_{n \rightarrow \infty} \frac{1}{(1+r)^n \exp\{n\Sigma_{i,i}\}} \quad (2.18)$$

Proof. Rewriting equation (2.14) yields:

$$\frac{r_n}{v_n} = \frac{1 - \frac{1}{(1+r)^n}}{(1+r)^n(\exp\{n\Sigma_{i,i}\} - 1)}.$$

Taking limits we obtain:

$$\lim_{n \rightarrow \infty} \frac{v_n}{r_n} = \lim_{n \rightarrow \infty} \frac{1}{(1+r)^n(\exp\{n\Sigma_{i,i}\} - 1)} = \lim_{n \rightarrow \infty} \frac{1}{(1+r)^n \exp\{n\Sigma_{i,i}\}}, \quad (2.19)$$

which proves the statement. $\square$

The ratio of expectation of return over variance of return $\frac{r_n}{v_n}$, given by equation (2.14), is to be optimized in mean-variance optimization and needs to be as large as possible. By Lemma 2.4, we have that the ratio decreases as n increases. For assets with higher one-period expected return r and variance v , the expected return-variance ratio decreases faster and hence low-return, low-variance assets get more attractive as the investment horizon gets longer. This is in contradiction with common knowledge that longer investment horizons increase risk-taking. We will illustrate the effect by an example.




Figure 2.1: Fraction of invested capital allocated towards the risky asset (equity) for different investment horizons (periods)




Figure 2.2: Ratio of expected return over variance for different investment horizons

Example 2.1. Suppose there are two risky assets, equity and bonds, of which the log-returns are normally distributed, the assets are uncorrelated. The log-returns have mean-vector and covariance-matrix:

$$\vec{\mu} = [0.06, 0.045]^\top, \Sigma = \begin{bmatrix} 0.2^2 & 0 \\ 0 & 0.1^2 \end{bmatrix}.$$

The aim is to optimize the total return of a portfolio consisting of equity and bonds. The optimal allocation is determined by solving the problem defined by equation (2.3). We repeat this optimization for different horizons. We assume a *constant asset allocation* over time, that is, the fraction of wealth invested in either asset is constant. This means that rebalancing might be necessary, as the random returns on the assets will change allocation.

Optimization of (2.3) with risk-aversion parameter $\lambda = 1$ is a constrained convex optimization problem. We optimize using the Python package *cvxopt*. The results for the equity allocation fractions are presented in Figure 2.1. The figure shows a decreasing allocation towards equity as the number of periods, that is, the investment horizon, increases. Note that Figure 2.1 shows the fixed equity allocation over the entire horizon, for portfolios with different investment horizons.

The result presented in Example 2.1 is consistent with the analysis of the ratio of expectation and variance of cumulative returns; as the investment term gets longer, that is, we invest for more periods, the asset allocation towards, the more risky, equity diminishes. This is inherent to mean-variance optimization. As the investment term increases, cumulative returns do too, which makes the variance of the returns increase more than the expectation of the returns and hence a less risky strategy is optimal. This result is demonstrated in Figure 2.2. This figure shows the ratio of expected return over variance of the optimal portfolio. This ratio clearly decreases as the horizon gets longer.

Example 2.1 does not fall under the power utility case, despite log-returns being assumed to be normal and therefore returns being log-normal. When returns are assumed log-normal, the CRRA log utility function implies mean-variance. The CRRA property would then imply that the size of returns does not matter for the optimization. A linear combination of two log-normally distributed variables, however, is *not* log-normally distributed. We are therefore in the quadratic utility situation, with the increasing absolute and relative risk-aversion property. So, size of returns does influence the optimization.

In Example 2.1, we fixed the asset allocation for the entire investment term. This assumption is in line with literature. In [38], Samuelson shows that, for a CRRA utility function, independent asset returns over periods, and the presence of a risk-free rate, asset allocation for the multi-period problem is equal to that of the single-period one, and is therefore constant over time. Samuelson's result can't be extended to mean-variance optimization in general but the next section will show that fixing the asset allocation is not overly restrictive.

2.4. Multi-period mean-variance optimization

This section shows numerical examples of mean-variance optimization for multi-period problems. These examples vary in length of investment period, and size & timing of cash flows going into the portfolio. We start with the definition of multi-period mean-variance optimization.

2.4.1. Definition of multi-period mean-variance optimization

We have to choose between optimization of returns or log-returns for multi-period portfolio selection. In a framework where returns are available for a given frequency, for example yearly, and the quantity of interest is total return, log-returns are preferred. Log-returns are preferred because the total log-return is calculated by summing the log-returns of every period, whereas for returns we have to take the product to obtain total return. When returns of multiple assets have to be combined, however, we cannot use log-returns: $w_1 \exp(R_1) + w_2 \exp(R_2) \neq \exp(w_1 R_1 + w_2 R_2)$, where w_1, w_2 are asset weights and R_1 and R_2 are log-returns. For returns we can easily calculate the combined return: $w_1 R_1 + w_2 R_2$. Summarized: log-returns aggregate easily through time, simple returns aggregate easily through assets.

As log-returns have to be transformed to simple returns to aggregate over assets, we opt to use returns for all optimization in this section. We introduce the multi-period mean-variance problem.

Problem 2.1 (Multi-period mean-variance optimization). The multi-period mean-variance optimization problem defines the optimal risk-return trade-off for a portfolio with multiple rebalancing moments, i.e., the asset allocation can change over time. The multi-period mean-variance optimal asset allocation is given by the solution to:

$$\max_{\vec{w}} \quad \mathbb{E}[W] - \lambda \text{Var}[W] \quad (2.20a)$$

$$\text{s.t.} \quad \sum_{j=1}^N w_{i,j} = 1, \quad \forall i \in \{1, \dots, M\}, \quad (2.20b)$$

$$w_{i,j} \geq 0, \quad \forall i \in \{1, \dots, M\}, \forall j \in \{1, \dots, N\} \quad (2.20c)$$

where W is given by:

$$W = W_0 \prod_{i=1}^M (1 + \vec{w}_i^\top \vec{R}_i), \quad (2.20d)$$

where:

- W_0 := starting capital
- $\vec{w}_i$:= the vector of asset weights at time i ,
- $\vec{R}_i$:= the vector of asset returns at time i ,
- N := number of assets available for investment,
- M := number of rebalancing moments.

2.4.2. Numerical examples

In this section, we use Monte-Carlo simulation to give examples of Problem 2.1.

Example 2.2 (Single cash flow). Assume for M periods the returns on equity and bonds are given by:

$$\vec{R}_i \sim \mathcal{N}(\vec{\mu}, \Sigma), \quad i \in \{1, \dots, M\}, \quad (2.21)$$

$$\vec{\mu} = [0.08, 0.043]^\top, \Sigma = \begin{bmatrix} 0.028 & 0 \\ 0 & 0.020 \end{bmatrix}. \quad (2.22)$$

For a single cash flow into the portfolio at the start of the investment period we choose $W_0 = 1$ so that $W = W_0 R_{tot} = R_{tot}$. For this simulation $N = 10000$ samples for all M periods are generated. The samples and assets weights are plugged in a function that returns the resulting wealth for all 10000 scenarios. This function is then plugged into equation (2.20a), which is maximized over the asset weights. We maximize using the optimizer from the Python package SciPy. The results are the equity weights for every period. Table 2.1 presents the result for up to 4 periods, with risk-aversion parameter $\lambda = 2$.

Table 2.1: Equity weights for Example 2.2. The table shows optimal equity weights for a varying number of rebalancing moments (periods).

	w_1	w_2	w_3	w_4
One period	0.61			
Two periods	0.58	0.58		
Three periods	0.57	0.56	0.56	
Four periods	0.54	0.54	0.54	0.54

Table 2.2: Equity weights for Example 2.3. The table shows optimal equity weights for a varying number of rebalancing moments (periods).

	w_1	w_2	w_3	w_4
One period	0.61			
Two periods	0.50	0.57		
Three periods	0.47	0.49	0.54	
Four periods	0.45	0.46	0.46	0.49

The equity weights are constant for every period, the limitation introduced in the analysis in Section 2.3 is acceptable. These results furthermore confirm the results from Section 2.3, equity allocation is lower for longer investment horizons.

We extend Example 2.2 to one with multiple cash flows.

Example 2.3 (Multiple cash flows). The setting is as in Example 2.2, with the addition of intermediate cash flows. To take into account intermediate cash flows, we adjust the definition of wealth, which is given by:

$$W = \sum_{i=1}^{M-1} c_i \left(\prod_{j=i}^M (1 + \tilde{w}_j^\top \tilde{R}_j) \right), \quad (2.23)$$

where c_i is the cash flow into the portfolio at the start of investment period i and all other variables are as in Problem 2.1.

We set all cash flows to one. Optimization is performed as in Example 2.2. For risk-aversion $\lambda = 2$ the optimization results can be found in Table 2.2.

When comparing Table 2.1 with Table 2.2, the first observation which stands out is that the equity weights are lower for the multiple cash flow case. The reason for this is that in the multiple cash flow case the total wealth, the quantity to be optimized, is higher than for the single cash flow. This leads to less risky asset allocations.

The main result is that equity allocation is increasing in time. Looking at the three-period line in Table 2.2 we can see that we start with an equity allocation of 0.49 and end up at 0.54, an increase of 5 percent points. This result goes against common knowledge that, as we approach the end of our investment periods, we should decrease our allocation to risky assets. Decreasing risky asset allocation would then reduce variance of returns in the last period and hence the variance of the total wealth. These results contradict this notion. In the next section we aim to explain this fact by focussing on the two-period portfolio.

2.4.3. Two-period portfolio problem

The analysis in this section starts with the two-period problem. We have one asset, which is normally distributed and the returns are independent of each other. The expectation and variance of the return can be different for different periods. We have two cash flows, one enters before period one and the second before period two. The whole system can then be described as:

$$R_1 \sim \mathcal{N}(\mu_1, \sigma_1^2), \quad R_2 \sim \mathcal{N}(\mu_2, \sigma_2^2),$$

$$W = \underbrace{(1 + R_1)(1 + R_2)}_{w_1} + \underbrace{(1 + R_2)}_{w_2}.$$

The expected value of the portfolio and its variance are given by:

$$\mathbb{E}[W] = (1 + \mu_1)(1 + \mu_2) + (1 + \mu_2), \quad (2.24)$$

$$\text{Var}(W) = \text{Var}(R_1 + 2R_2 + R_1 R_2). \quad (2.25)$$

The current representation of variance lends itself to a vector representation:

$$\text{Var}(W) = \text{Var} \left(\begin{bmatrix} 1 & 2 & 1 \end{bmatrix} \begin{bmatrix} R_1 \\ R_2 \\ R_1 R_2 \end{bmatrix} \right), \quad (2.26)$$

$$= \begin{bmatrix} 1 & 2 & 1 \end{bmatrix} S \begin{bmatrix} 1 \\ 2 \\ 1 \end{bmatrix}, \quad (2.27)$$

where S is the covariance matrix of the vector with returns and their cross-term. The more involved terms in the covariance matrix are given by:

$$\text{Cov}(R_1, R_1 R_2) = \mathbb{E}[R_1^2 R_2] - \mathbb{E}[R_1] \mathbb{E}[R_1 R_2], \quad (2.28a)$$

$$= \mu_2 \sigma_1^2, \quad (2.28b)$$

$$\text{Cov}(R_2, R_1 R_2) = \mu_1 \sigma_2^2, \quad (2.28c)$$

$$\text{Var}(R_1 R_2) = \mathbb{E}[R_1^2 R_2^2] - \mathbb{E}[R_1 R_2]^2, \quad (2.28d)$$

$$= \sigma_1^2 \sigma_2^2 + \mu_1^2 \sigma_2^2 + \mu_2^2 \sigma_1^2. \quad (2.28e)$$

We now have all the elements of matrix S . To isolate the effect of the second cash flow we split the weight vector as follows:

$$\begin{bmatrix} 1 \\ 2 \\ 1 \end{bmatrix} = \underbrace{\begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix}}_{v_1} + \underbrace{\begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix}}_{v_2}. \quad (2.29)$$

Now v_1 shows the effect of the first cash flow and v_2 the effect of the second cash flow. So, for the variance we have:

$$\text{Var}(W) = v_1^\top S v_1 + 2v_1^\top S v_2 + v_2^\top S v_2. \quad (2.30)$$

The terms containing v_2 are a measure of the effect of a second cash flow on the variance:

$$2v_1^\top S v_2 = 2\sigma_2^2(1 + \mu_1), \quad (2.31a)$$

$$v_2^\top S v_2 = \sigma_2^2. \quad (2.31b)$$

As expected the second cash flow brings in the variance of the second return σ_2^2 , which is represented by equation (2.31b). Via the cross term it also brings in the variance of the second return, but scaled by $1 + \mu_1$. Intuitively this makes sense: when the (expected) return in the first period is higher, there is more to win/lose with the investment in the second period. This is not only true in the absolute case (total wealth) but also in the relative case (return on wealth) as the returns compound. Hence the scaling of variance with $1 + \mu_1$. Adding the effect of v_1 we get the total variance:

$$\begin{aligned} v_1^\top S v_1 &= \sigma_1^2 + \sigma_2^2 + \sigma_1^2 \sigma_2^2 + 2(\mu_1 \sigma_2^2 + \mu_2 \sigma_1^2) + \mu_1^2 \sigma_2^2 + \mu_2^2 \sigma_1^2, \\ \text{Var}(W) &= \sigma_1^2 + 4\sigma_2^2 + \sigma_1^2 \sigma_2^2 + 2\mu_2 \sigma_1^2 + 4\mu_1 \sigma_2^2 + \mu_1^2 \sigma_2^2 + \mu_2^2 \sigma_1^2. \end{aligned}$$

As before notice the increased impact of μ_1 and σ_2^2 , caused by the second cash flow. This theoretical exercise is an elementary example of interaction between returns over multiple periods. It has shown that the effect of a second period is mainly driven by the variance of this second return, but also by the expected return of

the first period.

The implications for mean-variance optimization, and hence the results from Table 2.2 are clear: the effect of the expected return in the first period, μ_1 , on the total wealth is smaller than the return in the second, or subsequent periods. Expected returns in the first periods, however, have a bigger effect on the variance of the total wealth than returns in subsequent periods. Hence, reducing expected return in the first period by investing less risky in the first period and increasing risk in following periods is efficient in a mean-variance sense.

Although surprising, this result does not seem to contradict any existing literature. In [4], Basak and Chabakauri show that for the dynamic mean-variance problem, the static part of the solution is indeed increasing in time.

2.5. Conclusion

In this chapter mean-variance optimization for portfolio selection was discussed. The framework as proposed by Markovitz has been presented and extended with its utility function basis. In the utility based analysis the major drawback of mean-variance optimization was explained: mean-variance optimization is influenced by the size of a portfolio and more capital implies less risk. This is a drawback as most real-world investors won't behave like this. Mean-variance optimization is however very suitable for optimizing returns, not wealth, and even wealth optimization can be done in practice by varying the risk-aversion coefficient.

We extended the mean-variance framework to multi-period portfolios. We showed that it is mean-variance optimal to invest more risky as time goes by. An analysis of the two-period two cash flow portfolio showed that the return in the first period has a relatively larger effect on the variance than on the return of the total portfolio, implying that it is efficient to invest less risky in the first period in favour of more risk in the second period.

All in all the findings in this chapter imply that mean-variance is far from optimal for multi-period portfolio optimization. There are however two strong arguments in favour of mean-variance optimization. The first argument is that we have looked from a static point of view; an investor has to make all his investment decisions at the start of the investment horizon. In the real world the investor can make investment decisions based on past performance, which may make more realistic investment profiles. This real-world approach can be replicated using dynamic mean-variance optimization, a strategy that will be used for life-cycle optimization. The second argument in favour of mean-variance optimization is its applicability for practitioners. Finding the portfolio with the lowest variance given some minimum return is a natural way of approaching portfolio optimization, and one that is possible in the mean-variance framework. This immediately solves the dependence-on-wealth problem as the risk-aversion parameter is implicitly defined by the minimum return/wealth requirement, adjusting it to the wealth and risk-aversion of the investor. In the next chapters we discuss how this framework applies to life-cycle investing for retirement.

3

Life-cycle investing

This chapter describes life-cycle investing. We define life-cycle investing, give an overview of the literature and the current standards within the field of life-cycle investing. We end with an example, where we use mean-variance optimization to create a life-cycle asset allocation.

3.1. Definition of life-cycle investing

In life-cycle investing, the asset allocation is a function of the investor's age. The general idea is that young investors should allocate most of their capital to stocks and as they grow older they should allocate more of their capital to cash. This approach is in agreement with traditional financial planning [46] and for example John Bogle's investment advice to allocate $(100 - \text{age})\%$ in equity [7].

The asset allocations used for life-cycle portfolios can be characterized by the glide path. A glide path represents the fraction of the portfolio invested in different asset types over time. The glide path is a deterministic function of time, which only depends on the investor's current age. It is called the glide path because the allocation toward the risky asset goes down as time goes by, it glides towards a lower allocation. An example of a glide path can be found in Figure 3.1, which shows the fraction of a life-cycle portfolio invested in equity, the remainder is invested in cash.

Life-cycle investing and glide path construction are widely discussed topics in literature and no unambiguous definitions exist. We therefore given an overview of the relevant literature concerning life cycling investing.




Figure 3.1: An example of a glide path. It shows the fraction of wealth invested in equity over time, the remainder would be invested in cash.

3.2. Literature overview

The work of Markowitz [32], described in Chapter 2, is the basis of many portfolio selection problems. In earlier work on life-cycle investing, however, a CRRA utility approach is adopted. The work of Merton [33] and Samuelson [38] on investing over one's lifetime assumes investors to have CRRA utility. This leads to the conclusion that the investor should keep his asset allocation constant over time. Merton and Samuelson realize that this contradicts the standard approach where asset allocation becomes less risky as time goes by. Most of the work following the analysis by Merton and Samuelson attempts to justify a non-constant asset allocation.

A common approach to justify non-constant asset allocations is the introduction of human capital. The underlying assumption is that an investor who saves for retirement has two types of capital: human and financial. Financial capital is all cash plus investments, human capital is the present value of future labour income. Bodie, Merton and Samuelson [6] incorporate flexible labour income into their modelling of human wealth. Flexibility means that an individual can choose how much he works over his lifetime. Their main conclusion is that young individuals will invest more risky than older individuals because they have more labour flexibility and can therefore absorb financial losses by working more. There is no specific mentioning of retirement in their paper and it is possible for an individual to keep working indefinitely. Viceira [45] extends this model by modelling retirement as a period with no labour income. Viceira shows that young, employed people will invest more in risky assets, but also that they should try to hedge their labour income risk. When labour income is positively correlated with risky asset returns, this can imply that employees should invest more in the riskless asset.

The papers by Bodie et al. and Viceira show that introducing human capital indeed implies that young investors should invest more risky than older investors. An overview of the literature on the implications of labour income for life-cycle investing is given by Viceira [46].

The human capital based approach to glide path construction seems to be the preferred approach in practice. For example, TIAA-CREF Asset Management uses the human capital framework, combined with Monte-Carlo simulation, to design their life-cycle funds [41]. Vanguard creates their target-date portfolios' glide paths using a utility function approach, taking into account, among other factors, the retirement age. The higher the retirement age, the longer salary is earned, which is equivalent to holding a bond, according to Vanguard's model [44].

Finally, there exists literature disputing the conventional wisdom that glide paths should be decreasing. Shiller compared different glide paths on a long horizon and concluded that a 100% equity allocation works better than standard glide paths [39]. Shiller mentions that a reason for the underperformance of standard glide paths is that they invest risky when saved capital is low. Younger employees earn less money and have not yet saved much for their pension, when they only invest risky at this time, not much return will be generated. The argument that a too small a part of capital is allocated to the risky asset is used in the majority of papers arguing against the conventional life-cycles, an overview of their results and arguments is presented by Estrada [18].

3.3. Life-cycle optimization

From the literature overview, it is clear that no general definition of the life-cycle problem exists. We therefore define our own problem and possible solutions. Our point of departure is the development of a deterministic asset allocation (glide path) for pension savings.

We start with an investor who wants to invest his capital in a combination of stocks and bonds. The investor does not consume or invest additional capital during his investment horizon of T years. The allocation can be rebalanced periodically to a predetermined allocation. An example is an investor who rebalances his portfolio each year to hold a fixed ratio of stocks and bonds. In a mean-variance framework this leads to the following optimization problem:

Problem 3.1. The asset allocation $\{\bar{x}_t\}_{t \in \{0, \dots, T-1\}}$ with the optimal trade-off between risk and return of termi-

nal wealth satisfies:

$$\begin{aligned} \max_{\{\tilde{x}_t\}_{t \in \{0, \dots, T-1\}}} \quad & \mathbb{E}[W_T] - \lambda \text{Var}[W_T], \\ \text{s.t.} \quad & \sum_{i=1}^M x_{t,i} = 1, \quad \forall t \in \{0, 1, \dots, T-1\}, \\ & x_{t,i} \geq 0, \quad \forall t \in \{0, 1, \dots, T-1\}; \forall i \in \{1, \dots, M\}, \end{aligned} \quad (3.1)$$

where:

W_t := Investor's total wealth at time t ,

T := Final time of investment period in rebalancing frequency, e.g., years,

N := Number of asset allocation changes

M := Number of assets available for investment,

$\tilde{x}_t$:= Vector of length M containing fractions of wealth allocated towards every asset at time t ,

$x_{t,i}$:= The i -th element of $\tilde{x}_t$,

$\tilde{R}_t$:= Vector of length M containing returns on every asset from time t to $t+1$,

λ := Risk aversion coefficient.

Furthermore, wealth at time t , W_t , is given by:

$$W_t = W_{t-1} \tilde{x}_{t-1}^\top \tilde{R}_{t-1} + C_t \quad (3.2)$$

where C_t is defined as the cash flow into portfolio at time t and for wealth at time 0 we have:

$$W_0 = C_0. \quad (3.3)$$

The number of optimization parameters depends on the number of time periods N and the number of assets M . In practice N is often so large that optimization of all weights $\{\tilde{x}_t\}_{t \in \{0, \dots, T-1\}}$ is not feasible. By using functions that describe $\tilde{x}_t$ over time we reduce the number of optimization parameters.

It should be clear that life-cycle asset allocations are deterministic, i.e., they can vary over time, but their future values are set and do not depend on the course of events. In literature Problem 3.1 is often approached dynamically. When we solve Problem 3.1, however, we mean a deterministic problem where we find, at time $t = 0$, a deterministic asset allocation so that W_T is mean-variance optimal.

When going into retirement, an investor might not care so much about the total sum of money he has gathered but more about how much pension he can receive. The final wealth W_T can be transformed to a number of fixed, yearly payments, representing retirement income. We assume that at time $t = T$, the time of retirement, the investor can get a yearly cash flow C_r for N_{ret} years. The investor will be interested in the size of the yearly cash flow C_r relative to time-average yearly labour income $\bar{I}$. Therefore, we define the replacement ratio:

Definition 3.1 (Replacement ratio). The replacement ratio (RR) is defined as the retirement cash flow C_r as a fraction time-average labour income $\bar{I}$:

$$RR = \frac{C_r}{\bar{I}}.$$

The pension investor wishes to maximize his replacement ratio, for which he has to maximize the retirement cash flows C_r .

The retirement cash flows C_r are in the future and hence we can discount them, their present value is lower than their nominal value. Discounting the nominal value of the cash flows means that we calculate the money we need to have invested in a risk-less asset by now to obtain the nominal value in the future. This means dividing by the return of the risk-free asset at time $t = T$, $R_f(T)$. We can calculate the present value of all future cash flows at time T :

$$PV(T) = \sum_{i=0}^{M-1} \frac{C_r}{(1 + R_f(T))^i}. \quad (3.4)$$

Here we consider the same risk-free rate for all cash flows. This assumption is a simplification since the interest rate for an investment term of 1 year is lower than that for a term of 10 years. This difference in short-term long-term rates is visible in the yield curve, which represents the term-structure of interest rates. We can introduce a term-structure where the risk-free rates depend on the time in the future the cash flow is paid out:

$$\sum_{i=0}^{N_{ret}-1} \frac{C_r}{(1 + R_{f,i}(T))^i}, \quad (3.5)$$

where the $R_{f,i}(T)$ stands for the risk-free rate for a term of i years at time $t = T$. The present value of the future cash flows should be equal to the final wealth W_T and hence we can write:

$$W_T = \sum_{i=0}^{N_{ret}-1} \frac{C_r}{(1 + R_{f,i}(T))^i}, \quad (3.6)$$

$$C_r = \frac{W_T}{\sum_{i=0}^{N_{ret}-1} \frac{1}{(1 + R_{f,i}(T))^i}}. \quad (3.7)$$

The yearly retirement cash flow C_r , as defined by equation (3.7), depends on the interest rates. The risk-free interest rates are stochastic and correlated with bond returns, and possibly other asset returns. Therefore, optimizing wealth does not suffice for the pension investor. We therefore define an optimization problem in terms of the replacement ratio:

Problem 3.2. The asset allocation $\{\tilde{x}_t\}_{t \in \{0, \dots, T-1\}}$ with the optimal trade-off between risk and return of the replacement ratio satisfies:

$$\begin{aligned} \max_{\{\tilde{x}_t\}_{t \in \{0, \dots, T-1\}}} \quad & \mathbb{E}[RR] - \lambda \text{Var}[RR] \\ \text{s.t.} \quad & \sum_{i=1}^M x_{t,i} = 1, \quad \forall t \in \{0, 1, \dots, T-1\}, \\ & x_{t,i} \geq 0, \quad \forall t \in \{0, 1, \dots, T-1\}; \forall i \in \{1, \dots, M\}, \end{aligned} \quad (3.8)$$

where RR is given by Definition 3.1, C_r is given by equation (3.7) and all other variables are as in Problem 3.1.

3.3.1. Replacement ratio, interest rates and bond returns

The replacement ratio, as defined in Definition 3.1, with retirement cash flow C_r as defined by equation (3.7), is dependent on interest rates. When interest rates go up, the value of the denominator of C_r decreases and hence C_r increases. For declining interest rates, the value of C_r will decrease. When interest rates are volatile, C_r will be volatile. To reduce this volatility, and investor can consider investing in bonds.

A bond is a financial product that will pay the owner a periodical *coupon* and at maturity it will pay the owner the *principal*. The value of a bond is given by de present value of all future cash flows:

$$PV = \sum_{i=1}^N \frac{C_i}{(1 + y)^i},$$

where i is the time at which cash flow C_i is received by the owner of the bond and y is called the yield, the rate that determines by how much the future cash flows are discounted [25]. The yield is closely related to interest rates. Suppose we have a zero-coupon bond, i.e., a bond that only pays its principal at maturity and no coupons. The price P of such a bond with maturity at time T is given by:

$$P = \frac{C}{(1 + y)^T}, \quad (3.9)$$

where C is the principal amount. It follows from equation (3.9) that buying a zero-coupon bond at price P implies owning $C = P(1 + y)^T$ at time T , which is equivalent to putting P in a bank account with interest rate y for T years. So, when we talk about zero-coupon bonds, we can replace the yield y by interest rate R .

When interest rates are non-constant, bond prices fluctuate. Expanding the bond price around R_0 we obtain:

$$P = \frac{C}{(1 + R)^T} = \frac{C}{(1 + R_0)^T} - \frac{CT}{(1 + R_0)^{T+1}}(R - R_0) + \mathcal{O}\left((R - R_0)^2\right), \quad (3.10)$$

so, the bond return $\frac{\Delta P}{P}$ is given by:

$$\frac{\Delta P}{P} = \frac{-T}{1+R_0} \Delta R + \mathcal{O}((\Delta R)^2) \approx \frac{-T}{1+R_0} \Delta R. \quad (3.11)$$

Equation (3.11) shows that when interest rates increase, i.e., $\Delta R > 0$, bond returns are negative. When interest rates decrease, bond returns are positive. Furthermore, bonds with longer maturity T are more strongly influenced by interest rate changes than bonds with shorter maturity.

So, equation (3.11) shows that when interest rates increase, bond returns are negative. Conversely, equation (3.7) shows that when interest rates increase, pension payments increase. Therefore, there exists a negative correlation between bond returns and the replacement ratio, which makes bonds an investment that can decrease volatility in C_r .

Two remarks on the negative correlation described above. First, in equation (3.7), we discount every payment with a different rate, as interest rates are not constant for every maturity. The interest rate for a maturity of 10 years can be higher than the interest rate for a maturity of 1 year. These differences are made visible in a yield curve, which plots the interest rate for every maturity. The value of the pension payment C_r does not only change when interest rates in general go up or down, but also when the shape of the yield curve changes. Second, in equation (3.7), we use interest rates at time T to calculate C_r . Therefore, the negative correlation between bonds and the retirement cash flow C_r becomes more apparent as we approach the end of the investment horizon T . Both remarks indicate that, when using bonds as a way to decrease replacement ratio volatility, both timing of the moment bonds are introduced into the portfolio as well as the maturity of the bonds in the portfolio are important. Our focus will be on finding the optimal asset allocation over time, so, when and in which quantity bonds are added to the portfolio.

3.4. Example: a realistic base case

In this section, we aim to solve optimization Problem 3.2 numerically. Problem 3.2, however, is a high-dimension problem when there are more than a few rebalancing moments. High-dimensionality makes finding a solution to the problem computationally intensive. On top of that, it is expected that the asset weights have some correlation between periods, which is not modelled when optimizing them individually. We therefore use heuristics to approximate the solution.

3.4.1. Bogle's rule

A well-known approach to the life-cycle portfolio was suggested by John C. Bogle, founder of Vanguard [7]. He proposed the following: one should invest (100-age)% of wealth in equity and the rest in bonds. So the fraction invested in equity (or bonds) is a linear function of age, we will call this "Bogle's rule". We improve on Bogle's rule by proposing a linear function of time for the fraction invested in every asset class. These functions will have the form:

$$x_i(t) = \alpha_i + \beta_i t, \quad (3.12)$$

which leaves only two parameters to optimize for every asset class.

The no-shorting and no-leverage constraints impose bounds on the parameters α_i and β_i :

$$\begin{aligned} \sum_{i=1}^M x_i(t) &= 1, \quad \forall t, \\ 0 \leq x_i &\leq 1, \quad \forall i \in \{1, 2, \dots, M\} \end{aligned}$$

The first constraint is satisfied by setting $x_M(t) = 1 - \sum_{i=1}^{M-1} x_i(t)$. The second is implemented by bounding $\alpha_i \in [0, 1]$ and not demanding the function to be linear, but piecewise linear, giving it the following form:

$$x_i(t) = \min(\max(\alpha_i + \beta_i t, 0), 1) \quad (3.13)$$

Although it is no longer necessary to bound α_i , we still choose to do so to keep the results more interpretable: α_i indicates the allocation at time $t = 0$.

We find the optimal parameters α_i and β_i by plugging (3.13) into Problem 3.2. Problem 3.2 defines the mean-variance optimization for the replacement ratio where the parameter λ has to be chosen (or in classical portfolio optimization problems, varied to obtain the efficient frontier). Although varying λ will yield

sufficient results, it is easier and more easily interpretable to implement the problem as follows:

$$\begin{aligned} \min \quad & \text{Var}[RR] \\ \text{s.t.} \quad & \mathbb{E}[RR] > RR_{min} \end{aligned} \quad (3.14)$$

Implementation of this problem is done by varying the parameter RR_{min} between the expected RR for minimum variance and the maximum expected RR.

Unfortunately, even with a basic allocation rule given by equation (3.12) and using the optimization form given by equation (3.14), it is not possible to make the optimization analytically tractable. We therefore optimize numerically and proceed as follows:

1. A function that uses the optimization parameters $\{\alpha_i\}_{i \in \{1, \dots, M\}}$, $\{\beta_i\}_{i \in \{1, \dots, M\}}$ calculates the replacement ratio for a set of simulated returns and interest rates.
2. We calculate the mean and variance of the replacement ratios from step 1.
3. Under the conditions that the mean of the replacement ratios is bigger than RR_{min} , we use the optimizer from the ScyPi library for Python to minimize the variance, optimizing $\{\alpha_i\}_{i \in \{1, \dots, M\}}$ and $\{\beta_i\}_{i \in \{1, \dots, M\}}$.

3.4.2. Additional assumptions

The replacement ratio as given in Definition 3.1 depends on the yearly pension one receives and the time-averaged labour income. Equation (3.7) defines the yearly pension one receives, which depends on accumulated wealth. Both accumulated wealth and time-averaged labour income depend on the investor's wage. Furthermore, the accumulated wealth depends on the assets available for investment and the investment horizon.

Wages and ingoing cash flows

Accumulated wealth depends on how much money was saved. Every year a fraction of the investor's wages is saved and invested in the portfolio. For wages we use an average (Dutch) career path, plus a fixed percentage of wage inflation, 2.5%. The time-average value of these wages serves as the denominator in the replacement ratio calculation. Note that the fixed inflation assumption can be replaced by stochastic forecasting models for inflation.

The fraction of income saved for pension (premium η) is age-dependent and increases as one gets older. The percentages are summarized in Table 3.5 in Appendix 3.A. These are not percentages of income but of the pensionable amount, which is wage (I) minus franchise (F), where franchise is dependent on the current level of AOW (state pension). The ingoing cash flows are given by:

$$C(t) = (I(t) - F(t)) * \eta(t). \quad (3.15)$$

The ingoing cash flows form the wealth invested in a range of different assets.

Investment options

The asset-classes that are available for investment are:

- European equity;
- German government bonds (all tenors);
- Investment grade European corporate bonds (selected tenors);
- Listed European real-estate;
- Commodities (GSCI, in Euro).

All assets are denoted in Euro or converted to Euro.

The choice for these asset-classes is based on currently common choices in life-cycle portfolios. A risky part with equities, a less risky part with bonds and for diversification purposes we add real-estate and commodities. All assets are denominated in Euro because of the focus on the Dutch pension sector.




Figure 3.2: Bogle's rule: equity allocation over time and distribution of replacement ratio generated by following this rule

The returns on assets are simulated using the Ortec Finance Scenarios (OFS). A proprietary model capable of generating asset returns with fat tails, time-varying volatility etc. [26].

Investment horizon

In this example, the investment horizon for this portfolio is 40 years. Reweighting may take place once a year. After 40 years the accumulated wealth is used to buy the yearly retirement cash flow.

3.4.3. Optimization results

We optimize the replacement ratio using two asset classes: equity and 3-months government bonds. The 3-months government bond is the closest one can get to the risk-free asset, as interest rate changes have limited effect on short-term bonds, see Section 3.3.1. Presentation of results will be done visually for the most part, graphing asset allocation and distribution of replacement ratios, besides this, summary statistics will be presented. We find that mean-variance optimal linear asset allocations do not necessarily adhere to traditional investment wisdom, i.e., the optimal asset allocation does not always become less risky as time goes by.

Result for Bogle's rule

The first result is not one from an optimization but the implementation of Bogle's rule. For a 25 years old person who saves for 40 years, equity allocation will be 75% at the start and 35% at the end. What this and the distribution of the replacement ratios (RRs) look like can be seen in Figure 3.2. The mean of RRs is 1.183 and the variance 0.0718, a benchmark for the optimization result. It is important to notice that an RR of more than 1 seems rather high. This is caused by the assumption that there only needs to be a pension for 20 years. The ingoing cash flows are calculated based on probabilities on age of death, inflation etc., in such a way that they should generate an expected RR of 0.7. In general this means the cash flows are higher than necessary.




Figure 3.3: Equity fraction over time and replacement ratio distribution for minimization of variance with a minimum level of replacement ratio of 1.1

Linear allocation for two assets

The linear asset allocation for two assets has been optimized for different minimum levels of the expected RR; first we present the equivalent of Bogle's rule: the mean RR should be 1.18 or more. The similarity with Bogle's rule is remarkable, the equity allocation rule is $0.71 - 0.33t$ versus $0.75 - 0.40t$ for Bogle's rule and variance is slightly lower for the optimal version, but not significantly different from 0.0718 (which is rounded), see Table 3.1. We can conclude that improving on Bogle's rule is hardly possible, when asset allocations are linear in time.

For all other minimum expected RRs, the figure looks similar to Figure 3.2, with a decreasing allocation towards equity over time, with the exception of a minimum expected RR of 1.1, see Figure 3.3. This clearly shows an upward sloping equity allocation, which is at odds with conventional investing wisdom. Nevertheless, it is consistent with our findings in Section 2.4.3, i.e., investing less risky in the first period and increasing risk afterwards is mean-variance efficient.

The rest of the results for different minimum expected RRs can be found in Table 3.1. In this table we observe that the median is lower than the mean for all minimum expected RRs, which shows that all distributions of the replacement ratio have positive skew.

Linear allocation for six assets

We optimize for 6 different assets, the 5 mentioned in Section 3.4.2, with two tenors for government bonds, 3 months and 10 years, and 5 year corporate bonds. Our results show that for a minimum expected RR of 1.3, shown in Figure 3.5 in Appendix 3.A, diversification is not efficient; corporate bonds and equity are the only assets in the optimal portfolio. It does also show that moving from risky assets (equity) to less risky assets (corporate bonds) over time is efficient. This is one of the few cases where the result is in line with conventional investment wisdom.

For example, for a minimum expected RR of 1.1, we obtain Figure 3.6 in Appendix 3.A. Figure 3.6 shows a decreasing bond allocations as time goes by, whereas equity and other risky assets have an increasing allocation as time goes by. This result is unexpected and seems incorrect for two reasons. First, we expect that we invest more in bonds as time goes by, as these form a hedge for interest rate changes, see Section 3.3. Second,

Table 3.1: Estimated linear allocation parameters as in equation (3.13) and statistics of replacement ratios for two-asset mean-variance optimization of the replacement ratio

	α	β	mean	median	variance
Bogle	0,75	-0,4	1,18	1,12	0,0718
Min. exp. RR: 0.8	0	-0,82	0,997	0,95	0,0255
Min. exp. RR: 1	0,53	-2,10	1	0,95	0,0259
Min. exp. RR: 1.1	0,26	0,08	1,1	1,05	0,0426
Min. exp. RR: 1.18	0,71	-0,33	1,18	1,12	0,0718
Min. exp. RR: 1.2	0,76	-0,37	1,2	1,14	0,0793
Min. exp. RR: 1.3	1	-0,52	1,3	1,21	0,139
Min. exp. RR: 1.4	1	-0,34	1,4	1,28	0,227
Min. exp. RR: 1.5	1	-0,19	1,5	1,34	0,347

the OFS predicts rates going up strongly a few years from now. So, in the used sample, the average interest rates are going up. Which means bond returns, especially those with a long duration, will be low, mainly in the first years. The first years are the part of the time-window where allocation towards bonds is highest. Figure 3.6 shows that despite these two effects, bonds are still an attractive asset at the start of the investment horizon. A possible explanation is given in Section 2.4.3. Although the results in Section 2.4.3 are for a two-asset problem and for the optimization of wealth, not replacement ratio, it could explain why we invest more into risky assets as time goes by; that is, Section 2.4.3 concludes that risky assets get more attractive as time goes by.

Table 3.2 shows some statistics of the RRs for the different optimizations. Results are similar to those of the two assets example. Variance is lower than in the two-asset example, mainly for the lower minimum RRs, which is where one can benefit most from the diversification. For higher minimum RRs there is a bigger need for equity and hence smaller opportunity to diversify.

Table 3.2: Statistics of replacement ratios for six-asset mean-variance optimization of the replacement ratio

	mean	median	variance
Min. exp. RR: 0.8	1,02	0,99	0,021
Min. exp. RR: 1	1,03	0,99	0,021
Min. exp. RR: 1.1	1,1	1,06	0,024
Min. exp. RR: 1.2	1,2	1,16	0,036
Min. exp. RR: 1.3	1,3	1,24	0,077
Min. exp. RR: 1.4	1,4	1,31	0,162
Min. exp. RR: 1.5	1,5	1,36	0,305

3.5. Example: normal returns

For some insight in the results presented in the previous section we investigate what happens when we take asset returns and interest rates to be normally distributed, as in Chapter 2. For linear asset allocations, as in the previous section, we find that the mean-variance optimal equity allocation can be decreasing as well as increasing over time. When we remove the linear assumptions, we find that the mean-variance optimal asset allocation increases in risk over time, which explains the increasing equity allocation. On the other hand, we find that the correlation between bond returns and the replacement ratio increases bond allocation at the end of the investment horizon, which explains decreasing equity allocation.

3.5.1. Linear asset allocation

First of all, the same methodology as in Section 3.4 is applied: asset allocations are linear in time and bounded between zero and one, and the sum of allocations is one. There are two assets and one interest rate. They are generated by a multi-variate normal distribution. The mean-vector is determined by observed long-term means of equity and bond returns and the covariance matrix is estimated from yearly observations of the Russell 3000 index and the 10-years US Treasury bond (used for bond returns and interest rate). Given these assets we can optimize the replacement ratio in a mean-variance way to obtain the optimal linear allocation




Figure 3.4: Optimal equity asset allocation for linear asset allocation approximation to the mean-variance problem, as described in Section 3.4.1. Asset returns are normally distributed, ingoing cash flows are as in base case, described in Section 3.4.2. Legend indicates minimum value for expected replacement ratio.

towards equity.

The optimal equity allocation for the problem defined by equation (3.14), for 3 different levels of minimum expected RR, is shown in Figure 3.4. The red line in Figure 3.4 is an equity-only allocation, this allocation corresponds to a high minimum expected RR. The obtain this high RR, the investments should be risky, and hence we see an equity-only allocation. The green line in Figure 3.4 corresponds to a more moderate minimum expected RR and shows the classic life-cycle property: equity allocation, and therefore risk, is decreasing in time. Finally, the blue line in Figure 3.4 corresponds to the lowest level of minimum expected return and it characterizes the minimum-variance portfolio. This minimum-variance portfolio is not a portfolio of just bonds but also mixes in equity towards the end. This makes sense as we have seen in Chapter 2 that increasing risk at the end is mean-variance efficient.

Now instead of the base case cash flows, which grow over time, we optimize with a constant cash flow, i.e., $C(t)$ as given by equation (3.15) is constant. When the total of cash flows is spread equally over all time periods, results do not change notably. The resulting graph looks familiar to Figure 3.4, but with a smaller proportion of equity, the gradients are smaller for the less risky strategies. Smaller risk at the start makes sense as we have larger cash flows up front, which causes the variance to be relatively larger at the end. We can conclude that the cash flow pattern has little influence on the optimization result.

To explain the result as presented in Figure 3.4, we find the minimum-variance point for the one-period problem, i.e., no rebalancing is allowed. For the one-period problem, we find that the minimum-variance point is at an equity allocation of 0.07. The reason for this small equity allocation is the influence of interest rates on the replacement ratio. As described in Section 3.3.1, bond returns and the replacement ratio are negatively correlated, which is why we expect that as time passes, the allocation towards equity should come down in favour of bonds. The negative correlations between bond value and interest rates make it a valuable asset for reducing replacement ratio variance.

We search for the minimum variance allocation when there's no intermediate cash flows, i.e., $C(t) = 0$ for $t \neq 0$. The resulting optimal allocation for equity is $0.12 - 0.08t$, where t is from 0 to 1. As for the one-period

problem, we have that the fraction equity is low, and on top of that it is decreasing in time, exactly as expected.

We have for both the one-period and the no-intermediate-cash-flow problems that the minimum-variance asset allocation is as we expect: a low equity allocation that decreases over time. This result does not explain the minimum-variance allocation as seen in Figure 3.4, which brings us to the following conclusion.

We can conclude that there are two forces pushing the asset allocations in different directions. On the one hand, the hedging-property of bonds for the replacement ratio pushes equity allocation down in the last periods. On the other hand, the multi-period mean-variance property that prefers risky assets as the investment horizon becomes shorter, pushes the equity allocation up at the end of investment horizon, especially in the presence of intermediate cash flows. It depends on the choice of minimum required expected replacement ratio which effect is stronger.

3.5.2. Individual weight optimization

To gauge the effect of changing cash flows and that of the interest rate in the last period we reduce the investment horizon so that individual weight optimization is possible. Using individual weight optimization, we first show that the absolute size of cash flows does not matter for the optimal asset allocation. Next, we show that timing of the cash flows into the portfolio is important, the shape of the cash flow profile does have influence on the optimization result. Finally, we explain the results in Figure 3.4 and Section 3.4 by comparing the effect of cash flows and interest rates on individual asset weights.

The setting in this section is equal to the previous section with one exception: we keep the wage constant, which is a simplification. The same simplification applies to the franchise (F), which is a sum of money for retirement income the investor gets regardless of how much he saved for his retirement. This franchise is important for the cash flows going into the pension portfolio. These are calculated as a fraction of the wage minus franchise (the pensionable income). To check the effect of the size of cash flows on the optimization we vary the fraction of pensionable income (f) invested in the portfolio. For one period we obtain:

$$R_1 = w_1 R_{eq} + (1 - w_1) R_{bo}, \quad W = (I - F)f(1 + R_1), \quad (3.16)$$

$$DF = \sum_{n=1}^{20} \frac{1}{(1 + r)^n}, \quad RR = \frac{\frac{W}{DF} + F}{I}. \quad (3.17)$$

The representation in equation (3.17) is how we implement the optimization in Python; R_{eq} , R_{bo} and r (return on equity, bonds and the interest rate) get drawn from a multivariate normal 10000 times, RR is calculated and we minimize $\text{Var}(RR)$ over w_1 .

Effect of size of cash flows

Varying the fraction invested f shows that the optimization result is independent of the magnitude of the cash flow, see Table 3.3. The franchise has a major effect on the expected RR as we only look at one period, increasing the cash flow tenfold only increases the expected RR a fraction, hence we also present results for a zero franchise. The fraction invested in equity is the same in every situation, even for a tenfold increase of RR , which means an increase of a factor 100 in variance. Hence we can safely conclude that the magnitude of wealth invested is not of importance to the asset weightings (as expected with normal returns).

All other results presented in this section have been computed for the different cash flows invested and different franchises but as results are the same in each case we only present the case $f = 0.1$ and zero franchise.

Effect of cash flow pattern

Adding a second period to the investment horizon but *no* intermediate cash flow only changes the formulation of wealth, given by equation (3.16), as follows:

$$R_2 = w_2 R_{eq} + (1 - w_2) R_{bo}, \quad (3.18)$$

$$W = (I - F)f(1 + R_1)(1 + R_2), \quad (3.19)$$

with all other parameters unchanged. Note that R_1 and R_2 are independent, meaning that the equity and bond returns are independent for every period. Minimizing variance of the RR yields $w_1 = 0.383$, $w_2 = 0.060$ and $\mathbb{E}[RR] = 0.0083$. w_2 is of the same order as the single weight for the one period problem, see Table 3.3. This is for the same reason as before: the hedging-property of bonds. The drop in equity allocation is only present

Table 3.3: Mean-variance optimal equity allocation (w) and expected replacement ratio (RR) for different fractions of income (f) invested, all as in equations (3.16) and (3.17).

	wage: 45000, franchise: 15000		wage: 45000, franchise: 0	
f	w	RR	w	RR
0.05	0.0673	0.3359	0.0673	0.0039
0.10	0.0673	0.3386	0.0673	0.0079
0.15	0.0673	0.3411	0.0673	0.0118
0.20	0.0673	0.3438	0.0673	0.0157
0.25	0.0673	0.3464	0.0673	0.0196
0.30	0.0673	0.3490	0.0673	0.0236
0.35	0.0673	0.3517	0.0673	0.0275
0.40	0.0673	0.3543	0.0673	0.0314
0.45	0.0673	0.3569	0.0673	0.0354
0.50	0.0673	0.3595	0.0673	0.0393

in the second weight as, by independence of returns, only the interest rates in the last period matter. The first weight is bigger, but at 0.383 it is still lower than the minimum-variance point of the linear combination of equity and bonds, which is 0.420. Note that the differences are not caused by the randomness in R_{eq} etc., we use the same random sample for all results.

By adding a second cash flow the formulation for W changes:

$$W = (I - F)f(1 + R_1)(1 + R_2) + (I - F)f(1 + R_2). \quad (3.20)$$

The minimum-variance parameters are $w_1 = 0.350$, $w_2 = 0.0646$, $\mathbb{E}[RR] = 0.0161$. Adding the extra cash flow more than doubles the RR, but more notably: w_1 has decreased, where w_2 has increased. For comparison; when optimizing the one-period problem, given by equations (3.16) and (3.17), with the data used to generate R_2 , the optimum equity weight is 0.0645, hence the size of w_2 makes sense. The questions that remain are: why was w_2 not already higher in the situation with only one cash flow and why does w_1 reduce when we add a second cash flow.

The result may become clearer when we add a third period and remove the second cash flow:

$$R_3 = w_3 R_{eq} + (1 - w_3) R_{bo} \quad (3.21)$$

$$W = (I - F)f(1 + R_1)(1 + R_2)(1 + R_3) \quad (3.22)$$

The minimum-variance parameters are $w_1 = 0.377$, $w_2 = 0.386$, $w_3 = 0.0460$ and $\mathbb{E}[RR] = 0.0088$. The first conclusion from these results is that the last period will always have little equity because of the correlation between interest rates and bond returns, which only matters in the last period as periods are independent. The last-period weight w_3 is lower than for shorter investment horizons, which is caused by increasing the investment horizon. This generates more return and hence more risk-aversion, causing a smaller risky allocation. This could also hold for w_1 , it is lower than w_2 and also lower than in the two-period case. This is speculation, however, as the difference is relatively small. Let's examine the effect of adding a second and third cash flow to the 3-period problem:

$$W = (I - F)f(1 + R_1)(1 + R_2)(1 + R_3) + (I - F)f(1 + R_2)(1 + R_3) + (I - F)f(1 + R_3) \quad (3.23)$$

Minimization of variance yields: $w_1 = 0.298$, $w_2 = 0.378$, $w_3 = 0.0594$, $\mathbb{E}[RR] = 0.0253$. Adding a third cash flow has exactly the same effect as adding a second cash flow earlier. It increases w_3 while it decreases w_1 and w_2 . To check this we add a fourth period and fourth cash flow. The results: $w_1 = 0.179$, $w_2 = 0.325$, $w_3 = 0.373$, $w_4 = 0.0550$, $\mathbb{E}[RR] = 0.0341$. The decrease in w_1 is remarkable, but w_2 and w_3 are also small compared to what we saw for the last and second to last cash flow in the other cases. Even w_4 is lower than the last period return in every other case. All results are combined in Table 3.4. We can conclude that adding cash flows in later periods, decreases risk-taking in the first periods. In the next section, we explain this result.

Equity allocation: increasing or decreasing in time?

Table 3.4 is explanatory for the results in Figure 3.4 and in Section 3.4; equity allocation increases in every period before the last period and sharply decreases in the last period. If linear regression would be performed, it really depends on the exact value of the weights whether this line will slope upwards or downwards.

Table 3.4: Optimal equity allocation and expected replacement ratio for different cases.

Situation	w_1	w_2	w_3	w_4	$\mathbb{E}[RR]$
One period, one cash flow	0.067				0.0079
Two periods, one cash flow	0.383	0.060			0.0083
Two periods, two cash flows	0.350	0.065			0.0161
Three periods, one cash flow	0.377	0.386	0.046		0.0088
Three periods, three cash flows	0.298	0.378	0.059		0.0253
Four periods, four cash flows	0.179	0.325	0.373	0.055	0.0341

The upward part of the equity allocation can be explained by what was explained in Chapter 2; multiple cash flows flowing into the portfolio during the investment period cause an increasing risky asset allocation over time.

The downward part is explained by the negative correlation between bond returns and interest rates and hence the hedging property of bonds for the replacement ratio. In the scenario where returns and interest rates are independent between periods this only has an effect in the last period, but the scenarios used in the base case set out in Section 3.4 have dependence between periods, and hence the effect will be larger there.

What all previous examples show, from this section and from Section 3.4, is that there is a trade-off in life-cycle glide path design. This trade-off is between the mean-variance property described in Chapter 2, risk-taking increases as the investment horizon decreases, and the hedging property of bonds for the replacement ratio. The mean-variance property causes risky assets to have more presence in the portfolio as time goes by, the hedging property of bonds causes exactly the opposite. The hedging property of bonds supports traditional financial planning; as one gets older, a bigger share of capital should be invested in riskless assets. The mean-variance property, however, is at odds with traditional financial planning. This result is consistent with literature: most papers try to find justification for the traditional approach, but more recently papers disputing this approach have been published. Our research shows that both sides have their merit, but we do not yet have an answer to which approach is better.

3.6. Conclusion

This chapter set out the main problem for this research, that of optimizing the replacement ratio: pension income as a fraction of lifetime average income. To this end, we tried to find the optimal glide path: a deterministic function of time for the asset allocation. We chose mean-variance optimization as the method for finding such a glide path. Under the assumption of linear allocation over time, we optimized using simulations. The results were surprising: instead of a decreasing allocation towards risky assets, we often found the opposite: the optimal allocation becomes more risky as time goes by. Although at odds with conventional wisdom, this result is consistent with the findings in Chapter 2.

To gain some insight, we restricted the optimization and used a basic model for asset returns and interest rates. Asset returns and interest rates were generated using a multivariate normal distribution. Repeating the optimization showed the same results as for the realistic case. Restricting the problem even more so that individual weight optimization was possible provided clarification. It became clear why equity (or other more risky assets) allocation is sloping upward through time for some optimization parameters, and downward for others. There are two main drivers for this. The first is the property of the mean-variance optimization that the optimal portfolio is more risky for shorter horizons. The second is the negative correlation between bonds and interest rates: the hedging property of bonds. These two have opposite effects: the mean-variance property causes risky asset allocation to increase as time goes by, i.e. the investment horizon becomes shorter, whereas the hedging property of bonds causes risky assets allocation to decrease as time goes by. This contrast is the main result of this chapter: mean-variance optimization does not adhere to conventional life-cycle investing wisdom.

The conclusion of this chapter raises a question: is mean-variance optimization suitable for glide path construction. The answer is partly found in the literature, from which it is clear that standard utility approaches do also not yield desirable results unless concepts as human capital are introduced. Still, mean-variance optimization is the standard in portfolio optimization. A promising alternative, however, exists: dynamic mean-variance optimization. When using dynamic mean-variance optimization the asset allocation is no longer deterministic, it becomes dynamic. By using a dynamic asset allocation, we can adapt to

changes over the investment horizon. For example, when very high returns have been realized, we can afford to invest less risky afterwards. This can be done using a dynamic asset allocation, not with a deterministic one. Being able to adapt to changing market circumstances can be beneficial in life-cycle investing.

In the next chapter we will work out how dynamic mean-variance optimization can be used for life-cycle investing.

3.A. Appendix to Chapter 3




Figure 3.5: Asset allocation over time and replacement ratio distribution for minimization of variance with a minimum level of replacement ratio of 1.3

Table 3.5: Pension premium for every age

Age	Premium
21	8%
22	8%
23	8%
24	8%
25	9,30%
26	9,30%
27	9,30%
28	9,30%
29	9,30%
30	10,80%
31	10,80%
32	10,80%
33	10,80%
34	10,80%
35	12,50%
36	12,50%
37	12,50%
38	12,50%
39	12,50%
40	14,60%
41	14,60%
42	14,60%
43	14,60%
44	14,60%
45	17,00%
46	17,00%
47	17,00%
48	17,00%
49	17,00%
50	19,80%
51	19,80%
52	19,80%
53	19,80%
54	19,80%
55	23,30%
56	23,30%
57	23,30%
58	23,30%
59	23,30%
60	27,70%
61	27,70%
62	27,70%
63	27,70%
64	27,70%
65	31,50%
66	31,50%




Figure 3.6: Asset allocation over time and replacement ratio distribution for minimization of variance with a minimum level of replacement ratio of 1.1

4

Dynamic mean-variance optimization

In Chapter 3, we found that using multi-period mean-variance optimization for deterministic asset allocations does not yield glide paths consistent with conventional investing wisdom. In this chapter we introduce an alternative that does yield results consistent with conventional investing wisdom: dynamic mean-variance optimization.

Dynamic mean-variance (DMV) optimization is multi-period mean-variance optimization where at every rebalancing moment the information of what happened before is used to optimize the asset allocation. Optimization is done conditional on the past. We show how we can use DMV for pension portfolios.

4.1. Dynamic versus deterministic asset allocation

The DMV problem related to life-cycle investing is profoundly different from the deterministic problem. The main difference is that the asset allocation decision at time $t + 1$ depends on what happened between t and $t + 1$. The solution to the deterministic problem is fixed, whatever happens during the investment horizon, the asset allocation has already been decided at time $t = 0$ and will not change. The dynamic solution, however, dynamically changes depending on the course of events, e.g., the stochastic accumulated wealth influences the asset allocation.

For life-cycle pension plans the current standard is a deterministic asset allocation, often called a glide path. As we have seen in Chapter 3, however, optimizing with the deterministic definition does not lead to consistent results. Furthermore, Forsyth, Li and Vetzal show in [20] that Target Date Funds, which have a deterministic asset allocation (glide path), underperform adaptive strategies (obtained by dynamic optimization) and even constant allocation strategies. Because of these convincing arguments we investigate how DMV optimization can be used for pension portfolios.

4.2. Definition of the dynamic mean-variance problem

The dynamic mean-variance problem for wealth is given by:

Problem 4.1. The asset allocation $\{\tilde{x}_s\}_{s \in \{t, t+1, \dots, T-1\}}$, for each time $t \in \{0, 1, \dots, T-1\}$, with the optimal dynamic trade-off between risk and return of wealth satisfies:

$$\begin{aligned} \max_{\{\tilde{x}_s\}_{s \in \{t, \dots, T-1\}}} \quad & \mathbb{E}[W_T | Z_t] - \lambda \text{Var}[W_T | Z_t] \\ \text{s.t.} \quad & \sum_{i=1}^M x_{s,i} = 1, \quad \forall s \in \{t, \dots, T-1\} \\ & \tilde{x}_s \in [0, 1]^M, \quad \forall s \in \{t, \dots, T-1\}, \end{aligned} \tag{4.1}$$

where Z_t denotes the (vector of) state variable(s) at time t and all other variables are as in Problem 3.1.

Examples of state variables are the wealth at time t and the returns on assets in the period before t . It is possible to add other state variables, but they have to be related to future portfolio returns to be useful.

Problem 4.1 is as before; we aim to maximize expected return minus a constant times the variance. The constraints are that allocation to any asset has to be non-negative and the sum of investment fractions should

be one. Both constraints can be relaxed if desired.

The main difficulty in solving the DMV optimization problem is the non-linearity of conditional variance. The variance operator does not possess the tower property:

Lemma 4.1 (Tower property). *Let X be a random variable and $\mathcal{H} \subset \mathcal{G}$ then $\mathbb{E}[\mathbb{E}[X|\mathcal{G}]|\mathcal{H}] = \mathbb{E}[X|\mathcal{H}]$.*

Proof. See [49], page 90. □

Lemma 4.1 only applies to conditional expectation, not conditional variance. Because of this, it is not possible to write Problem 4.1 in a recursive way and hence the standard (stochastic) dynamic programming approach is not applicable:

Lemma 4.2 (Conditional variance has no recursive representation). *Let $X_i = X_{i-1} + \epsilon_i$ where $\epsilon \sim \mathcal{N}(0, \sigma^2)$ and independent of all other terms. Then conditional expectation has a recursive representation: $J_n(X_n) = \mathbb{E}[X_N|X_n]$, and we can write $J_n(X_n) = \mathbb{E}[J_{n+1}(X_{n+1})|X_n]$ where $J_N(X_N) = X_N$. Conditional variance, however, does not have a recursive representation.*

Proof. By Lemma 4.1, $\mathbb{E}[\mathbb{E}[X_N|X_{n+i}]|X_n] = \mathbb{E}[X_N|X_n]$ for all $i > 0$. For conditional variance we have $J_n(X_n) = \text{Var}[X_N|X_n] = (N-n)\sigma^2$. We define $J_n(X_n) = \text{Var}[J_{n+1}(X_{n+1})|X_n]$ with $J_N = X_N$. Then $J_{N-1}(X_{N-1}) = \text{Var}[J_N(X_N)|X_{N-1}] = \text{Var}[X_N|X_{N-1}] = \sigma^2 \neq (N-n)\sigma^2$, which concludes the proof. □

Lemma 4.2 shows that conditional variance has no recursive representation and hence we cannot recursively solve Problem 4.1. Nevertheless, a recursive alternative for the DMV optimization problem does exist; a second moment optimization, called the pre-commitment problem, given by:

Problem 4.2 (Pre-commitment problem). The optimal asset allocation $\{\tilde{x}_s\}_{s \in \{t, t+1, \dots, T-1\}}$, for each time $t \in \{0, 1, \dots, T-1\}$, satisfies:

$$\begin{aligned} \min_{\{\tilde{x}_s\}_{s \in \{t, \dots, T-1\}}} \quad & \mathbb{E}[(W_T - \gamma)^2 | Z_t] \\ \text{s.t.} \quad & \sum_{i=1}^M x_{s,i} = 1, \quad \forall s \in \{t, \dots, T-1\} \\ & \tilde{x}_s \in [0, 1]^M, \quad \forall s \in \{t, \dots, T-1\}, \end{aligned} \tag{4.2}$$

where Z_t denotes the (vector of) state variable(s) at time t , γ specifies a wealth-target and all other variables are as in Problem 3.1.

Li and Ng [27] showed that Problem 4.1 is equivalent to Problem 4.2 when

$$\gamma = \frac{1}{2\lambda} + \mathbb{E}[W_T^* | Z_t], \tag{4.3}$$

where W_T^* denotes the final wealth if wealth is invested according to the optimal strategy.

The relation between risk-aversion λ in Problem 4.1 and the target γ in Problem 4.2 given by equation (4.3) describes an important property of the pre-commitment problem. The relation between λ and γ is clear: when risk-aversion λ is large, the target γ will be lower, this will result in a less risky allocation as a lower target requires less return and therefore less risk. The second element in equation (4.3) is the conditional expectation of optimal wealth. The conditional part is crucial: when at time t the wealth is high, the conditional expectation of final wealth will be high, and hence we increase the target γ . By adapting the target to the circumstances, the riskiness of the asset allocation will not change much over time. So, in Problem 4.1 we have constant risk-aversion, and therefore changes in the riskiness of the asset allocation are small. In Problem 4.2, we have a constant target, and therefore riskiness of the asset allocation changes when the distance to target changes. Problem 4.2, the pre-commitment problem, is a *target-based* problem.

The representation in Problem 4.2 is useful because it allows us to write down the problem in a recursive fashion. Define $V_T = (RR - \gamma)^2$ and remember that RR is a function of W_T (wealth at the final time) and the term-structure at time T . As the control variables (asset allocation) have no influence on the term-structure (Note: this may be invalid if a fund is able to invest market moving sums of money), we can state that V_T is

a function of W_t , the only variable dependent on the state and control variables. The optimization problem can then be defined as:

$$V_t(W_t) = \min_{\{\vec{x}_s\}_{s \in [t, \dots, T-1]}} \mathbb{E}[(RR - \gamma)^2 | Z_t], \quad (4.4)$$

where W_t is part of Z_t . With the definition of V_T and Lemma 4.1, we can rewrite (4.4) in a recursive fashion:

$$V_t(W_t) = \min_{\vec{x}_t} \mathbb{E}[V_{t+1}(W_{t+1}) | Z_t]. \quad (4.5)$$

The recursive problem definition allows us to solve the problem analytically in some cases, and numerically efficient in others, as we will see in the next sections.

4.3. Analytic solutions to the dynamic mean-variance and pre-commitment problem

The literature on analytic solutions to the DMV problem is pointed towards a couple of specific cases. First of all there is a distinction between the continuous-time and discrete-time case. The discrete case is under consideration in this thesis, but the continuous-time case is worth studying as there are solutions available to constrained problems. In discrete-time, analytic solutions only exist for the unconstrained case.

The other distinction is between Problem 4.1 and Problem 4.2. Problem 4.1 is referred to as the time-consistent problem and Problem 4.2 is called the pre-commitment problem. These solutions are in general different from each other, which seems to contradict the statement that both problems are equivalent. There is no contradiction however, as the problems are equivalent when $\gamma = \frac{1}{2\lambda} + \mathbb{E}[W_T^* | Z_t]$. Here γ depends on the state variables and can therefore take different values for the same λ .

4.3.1. Continuous-time solutions

As this thesis is concerned with the discrete-time problem, we only briefly discuss the continuous-time literature. For the unconstrained problem, Bajeux-Besnainou and Portait [1] employ the martingale method to calculate the dynamic mean-variance efficient frontier. In the presence of a risk-free asset this frontier is a straight line. Zhao and Ziemba [31] use the martingale method to solve the DMV problem, and compare it to the corresponding expected utility problem. The idea of the martingale method is finding an optimal wealth at time T , at a maximum cost of wealth at time zero. Given this optimal wealth at T , one finds an allocation strategy that replicates this optimal wealth process. For this method to work, a complete market is necessary. In a complete market, there exists an equilibrium price for every asset in every possible state of the world [35]. When we constrain the asset allocation, the market is possibly no longer complete.

Bielecki et al. [5] are able to solve a constrained DMV problem in continuous-time. The no-bankruptcy constraint is introduced: the optimal strategy may not allow the wealth to go below zero. They solve this problem as a variance minimization problem, where Lagrange multipliers are used to find optimal portfolios. As in the martingale approach they then use these portfolios to replicate the asset allocation strategy.

All methods above use variance minimization for a given level of return, which is not exactly as in Problem 4.1 or 4.2. The unconstrained pre-commitment problem, in older literature referred to as the 'embedded' linear-quadratic (LQ) problem, has been solved by Zhou and Li [50]. They show that the DMV problem is a stochastic LQ problem and that this can be solved by the solutions to a set of Riccati differential equations. Lim and Zhou extend this work to a problem with random parameters, where they show that the stochastic Riccati equations are non-linear backward stochastic differential equations. Exploiting the structure of the Riccati equation following from the mean-variance problem, they are able to find solutions [29]. Finally, Lim et al. solve the LQ problem with no-shorting constraints. The Riccati equations approach does no longer work, so they fall back to the Hamilton-Jacobi-Bellman equation, a standard approach for control problems [28].

This summarizes the analytic solutions to the continuous-time DMV problem. There are more solutions for problems with different constraints, but all methods rely on restricting assumption, such as market completeness, which makes the continuous-time methods unsuitable for our purpose.

4.3.2. Discrete-time solutions

In discrete-time the distinction between time-consistent and pre-commitment strategies is also present. The pre-commitment problem and its solution in discrete time were developed by Li and Ng in [27]. They

first prove equivalence between the time-consistent and pre-commitment problem when $\gamma = \frac{1}{2\lambda} + \mathbb{E}[W_T^*|Z_t]$. Afterwards they solve the pre-commitment problem for asset returns independent through time and unconstrained asset allocation. Their solution is elegant but involved. Essentially the optimal asset allocation consists of two parts. The first is depending on the current wealth and the expectation of returns and their (co)variance in the following period. The second part is involved and consists of the expectation and (co)variance of returns in all the periods after the next, this part however does not depend on current wealth.

The first solution to the time-consistent problem was introduced by Basak and Chabakauri [4]. They are able to define Problem 4.1 as a recursive problem, by adding an adjustment term:

$$U_t = \mathbb{E}[W_T|Z_t] - \lambda \text{Var}[W_T|Z_t], \quad (4.6)$$

$$= \mathbb{E}[U_{t+\tau}|Z_t] - \lambda \text{Var}[\mathbb{E}[W_T|Z_{t+\tau}]|Z_t]. \quad (4.7)$$

By splitting the problem, they are able to optimize for the wealth part, the only part depending on the control and the part for which they can write down an explicit Hamilton-Jacobi-Bellman equation. The solution to this equation does not depend on the current wealth. This last property is desirable for life-cycles, as this is a deterministic solution. When applied to models of asset returns, however, it turns out that the optimal allocation does depend on the current state of the asset price and the optimal allocation is decreasing in time horizon $(T - t)$, or increasing in time (t) . This result supports the findings from Chapter 3.

There are other analytic solutions for the discrete-time problem proposed in literature, they however only show different forms of the two mentioned above. An exception is the paper by Cong and Oosterlee [12], where a numerical method for the problem is proposed. For the unconstrained problem, their strategy turns out to yield an analytical solution to the pre-commitment problem. As this numerical approach is applicable to the subject of this thesis, we will elaborate on this method in the next section.

It is clear that analytic solutions do not suffice for our problem. We have not found any analytic solutions for the constrained discrete-time problem. Although in continuous-time some solutions to the constrained problem exist, they all make assumptions on asset returns or other variables that are too restrictive for our purpose. We therefore focus on numerical solutions.

4.4. Numerical solutions to the dynamic-mean variance problem

The methods for obtaining numerical solutions consist of two broad categories. One entails discretizing the Hamilton-Jacobi-Bellman partial differential equation, which can then be solved numerically. The other category is that of (Monte-Carlo) simulations, which can then be used to maximize (conditional) expectations. The latter category seems applicable to our problem; we aim to create a general framework for optimization, and as simulations are almost always available, this category seems appropriate. We will give a short overview of work done in both categories and propose a method suitable to the DMV problem for the replacement ratio.

4.4.1. Discretizing partial differential equations

The use of the Hamilton-Jacobi-Bellman (HJB) partial differential equation (PDE) in portfolio optimization has been in practice for a long time. The main application has however been to utility functions that do not, like mean-variance, depend on the current state of the wealth. An example is given by Brennan et al. [10], who discretize the state-space of 3 state variables and solve the PDE on this grid for each time-step. They repeat this process to update the asset allocation.

More recently, the literature has focused on the DMV problem, specifically the pre-commitment problem. Wang and Forsyth [47] solve the continuous-time pre-commitment problem with different types of constraints, discretizing the HJB PDE that corresponds to the pre-commitment problem. The same authors show in [48] that the constrained time-consistent problem can be solved using the HJB PDE.

Although the HJB PDE approach is promising for the DMV problem, it has one major drawback: the curse of dimensionality. In the case where there is just one risky asset, or where one state variable drives the entire wealth process, the approach is accurate and fast. As more state variables are added, however, the computation time increases exponentially. As we wish to solve the problem for multiple asset with complicated dynamics, the HJB PDE approach is not suitable. Simulation-based methods could be a solution.

4.4.2. Simulation-based methods

Simulation-based methods are popular in portfolio selection. Not just in academics, but also in practice. See for example Vanguard [44] and TIAA-CREF [41]. Although these reports mention they use simulation for glide path design, they do not describe the exact methodology. We therefore focus on the academic literature.

Brandt et al. [9] introduced the first general simulation-based method for portfolio selection. Their method is implemented for CRRA utility, which simplifies the problem as the results are no longer path-dependent. They do however propose a method applicable to general utility functions, and hence applicable to the pre-commitment problem. Their method consists of applying a Taylor expansion of the utility function, simulating paths and then for each sample path maximizing the utility function, by solving the first-order conditions. This is done backwards in time.

There exist some extension to the method of Brandt et al., see for example van Binsbergen & Brandt [43] and Garlappi & Skoulakis [22]. The first uses grid-searching to optimize asset allocation, instead of solving the first-order conditions. Both papers discuss the way to evolve information through time steps. van Binsbergen and Brandt show that portfolio weights are best used to evolve information, while Garlappi and Skoulakis propose the use of a return measure, the certainty equivalent.

These method work well for CRRA utility, but are more involved when the CRRA-condition is not satisfied and the problem becomes wealth-dependent.

Cong and Oosterlee [12] proposed an alternative. They solve the problem of unknown wealth when solving backward in time by starting with a forward-in-time solution and updating this solution with backward iterations. The method is developed for the pre-commitment problem and can incorporate constraints. The method has been implemented for 1D and 2D problems and implementation on higher dimensional problems seems possible.

Cong and Oosterlee link their method to time-consistency in [13] and extend the algorithm to solve time-consistent dynamic mean-variance problems. As Basak and Chabakauri have shown that time-consistency does not necessarily lead to attractive asset allocations, we aim to solve the pre-commitment problem.

The methods of Brandt et al. [9] and Cong & Oosterlee [12] are applicable to Problem 4.2, so we will set out how they work and what their characteristics are. These and other simulation-based methods assume, however, for simplicity, the existence of a risk-free asset. Therefore we define the following problem, which is a variation on Problem 4.2:

Problem 4.3 (Pre-commitment problem with a risk-free asset). The optimal asset allocation $\{\tilde{x}_s\}_{s \in \{t, t+1, \dots, T-1\}}$, for each time $t \in \{0, 1, \dots, T-1\}$, satisfies:

$$\begin{aligned} \min_{\{\tilde{x}_s\}_{s \in \{t, \dots, T-1\}}} \quad & \mathbb{E}[(W_T - \gamma)^2 | Z_t] \\ \text{s.t.} \quad & \sum_{i=1}^M x_{s,i} \leq 1, \quad \forall s \in \{t, \dots, T-1\} \\ & \tilde{x}_s \in [0, 1]^M, \quad \forall s \in \{t, \dots, T-1\}, \end{aligned} \quad (4.8)$$

where Z_t denotes the (vector of) state variable(s) at time t , γ specifies a wealth-target and wealth at time t is defined by:

$$W_t = W_{t-1}(\tilde{x}_{t-1} \cdot \tilde{R}_{t-1} + R_f), \quad (4.9)$$

where $\tilde{R}_{t-1}$ represents the excess returns over the risk-free. All other variables are as in Problem 3.1.

The algorithm by Brandt et al. solves a more general version of Problem 4.3:

Algorithm 1 (Dynamic portfolio optimization algorithm by Brandt et al.). The following algorithm maximizes the expected utility of final wealth, $\mathbb{E}[u(W_T)]$, for a portfolio with dynamic asset allocation.

1. Expand the value function, the conditional expectation of utility: $V_t = \mathbb{E}[u(W_T) | Z_t]$, where $u(\cdot)$ is the utility function, into a Taylor series.
2. Simulate paths. This is straightforward sampling from a known distribution, resampling or any other method of generating samples.
3. Maximize the value function for each sample path, recursively, backwards in time. In order to do this, the expanded value function is maximized by solving the first-order conditions. To solve these first-order conditions, estimate the conditional expectations of the value function's derivatives and the asset returns.

In step 1 of Algorithm 1, Brandt et al. expand the CRRA utility function in a Taylor series, which is efficient as it makes the optimization independent of wealth. We elaborate on step 3 of Algorithm 1 in the following lemma:

Lemma 4.3. *The argument of the maximum of the second-order Taylor expansion of the value function as defined in step 1 of Algorithm 1 is defined by conditional expectations of the first and second derivative terms of value function's Taylor expansion.*

Proof. Observe that $\max_{\{\tilde{x}_s\}_{s \in \{t, \dots, T-1\}}} V_t = \max_{\{\tilde{x}_s\}_{s \in \{t, \dots, T-1\}}} \mathbb{E}[u(W_T)|Z_t]$ can be defined recursively as:

$$V_t(W_t) = \max_{\tilde{x}_t} \mathbb{E}[V_{t+1}(W_{t+1})|Z_t], \quad (4.10)$$

where $V_T = u(W_T)$. Plugging in wealth as defined in equation (4.9) into the value function as defined in equation (4.10) and expanding around $W_t R_f$ yields:

$$V_t(W_t) = \max_{\tilde{x}_t} \mathbb{E}[V_{t+1}(W_t R_f) + \frac{\partial V_{t+1}(W_t R_f)}{\partial W_{t+1}} (W_t \tilde{x}_t \tilde{R}_{t+1}) + \frac{1}{2} \frac{\partial^2 V_{t+1}(W_t R_f)}{\partial W_{t+1}^2} (W_t \tilde{x}_t \tilde{R}_{t+1})^2 | Z_t]. \quad (4.11)$$

Solving the first-order conditions of equation (4.11) yields:

$$\tilde{x}_t = -\mathbb{E} \left[\frac{\partial V_{t+1}(W_t R_f)}{\partial W_{t+1}} (W_t \tilde{R}_{t+1}) \middle| Z_t \right] \mathbb{E} \left[\frac{\partial^2 V_{t+1}(W_t R_f)}{\partial W_{t+1}^2} (W_t \tilde{R}_{t+1})^2 \middle| Z_t \right]^{-1}. \quad (4.12)$$

So, by approximating the conditional expectation in the numerator and denominator of equation (4.12), we have maximized the Taylor-expansion of the value function. $\square$

The conditional expectations that need to be determined in order to calculate the solution are obtained through *regression*. Calculating conditional expectations through regressions is straightforward. Suppose we want the conditional expectation of the asset returns:

$$\mathbb{E}[R_{t+1}|Z_t] = \phi(Z_t)' \theta_t, \quad (4.13)$$

where $\phi(Z_t)$ are basis functions of Z_t , for example $\phi(Z_t) = [1, Z_t, Z_t^2]'$. θ_t are the parameters that characterize the conditional expectation and need to be determined through regression. As many paths are available, θ_t can be estimated by regressing observations of R_{t+1} on the corresponding observations of $\phi(Z_t)$, this is called *cross-path regression*. Brandt et al. briefly discuss the choice of basis functions for $\phi(\cdot)$ and conclude that a polynomial basis suffices. From their numeric results it follows that a linear function is often sufficient, using quadratic functions is only useful on long horizons.

This completes the algorithm for unconstrained CRRA utility optimization. For constrained problems, Brandt et al. refer to methods able to solve high-dimensional constrained optimization problems. These methods can be used to incorporate the constraints when solving the problem given by equation (4.11).

For non-CRRA utility, among which the pre-commitment problem, the method is less straightforward. As can be seen in equation (4.12), the solution will depend on W_t . In the CRRA case the size of W_t does not matter and can hence be set to one. For other cases, W_t does matter, and as the solution is backwards in time, W_t is not known when solving for $\tilde{x}_t$. Brandt et al. [9] propose using a grid of wealth levels. For each wealth level they solve the problem by interpolating the portfolio choices found for every grid level at every future date. This is cumbersome and causes extra complications as choosing a proper grid and reduction of computational speed. The algorithm by Cong and Oosterlee [12] does not rely on a grid and is designed to solve the pre-commitment problem. It is therefore more suited to our purpose and introduced next.

Mean-variance portfolio optimization algorithm by Cong and Oosterlee

The algorithm by Cong and Oosterlee [12] is designed for the pre-commitment problem, Problem 4.3, and consists of two parts.

The first part of the method is a forward algorithm called the multi-stage strategy. This strategy aims to solve the following for every time $t \in \{0, \dots, T-1\}$:

$$\min_{\tilde{x}_t} \mathbb{E} \left[\left(W_t (\tilde{x}_t \cdot \tilde{R}_t + R_f) \prod_{s=t+1}^{T-1} (R_f) - \gamma \right)^2 \middle| Z_t \right]. \quad (4.14)$$

As this is solved in a forward manner, W_t is known and we can solve the first-order conditions. In equation (4.14), the wealth is assumed to be invested in the risk-free from $t + 1$ until the end of the investment horizon. This is different from the pre-commitment definition but it turns out that in the unconstrained case, the multi-stage solution is the same as the pre-commitment solution. In the constrained case, the solutions are not equivalent. Hence the algorithm has a second part: a backward algorithm used to update the forward solution:

Algorithm 2. This algorithm solves Problem 4.3 and consists of 4 steps, which can be repeated, so that we iterate to the optimal solution. The four steps of the algorithm are:

1. Generate the wealth values using an initial guess for the asset allocation and the asset returns from the simulated paths. The initial guess is the result of the forward step. Steps 2 to 4 are performed backward in time, starting at $t = T - 1$ and ending at $t = 0$.
2. Determine a function f_{t+1} that maps the wealth values W_{t+1} to the value function V_{t+1} . The value function V_t is given by $\min_{\{\tilde{x}_s\}_{s \in \{t, \dots, T-1\}}} \mathbb{E}[(W_T - \gamma)^2 | Z_t]$ or recursively by $\min_{\tilde{x}_t} \mathbb{E}[V_{t+1}(W_{t+1}) | Z_t]$, where $V_T = (W_T - \gamma)^2$. The function f_{t+1} is determined by regression. Since $f_{t+1}(W_{t+1}) = f_{t+1}(W_t(\tilde{x}_t \cdot \tilde{R}_t + R_f))$ approximates V_{t+1} , we can optimize by solving the first-order conditions. This yields a new asset allocation $\hat{x}_t$. With this new values we can calculate the updated value function $\hat{V}_t$. We calculate the conditional expectation $\hat{V}_t$ through cross-path regression.
3. Using the old asset allocation, calculate the values of the previous value function $\tilde{V}_t$, again through cross-path regression. Then, for every path, if $\tilde{V}_t > \hat{V}_t$ we choose $\hat{x}_t$ as the new allocation.
4. Calculate the new value function V_t using the updated asset allocations $\tilde{x}_t$. Calculation of conditional expectation V_t is done through cross-path regression. The new values of V_t are used in step 2.

This algorithm seems suitable for the replacement ratio. For the replacement ratio the wealth process is essentially the same, there is only a factor consisting of rates and wage to take into account. In the next section we will show how this can be incorporated.

There is one drawback associated with this method. The backward algorithm contains four regression steps in every time step. Due to the recursive manner of the algorithm, errors in these regression steps can accumulate. Cong and Oosterlee use a technique called regress-later to remedy this problem. This technique consists of using known conditional expectations of the basis functions instead of calculating them from the paths. Details can be found in [11], section 3.1.

In [12], bundling is used to make the global regression in the algorithm into local regression. Bundling is done by dividing the simulated paths into equal-sized partitions, where each partition contains paths with the similar wealth values at time t . By using local regression we can obtain good regression fits, even when the function is non-smooth on the global domain.

The conditional expectations of basis functions are not always available, in the Ortec Finance scenarios for example, we do not have explicit expression for these expectations. A solution is using regress-now, in which case we regress on observed values of basis functions, instead of their conditional expectations. This is the method used by Brandt et al [9] and generally yields sufficient results.

The algorithm developed by Cong and Oosterlee seems promising for use in replacement ratio optimization. We expect that calculation of the conditional expectations can be done precise enough. We therefore introduce an adjusted version of their algorithm for the replacement ratio.

4.5. A simulation-based algorithm for mean-variance optimization of the replacement ratio

We introduce the pre-commitment problem for the replacement ratio:

Problem 4.4. The asset allocation $\{\vec{x}_s\}_{s \in \{t, t+1, \dots, T-1\}}$, for each time $t \in \{0, 1, \dots, T-1\}$, with the optimal dynamic trade-off between risk and return of the replacement ratio satisfies:

$$\min_{\{\vec{x}_s\}_{s \in \{t, \dots, T-1\}}} \mathbb{E}[(RR - \gamma)^2 | Z_t] \quad (4.15a)$$

$$\text{s.t.} \quad \sum_{i=1}^M x_{s,i} = 1, \quad \forall s \in \{t, \dots, T-1\}, \quad (4.15b)$$

$$\vec{x}_s \in [0, 1]^M, \quad \forall s \in \{t, \dots, T-1\}, \quad (4.15c)$$

where Z_t denotes the (vector of) state variable(s) at time t , γ is the replacement ratio target, RR is given by Definition 3.1 and C_r is given by equation (3.7). All other variables are as in Problem 3.1.

We introduce an algorithm to find the optimal asset allocation for Problem 4.4. For the general algorithm we assume that interest rates are stochastic and that intermediate contributions can be made. The algorithm works using sample paths of asset returns and interest rates. Every sample path should contain returns on the available assets and a term-structure of interest rates, in at least the same frequency as the rebalancing moments. An incomplete term-structure can be made complete by interpolation and extrapolation so completeness is not required, but desirable. The algorithm consists of two parts, a forward part producing a sub-optimal solution and a backward part that iteratively updates the sub-optimal solution, in order to obtain the optimal solution. We start by introducing the forward part of the algorithm.

4.5.1. Generating an initial guess: the forward algorithm

The forward algorithm solves for $t \in \{0, \dots, T-1\}$:

$$\vec{x}_t = \arg \min_{\vec{x}_t} \mathbb{E}[(W_t(\vec{x}_t^\top \vec{R}_t) + C_{t+1} - \delta_{t+1})^2 | Z_t], \quad (4.16)$$

where δ_t is an intermediate target given by:

$$\delta_t = \frac{\Gamma - \sum_{s=t+1}^{T-1} C_s \mathbb{E}[R_{T-s}^f(s)]^{T-s}}{\mathbb{E}[R_{T-t}^f(t)]^{T-t}}, \quad (4.17)$$

where $\Gamma := \gamma \bar{I} \mathbb{E}[\sum_{i=0}^{19} (R_i^f(T))^{-i}]$ is the final wealth target. $\bar{I}$ is the (deterministic) time-averaged labour income. $R_i^f(t)$ denotes the return on the (practically) riskless bond with tenor i at time t . This means that when capital is invested in this asset at time t and retained until time $t+i$, the return $R_i^f(t)$ is guaranteed. The sum $\sum_{s=t+1}^{T-1} C_s \mathbb{E}[R_{T-s}^f(s)]^{T-s}$ represents wealth from contributions if every future cash flow will be invested in the riskless asset.

Equations (4.16) and (4.17) follow from the following derivation:

The replacement ratio is given by:

$$RR = \frac{W_T}{\bar{I} \sum_{i=0}^{19} (R_i^f(T))^{-i}}, \quad (4.18)$$

where W_T is the final wealth, $\bar{I}$ is the time-average labour income and $R_i^f(t)$ denotes the return on the riskless bond with tenor i at time t .

For convenience we wish to have a deterministic wealth target, and we therefore assume we do not minimize expectation but only the quadratic difference between the replacement ratio and the target:

$$\min(RR - \gamma^2) = \min \left(\frac{W_T}{\bar{I} \sum_{i=0}^{19} (R_i^f(T))^{-i}} - \gamma \right)^2 \quad (4.19)$$

$$= \min \left(W_T - \gamma \bar{I} \sum_{i=0}^{19} (R_i^f(T))^{-i} \right)^2. \quad (4.20)$$

So we find the target wealth to be $\gamma \bar{I} \sum_{i=0}^{19} (R_i^f(T))^{-i}$, which is not deterministic as $R_i^f(T)$ can be stochastic. We therefore take expectations to obtain the deterministic wealth target $\Gamma := \gamma \bar{I} \mathbb{E}[\sum_{i=0}^{19} (R_i^f(T))^{-i}]$.

Using the deterministic wealth target Γ in the minimization yields:

$$\min_{\{\tilde{x}_s\}_{s \in \{t, \dots, T-1\}}} \mathbb{E} \left[(W_T - \Gamma)^2 \mid Z_t \right]. \quad (4.21)$$

In the forward algorithm we wish to find $\tilde{x}_t$ while moving forward through time. For every period after $t+1$ we assume the wealth to be invested in the riskless asset. We can then optimize for $\tilde{x}_t$:

$$\min_{\tilde{x}_t} \mathbb{E} \left[\left(W_t \tilde{x}_t^\top \tilde{R}_t \left(R_{T-t-1}^f(t+1) \right)^{T-t-1} - \Gamma \right)^2 \mid Z_t \right], \quad (4.22)$$

where all intermediate contributions have been set to zero. Introducing intermediate contributions is straightforward: all future intermediate contributions are assumed to be invested in the riskless asset. The optimization problem with non-zero intermediate contributions, given by C_t , is given by:

$$\min_{\tilde{x}_t} \mathbb{E} \left[\left(\left(W_t \tilde{x}_t^\top \tilde{R}_t + C_{t+1} \right) \left(R_{T-t-1}^f(t+1) \right)^{T-t-1} + \sum_{s=t+2}^{T-1} C_s (R_{T-s}^f(s))^{T-s} - \Gamma \right)^2 \mid Z_t \right]. \quad (4.23)$$

To obtain formulation (4.16), we introduce the intermediate target δ_{t+1} . Like the final wealth target Γ , we need the intermediate target to be deterministic. We follow the procedure laid out in equations (4.19) and (4.20). This gives an intermediate target

$$\frac{\Gamma - \sum_{s=t+1}^{T-1} C_s \left(R_{T-s}^f(s) \right)^{T-s}}{\left(R_{T-t}^f(t+1) \right)^{T-t}}.$$

To obtain a deterministic target, we use expectations of the return on the riskless asset, by which we obtain equation (4.17). Using the expectation of the return on the riskless asset makes the intermediate target deterministic, but it has a drawback. In a constant risk-free environment, there is no drawback as $\mathbb{E}[R_i^f(t)] = R^f$ for every t, i . In a stochastic rates environment, using expectations will hurt the quality of the solution. This issue is resolved in the backward algorithm, where assumptions like these are not made.

Finally, taking δ_{t+1} to be the target when optimizing $\tilde{x}_t$, we find equation (4.16).

Lemma 4.4. *The solution to the first-order condition for equation (4.16) is given by:*

$$\tilde{x}_t = (W_t \mathbb{E}[\tilde{R}_t \tilde{R}_t^\top \mid Z_t])^{-1} (\delta_{t+1} - C_t) \mathbb{E}[R_t \mid Z_t]. \quad (4.24)$$

Proof. Setting the first-order derivative w.r.t. $\tilde{x}_t$ of $\mathbb{E}[(W_t(\tilde{x}_t \cdot \tilde{R}_t) + C_{t+1} - \delta_{t+1})^2 \mid Z_t]$ to zero yields the desired result. $\square$

If $W_t \mathbb{E}[\tilde{R}_t \tilde{R}_t^\top \mid Z_t]$ in Lemma 4.4 is positive-definite, the quadratic program given by equation (4.16) is convex. For constrained convex quadratic programs, efficient numerical solvers exist. The existence of these solvers makes it possible to solve high-dimensional (multiple asset) problems.

The forward algorithm starts at time $t = 0$ with calculation of $\tilde{x}_0$ for every path. It will be equal for every path as W_0 is the same for every path. Then using the returns and the asset allocation we can calculate W_1 for every path. Given this wealth we can calculate $\tilde{x}_1$ for every path. Using $\tilde{x}_1$ and the returns we calculate the next wealth values and so we can continue until $t = T - 1$. This concludes the forward algorithm.

4.5.2. Updating towards the optimum: the backward algorithm

The backward algorithm consists of more steps than the forward algorithm and is an iterative algorithm, every iteration of the algorithm should improve the solution.

Algorithm 3. This algorithm solves Problem 4.4 and consists of the following steps:

1. Generate wealth values $\{W_t^i\}_{i=1}^{N_s}$ for $t \in \{0, 1, \dots, T-1, T\}$, following equation (3.2), using the returns from N_s sample paths. The asset allocation to generate these values is obtained from the forward algorithm or a previous backward algorithm iteration. Using $\{W_T^i\}_{i=1}^{N_s}$ we can calculate a set of replacement ratios $\{RR^i\}_{i=1}^{N_s}$, following equation (4.18). We then calculate $V_T = \mathbb{E}[(RR - \gamma)^2 | Z_T]$ by calculating $(RR^i - \gamma)^2$ for every path.

The following steps are performed backwards in time for every $t \in \{0, \dots, T-1\}$.

2. Determine a function f_{t+1} such that $V_{t+1} = f_{t+1}(W_{t+1})$. We can approximate such a function using regression. As we know that V_T is quadratic in wealth, we use the following regression basis: $\left[1, W_{t+1}, W_{t+1}^2\right]$. We have $\{V_{t+1}^i\}_{i=1}^{N_s}$ from step 1 (for $t = T-1$) or step 4 and $\{W_t^i\}_{i=1}^{N_s}$ has been calculated in step 1. We approximate $f_{t+1}(W_{t+1})$ using cross-path regression.

When the function $f_{t+1}(W_{t+1})$ has been learned, we can use that $W_{t+1} = W_t(\tilde{x}_t^\top \tilde{R}_t) + C_t$ to minimize f_{t+1} w.r.t $\tilde{x}_t$, for every path:

$$f_{t+1}(W_{t+1}^i) = a(W_{t+1}^i)^2 + bW_{t+1}^i + c, \quad (4.25)$$

$$= a(W_t^i)^2 (\tilde{x}_t^i)^\top \mathbb{E}[\tilde{R}_t \tilde{R}_t^\top | Z_t] \tilde{x}_t^i + (2aC_t + b)W_t^i (\tilde{x}_t^i)^\top \mathbb{E}[\tilde{R}_t | Z_t] + aC_t^2 + bC_t + c. \quad (4.26)$$

Here a, b, c are the parameters found during the estimation of f_{t+1} . If $aW_t^i \mathbb{E}[\tilde{R}_t \tilde{R}_t^\top | Z_t]$ is positive definite, we can make use of the convexity and fast numerical solvers to add constraints to the minimization. The solution is a new asset allocation $\hat{x}_t^i$. Using the new allocations $\{\hat{x}_t^i\}_{i=1}^{N_s}$ we can calculate new wealth values $\{\hat{W}_{t+1}^i\}_{i=1}^{N_s}$. Using portfolio weights from future times $\{t+1, \dots, T-1\}$ we can calculate the new wealth, replacement ratio and corresponding values of $(\hat{R}R^i - \gamma)^2$ for every path.

Using the set $\{(\hat{R}R^i - \gamma)^2\}_{i=1}^{N_s}$ we can calculate $\hat{V}_t = \mathbb{E}[(\hat{R}R - \gamma)^2 | Z_t]$. $\hat{V}_t$ is the expectation of the value function conditional on W_t and $\hat{x}_t$. We calculate this conditional expectation through cross-path regression. For this we regress $(\hat{R}R - \gamma)^2$ on the regression basis $B_t = \left[1, W_t, W_t^2, R_{i \in \{1, \dots, 19\}}^f(t)\right]$. We have to include the return on the riskless asset as the replacement ratio is dependent on this. Calculation of $\hat{V}_t$ concludes step 2.

3. Use the old values for the asset allocation to estimate the old values of $\tilde{V}_t = \mathbb{E}[(\tilde{R}R - \gamma)^2 | Z_t]$, where $\tilde{R}R$ is the replacement ratio calculated using the old asset allocation. Calculation of $\tilde{V}_t$ is done in the same way as in step 2, through cross-path regression on regression basis B_t . We then compare the estimated value functions for every path i and when $\tilde{V}_t^i > \hat{V}_t^i$, we update the asset allocation $\tilde{x}_t^i$ by $\hat{x}_t^i$.
4. Calculate the updated values of $V_t = \mathbb{E}[(RR - \gamma)^2 | Z_t]$, using the updated values of the asset allocation $\tilde{x}_t$. We update the wealth values and recalculate the replacement ratio. Through regression we can calculate the new values of V_t , using regression basis B_t . The values for V_t can now be used in step 2 for the next time step.

The backward algorithm can be repeated to update the asset allocation iteratively. As Cong and Oosterlee [12], we can apply bundling for the regressions in this algorithm. For the wealth case, Cong and Oosterlee [12] prove that this iterative process is convergent. In the next sections we will show that their proof also holds for the replacement ratio. First we analyse the forward part of the algorithm.

4.5.3. Analysis of the forward algorithm

The multi-stage strategy as proposed by Cong and Oosterlee in [12] is equivalent to solving the unconstrained pre-commitment problem. The only strong assumptions made besides the unconstrained allocation is that there is a risk-free interest rate and that asset returns are independent over periods. Does such a property hold for the replacement ratio? In some cases we can show it does indeed hold:

1. There is one constant, non-stochastic, risk-free interest rate R_f .
2. There is a non-stochastic term-structure of interest rates.

For both cases, we will show the assumption essentially reduces the problem to the one in [12], which we will present as well.

For one non-stochastic risk-free rate R_f we have that the expectations in equation (4.17), for the intermediate target, are no longer necessary. R_f is not a random variable and hence the intermediate target is automatically deterministic. The final wealth target Γ is also deterministic. Rewriting equation (4.16) yields:

$$\tilde{x}_t = \arg \min_{\tilde{x}_t} \mathbb{E} \left[\left(W_t(\tilde{x}_t \cdot R_t^e + R_f) \prod_{s=t+1}^{T-1} (R_f) - \Gamma \right)^2 \middle| Z_t \right], \quad (4.27)$$

where $R_t^e = \tilde{R}_t - R_f$ and we have set all intermediate contributions to zero. Equation (4.27) is equivalent to equation (8) from [12], hence, from this point on, we follow their proof. Note that the goal is to show that the strategy resulting from equation (4.27) is equivalent to the strategy resulting from the unconstrained version of the pre-commitment problem, Problem 4.3, given by:

$$J_t(W_t) = \min_{\{\tilde{x}_s\}_{s \in \{t, \dots, T-1\}}} \mathbb{E} \left[\left(W_t(\tilde{x}_t^\top R_t^e + R_f) \prod_{s=t+1}^{T-1} (\tilde{x}_s R_s^e + R_f) - \Gamma \right)^2 \middle| Z_t \right]. \quad (4.28)$$

To show the equivalence, we show that $J_t(W_t)$ can be reformulated in a useful way. The following lemma is Lemma 3.4 from [12], note that this lemma is set up for the one-dimensional case, but can be extended.

Lemma 4.5. *Suppose there is a constant, non-stochastic, risk-free rate and one risky asset with independent returns over time, then the value function $J_t(W_t)$, given by equation (4.28) can be formulated as:*

$$J_t(W_t) = L_t(W_t(R_f)^{T-t} - \Gamma)^2, \quad (4.29)$$

where $L_t = \prod_{s=t}^T l_s$, with l_t defined as follows:

$$l_t = 1 - \frac{\mathbb{E}[R_t^e]^2}{\mathbb{E}[(R_t^e)^2]}, \quad t \in \{0, 1, \dots, T-1\}$$

$$l_T = 1.$$

Proof. At time T we have:

$$J_T(W_T) = (W_T - \Gamma)^2,$$

which satisfies equation (4.29). At time $T-1$ we have:

$$J_{T-1}(W_{T-1}) = \min_{x_{T-1}} \mathbb{E} \left[\left(W_{T-1}(x_{T-1} R_{T-1}^e + R_f) - \Gamma \right)^2 \middle| Z_{T-1} \right]. \quad (4.30)$$

To obtain the analytic form of $J_{T-1}(W_{T-1})$, we solve the first-order conditions, which gives us that the optimal allocation $\tilde{x}_{T-1}^*$ is the solution to:

$$\mathbb{E} \left[\left(W_{T-1}(x_{T-1} R_{T-1}^e + R_f) - \Gamma \right) W_{T-1} R_{T-1}^e \middle| Z_{T-1} \right] = 0.$$

The optimal allocation is therefore given by:

$$x_{T-1}^* = \frac{(\Gamma - W_{T-1} R_f) \mathbb{E}[R_{T-1}^e]}{W_{T-1} \mathbb{E}[(R_{T-1}^e)^2]}. \quad (4.31)$$

Plugging equation (4.31) into equation (4.30) yields:

$$\begin{aligned} J_{T-1}(W_{T-1}) &= \mathbb{E} \left[\left(1 - \frac{\mathbb{E}[R_{T-1}^e] R_{T-1}^e}{\mathbb{E}[(R_{T-1}^e)^2]} \right)^2 \right] (W_{T-1} R_f - \Gamma)^2 \\ &= \left(1 - \frac{\mathbb{E}[R_{T-1}^e]^2}{\mathbb{E}[(R_{T-1}^e)^2]} \right) (W_{T-1} R_f - \Gamma)^2, \end{aligned}$$

which has the same form as equation (4.29). For all other time steps, we can use backward induction. Assume that at time $t + 1$ we have:

$$J_{t+1}(W_{t+1}) = L_{t+1}(W_{t+1}(R_f)^{T-(t+1)} - \Gamma)^2.$$

Then, at time t the value function $J_t(W_t)$ is given by:

$$\begin{aligned} J_t(W_t) &= \min_{x_t} \mathbb{E} \left[J_{t+1} \left(W_t(x_t R_t^e + R_f) \right) \middle| Z_t \right], \\ &= \min_{x_t} \mathbb{E} \left[L_{t+1}(W_t(x_t R_t^e + R_f)(R_f)^{T-t-1} - \Gamma)^2 \middle| Z_t \right] \\ &= \mathbb{E}[L_{t+1}] \min_{x_t} \mathbb{E} \left[(W_t(x_t R_t^e + R_f)(R_f)^{T-t-1} - \Gamma)^2 \middle| Z_t \right], \end{aligned}$$

where the first equality follows from the recursiveness of the value function, see equation (4.5), and the final equality follows from the independence of excess returns. Solving the first-order conditions yields:

$$x_t^* = \frac{(\Gamma - W_t(R_f)^{T-t})\mathbb{E}[R_t^e]}{W_t R_f^{T-t-1} \mathbb{E}[(R_t^e)^2]}. \quad (4.32)$$

After plugging the optimal allocation into the value function we have:

$$\begin{aligned} J_t(W_t) &= L_{t+1} \left(1 - \frac{\mathbb{E}[R_t^e]^2}{\mathbb{E}[(R_t^e)^2]} \right) (W_t(R_f)^{T-t} - \Gamma)^2 \\ &= L_t(W_t R_f^{T-t} - \Gamma)^2. \end{aligned}$$

This finalizes the proof. $\square$

With Lemma 4.5 we can prove the following theorem, which is Theorem 3.6 from [12].

Theorem 4.1. *Suppose there exists a constant, non-stochastic risk-free rate, excess asset returns are independent over time and asset allocations are unconstrained. Then the optimal strategy for the pre-commitment problem, Problem 4.3, is equivalent to the optimal strategy given by equation (4.27).*

Proof. For the pre-commitment problem, Problem 4.3, the optimal strategy is given by:

$$x_t^{pc} = \arg \min_{x_t} \mathbb{E} \left[J_{t+1} \left(W_t(x_t R_t^e + R_f) \right) \middle| Z_t \right],$$

where $J_t(W_t)$ is given by equation (4.28). By using the form of $J_{t+1}(\cdot)$ given by Lemma 4.5, we have:

$$x_t^{pc} = \arg \min_{x_t} \mathbb{E} \left[L_{t+1} \left(W_t(x_t R_t^e + R_f)(R_f)^{T-t-1} - \Gamma \right)^2 \middle| Z_t \right].$$

By the independence of excess asset returns we can treat L_{t+1} as a constant and therefore the optimal strategy is given by:

$$x_t^{pc} = \arg \min_{x_t} \mathbb{E} \left[\left(W_t(x_t R_t^e + R_f)(R_f)^{T-t-1} - \Gamma \right)^2 \middle| Z_t \right],$$

which is equivalent to equation (4.27). $\square$

Theorem 4.1 shows that our forward algorithm for optimization of the replacement ratio is equivalent to the unconstrained pre-commitment problem, under the assumption that there is one constant risk-free rate. This finishes the first case.

For the second case, there is a non-stochastic term-structure: for every tenor i we have a riskless rate R_i^f . We again have deterministic intermediate and final targets. For the wealth target, the situation changes

slightly from the previous case, as we now discount by different rates for every tenor. For the proof of equivalence this does not matter, we still obtain a fixed wealth target, which is what we need. The intermediate targets however, are deterministic but not constant, as we will show below. We aim to solve:

$$\tilde{x}_t = \arg \min_{\tilde{x}_t} \mathbb{E} \left[\left(W_t(\tilde{x}_t \cdot \tilde{R}_t) (R_{T-t-1}^f)^{T-t-1} - \Gamma \right)^2 \middle| Z_t \right], \quad (4.33)$$

which is different from equation (4.27). Nevertheless, as all rates are deterministic, we can divide by $\left(\frac{(R_{T-t-1}^f)^{T-t-1}}{(R_1^f)^{T-t-1}} \right)^2$.

As R_1^f is the one-period rate and is constant, this rate is essentially the same as R_f in (4.27). So, now we solve:

$$\tilde{x}_t = \arg \min_{\tilde{x}_t} \mathbb{E} \left[\left(W_t(\tilde{x}_t \cdot R_t^e + R_1^f) \prod_{s=t+1}^{T-1} (R_1^f) - \frac{\Gamma}{(R_{T-t-1}^f)^{T-t-1}} \right)^2 \middle| Z_t \right], \quad (4.34)$$

again with $R_t^e = \tilde{R}_t - R_f$ and all intermediate contributions set to zero. We have a problem with one risk-free rate (R_1^f) and a deterministic target. Unfortunately, the target is not constant, it changes for every time t . This makes it impossible to use the same proof as for the first case. We therefore chose to only use the term-structure in discounting the cash flows that represent the pension payments. This only influences the wealth target Γ , but insures it is constant over time. For all other instances where the interest rate is used, we use the 1-year rate. In this case we simplify to equation (4.27), where δ_t has changed to:

$$\delta_t = \frac{\Gamma - \sum_{s=t+1}^{T-1} C_s (R_1^f)^{T-s}}{(R_1^f)^{T-t}}. \quad (4.35)$$

This concludes the examples. We have tried to introduce stochastic interest rates, but even under the most simple assumptions, proving equivalence has not been successful. The main problem is that the target cannot be discounted by fixed interest rates, and we can therefore not obtain a wealth optimization problem.

Although we cannot prove equivalence between the pre-commitment and multi-stage solution (forward algorithm) in a stochastic rates framework, we can assume that a solution by the forward algorithm is a good initial guess to start the backward algorithm. A converging backward algorithm would then yield the optimal solution.

4.5.4. Convergence of the backward algorithm

The convergence of the backward algorithm has been proven by Cong and Oosterlee. This proof is based on the monotonicity property of the Bellman operator, the same operator is used in the Algorithm 3; we rewrite equation (4.5):

$$V_t(W_t) = \min_{\tilde{x}_t} \mathbb{E} \left[V_{t+1} \left(W_t(\tilde{x}_t \cdot \tilde{R}_t) + C_{t+1} \right) \middle| Z_t \right] \quad (4.36)$$

$$V_t = \Psi_t V_{t+1}. \quad (4.37)$$

Here Ψ_t is the Bellman operator:

$$(\Psi_t h)(W_t) = \min_{\tilde{x}_t} \mathbb{E} \left[h \left(W_t(\tilde{x}_t \cdot \tilde{R}_t) + C_{t+1} \right) \middle| Z_t \right] \quad (4.38)$$

From this point we can use Lemma 4.2 and Proposition 4.3 from [12] to show convergence of Algorithm 3.

Convergence is only guaranteed under proper estimation of the value function V_t . As regression will introduce errors into the estimations of these functions, convergence is not assured.

4.6. Conclusion

In this chapter we have introduced dynamic mean-variance portfolio optimization. Dynamic mean-variance optimization is used for multi-period portfolio optimization. At every portfolio rebalancing time, the asset

allocation is optimized conditional on the past. It has been shown that dynamic asset allocations perform better than deterministic asset allocations, currently used in life-cycle investing.

We defined the DMV problem for the replacement ratio. Instead of using the original definition of dynamic mean-variance, we used a second moment optimization, a target-based approach, as this allows for a recursive representation. This formulation of the optimization problem is called the pre-commitment problem.

Building on the work of Cong and Oosterlee, we developed an algorithm to solve the pre-commitment problem for optimization of the replacement ratio. This algorithm has a two-stage approach. First we generate an initial, sub-optimal, solution using a forward algorithm. The forward algorithm works by assuming that capital is invested in the risk-free asset for every period after the period for which the asset allocation is optimized. We have shown that the forward algorithm is equivalent to solving the pre-commitment problem when the asset allocation is unconstrained and some assumptions are made on the asset returns.

After the forward part we update the (sub-optimal) solution using a backward algorithm. This algorithm iteratively improves the initial solution. The algorithm is shown to be converging, under the assumption that the conditional expectation of the value function is correctly estimated.

The estimation of conditional expectations is done through regression. Standard OLS regression may introduce errors into the estimation. Due to the recursive nature of the backward algorithm, these errors can accumulate, yielding sub-optimal asset allocations as result.

We need to estimate conditional expectations often in our algorithm, not only for the value function, but also for the returns. The forward and backward parts of the algorithm both assume knowledge of the conditional expectations of returns. When using basic models, this knowledge is available. We, however, wish to make our algorithm applicable to a wide range of models, even models where the exact drivers of returns are unclear. We need to estimate these conditional expectation, e.g. through regression. Precise estimation of conditional expectation is crucial for convergence of our algorithm. Therefore, we will focus on the estimation of conditional expectation in the next chapter.

Regression and conditional expectation

In the previous chapter we saw that the algorithm for mean-variance optimization of the replacement ratio, Algorithm 3, contains multiple conditional expectation calculations. These expectations have to be estimated numerically using the available sample paths.

In this chapter we discuss methods to estimate conditional expectations from samples. We want to avoid using simulation in simulation, i.e., for every sample we start a new simulation conditional on that sample. Using simulation in simulation is computationally intensive as the required computing power grows exponentially. We therefore focus on two other methods for estimating conditional expectation.

The first method is regression and the second is stratified state aggregation (SSA). The first is well known, has many variants and is by definition linked to conditional expectation. SSA is less known but in a simulation context it has an intuitive foundation. We discuss both categories, but within regression we only discuss methods which we have found to be relevant to this thesis. Those include the distinction between regress-now and regress-later and three implementations of regression: ordinary least squares (OLS), the lasso and support vector regression (SVR).

We first introduce the mathematics of estimating conditional expectation, and then present tests of the methods described.

5.1. Definition of conditional expectation through regression

In Chapter 4 we have focused on simulation-based optimization, we are therefore interested in conditional expectations with respect to an event, i.e., the realization of a random variable. For such conditional expectations, we have the following lemma:

Lemma 5.1 (Decomposition of random variable into conditional expectation). *Let Y_i be a realization of the random variable Y and X_i of the random variable X . Then*

$$Y_i = \mathbb{E}[Y|X_i] + \epsilon_i, \quad (5.1)$$

where $\mathbb{E}[\epsilon_i|X_i] = 0$ and $\mathbb{E}[h(X_i)\epsilon_i] = 0$ for any function $h(\cdot)$.

Proof. When $Y_i = \mathbb{E}[Y|X_i] + \epsilon_i$, we have $\mathbb{E}[\epsilon_i|X_i] = \mathbb{E}[Y_i - \mathbb{E}[Y|X_i]|X_i] = 0$ and $\mathbb{E}[h(X_i)\epsilon_i] = \mathbb{E}[h(X_i)\mathbb{E}[\epsilon_i|X_i]] = 0$, which proves the lemma. $\square$

In case of regression, we estimate a function $g(\cdot)$ such that $Y = g(X) + \epsilon$ where ϵ is the residual term. The residual is a random variable that should be independent of X and should have mean zero. For realizations Y_i and X_i we have that $Y_i = g(X_i) + \epsilon_i$ and by Lemma 5.1: $\mathbb{E}[Y|X_i] = g(X_i)$.

So, we have shown the relation between regression and conditional expectation, our main purpose now is to find the function that represents the conditional expectation, using various regression techniques.

5.2. Regression techniques

In this section we describe how we can use regression to estimate conditional expectation in a simulation setting. To do so, we introduce two methods: regress-now and regress-later, both capable of estimating

conditional expectation in a simulation setting. First, we describe three types of regression: ordinary least squares, the lasso and support vector regression.

Ordinary least squares (OLS) regression is the widely used standard form of regression. OLS has the tendency to *overfit* in high-dimensional problems, i.e., the regression result partly describes the noise. Therefore, we use the lasso. The lasso is designed to automatically select features from a high-dimensional dataset. Both OLS and the lasso rely on the user to specify the regression basis and therefore the functional form that is fitted. SVR can be used without specifying a functional form and is therefore used as regression technique. A comprehensive description of all three methods is given in Appendix 5.A.

5.2.1. Regression applied to scenarios

Having established which regression techniques we will use, we can discuss how these can be applied in a simulation setting, specifically a scenario setting. In a scenario setting, we have multiple realizations of a set of time-series. We are interested in the expectation of some of these time-series, conditional on the past. The literature regarding conditional expectation in a scenario setting is mainly focused on valuing American options and on dynamic portfolio optimization. We present two methods for conditional expectation estimation: regress-now and regress-later. Both methods have the same underlying idea: regress the variable of which conditional expectation is calculated on the conditioning variables.

Longstaff and Schwartz [30] use least squares regression to value American options, they use the *regress-now* method. We define regress-now as a general method for approximating conditional expectation:

Definition 5.1 (Regress-now). By regress-now, we estimate $\mathbb{E}[Y_{t+1}|X_t]$, by estimating

$$Y_{t+1} \approx \sum_{i=0}^N \alpha_i \psi_i(X_t), \quad (5.2)$$

where we regress the value of random variable Y_{t+1} on a basis formed by the previous value of state variables X_t . $\psi_k(\cdot)$ indicates a basis function of the argument. Now, using the properties of conditional expectation:

$$\mathbb{E}[Y_{t+1}|X_t] \approx \mathbb{E}\left[\sum_{i=0}^N \alpha_i \psi_i(X_t)|X_t\right] = \sum_{i=0}^N \alpha_i \psi_i(X_t). \quad (5.3)$$

Equation (5.3) shows how we approximate conditional expectation using regress-now.

As described in Section 4.4.2, Cong and Oosterlee use *regress-later* for the estimation of conditional expectation.

Definition 5.2 (Regress-later). By regress-later, we estimate $\mathbb{E}[Y_{t+1}|X_t]$, by estimating:

$$Y_{t+1} \approx \sum_{i=0}^N \alpha_i \psi_i(X_{t+1}), \quad (5.4)$$

through regression. Then using the properties of conditional expectation:

$$\mathbb{E}[Y_{t+1}|X_t] \approx \mathbb{E}\left[\sum_{i=0}^N \alpha_i \psi_i(X_{t+1})|X_t\right] = \sum_{i=0}^N \alpha_i \mathbb{E}[\psi_i(X_{t+1})|X_t], \quad (5.5)$$

where knowledge of the conditional expectations $\mathbb{E}[\psi_k(X_{t+1})|X_t]$ should be available.

Regress-later is more stable than regress-now, according to [23]. Using regress-later, however, is not always possible. The reason for this is that explicit formulas for conditional expectations of the basis function are not always available. For example, in our forward algorithm, we do need to calculate the conditional first and second moment of the asset returns, see equation (4.24). To calculate these, we could use regress-now. This can however be unstable, so we will investigate with numerical examples.

5.3. Stratified state aggregation

Another method for calculation of conditional expectation is *stratified state aggregation*. Introduced by Baraquand and Martineau [3] for valuation of American options, stratified state aggregation (SSA) is a technique

that uses partitioning. By dividing the state space into non-overlapping partitions, the future expectations of paths in the partition can be calculated. These future expectations are then used as approximation of the conditional expectation for every path in the partition. So although paths may have different values at time t , if they are in the same partition, the expectation of the path at time $t + 1$ conditional on t will be the same for both paths.

To show the idea mathematically we follow the description of SSA by Coyle and Yang [15].

We partition the state-space, for example the asset returns and interest rates, at every time $t \in \{0, \dots, T-1\}$ into $N(t)$ cells $Q_k(t)$, $k \in \{1, \dots, N(t)\}$. State aggregation means that we have a constant conditional expectation on every cell:

$$\mathbb{E}[Y_{t+1}|X_t](k) = \mathbb{E}[Y_{t+1}|X_t \in Q_k(t)]. \quad (5.6)$$

The stratification part is concerned with mapping the n -dimensional state space to an l -dimensional space, where $l < n$. In our case, we will map to $\mathbb{R}_+$. We pick a sorting value, for example returns in the previous period R_{t-1} , to create a partition. This means a cell could be:

$$Q_k(t) = \{X_t : R_{t-1,(kN(t)+1)} \leq R_{t-1}(X_t) < R_{t-1,((k+1)N(t)+1)}\}, \quad (5.7)$$

where $R_{t-1,(i)}$ is the i -th order statistic of the R_{t-1} s of all paths. This stratification choice guarantees that $\mathbb{P}(X_t \in Q_k(t)) = \frac{1}{N(t)}$.

The idea is that when there is a large number of sample paths and sufficiently small cells, the conditional expectation which is constant over the cell, converges to the actual conditional expectation of all paths in the cell. Although the idea behind SSA is intuitively appealing, Coyle and Yang [15] have proven that no matter how small the cells are, the different points in the cell will have different expected future values. For American option pricing this implies that incorrect exercise decisions will be made. This could mean for our problem that wrong portfolio decision will be made. Given the necessity to estimate the conditional expectations of asset returns, however, we will use SSA, comparing the path estimator SSA to the direct regression-based methods.

5.4. Numerical tests

In this section we compare the performance of regression techniques and SSA for the estimation of conditional expectation. We focus on estimation of conditional expectation of asset returns as well as of the value function, as defined in equation (4.4).

We use three different models for asset returns in these numerical tests: the geometric Brownian motion, an autoregressive (AR) model with one lagging component and a GARCH-in-mean (MGARCH) model with single lags.

In general, we find that regression techniques outperform SSA. Furthermore, we show that SVR is capable of approximating correct functional forms of the conditional expectations, but performs slightly worse than OLS and the lasso. Finally, we show that OLS and the lasso are approximately equal in performance, but as OLS has the tendency to overfit, the lasso is in general preferred.

5.4.1. Test procedure

For every model of asset returns, we test the performance of OLS, the lasso, SVR, and SSA on the estimation of the conditional expectation of the asset returns.

We assess the different methods by calculating the mean squared error (MSE) given by:

$$MSE := \frac{1}{N} \sum_{i=1}^N (\hat{R}_{t,cond}^i - \mathbb{E}[R_t|R_{t-1}^i])^2, \quad (5.8)$$

where $\hat{R}_{t,cond}^i$ is the estimated value of the conditional expectation for path i and $\mathbb{E}[R_t|R_{t-1}^i]$ is the analytic conditional expectation given the previous return of path i . That is, R_{t-1}^i is a sample drawn from the asset return distribution we use and using R_{t-1}^i , we can calculate $\mathbb{E}[R_t|R_{t-1}^i]$ analytically.

For the regression methods we need to choose a regression basis. For OLS and the lasso we use a quadratic regression basis. We can use a regression basis consisting of only the previous time values: R_{t-1} , when estimating $\mathbb{E}[R_t|Z_{t-1}]$, but we can also use $\{R_s\}_{s \in \{t-1, \dots, t-n\}}$. We will evaluate the MSE for increasing n .

Finally, for the stratified state aggregation, we use the previous value R_{t-1} for sorting the values into cells.

For the estimation of the value function given by equation (4.4), we assume a constant risk-free rate of 4.3%. We have $N = 5$ time steps and we estimate the value function conditional on the state space after time step 3. Additionally, at time 0 not all paths start with the same wealth, but we draw from a uniform distribution. The value function is obtained through calculating the wealth paths with asset allocations obtained from the forward algorithm, given by equation (4.16). In the forward algorithm we use the analytically known conditional expectations of the asset returns. We apply the no-shorting and no-leverage constraints as in Problem 4.1.

The value function is not analytically tractable when the asset allocation is constrained, calculation of the MSE is therefore a problem. We can solve this problem by simulating conditionally:

Definition 5.3 (Conditional simulation). By conditional simulation, we estimate

$$\mathbb{E}[Y|X = c],$$

by sampling from the conditional probability distribution $f(x = c, y)$ and taking the average over the sample. Here, $f(x, y)$ is the joint probability density functions of random variables X and Y .

We can estimate $V_t = \mathbb{E}[(RR - \gamma)^2 | Z_t = W_t]$ through conditional simulation by setting wealth W_t to a specific starting value and simulate paths using that starting value. In this way we can obtain accurate estimates of V_t for different values of W_t so that we can assess the quality of our methods.

We present our results per asset return model, starting with the basic model: the geometric Brownian motion.

5.4.2. Geometric Brownian Motion

Analytically the geometric Brownian motion (GBM) is given by:

$$dS_t = \mu S_t dt + \sigma S_t dW_t,$$

with exact solution:

$$S_t = S_0 \exp\left(\left(\mu - \frac{1}{2}\sigma^2\right)t + \sigma\sqrt{t}Z_t\right),$$

where $Z_t \sim \mathcal{N}(0, 1)$. For simulation, we use the discrete version of the stochastic differential equation:

$$\Delta S = \mu S_t \Delta t + \sigma S_t \sqrt{\Delta t} Z_t,$$

where $Z_t \sim \mathcal{N}(0, 1)$. The discrete return is given by:

$$\frac{\Delta S}{S_t} \sim \mathcal{N}(\mu \Delta t, \sigma^2 \Delta t) \quad (5.9)$$

We can therefore simulate returns by drawing from a normal distribution with mean μ and variance σ^2 (we set our time step to 1 year). We set μ to 0.08 and σ to 0.2. The sample size for this test is 10^4 .

Conditional expectation of returns

In this subsection we evaluate our methods for estimation of conditional expectation of GBM returns. We show that both OLS and SVR overfit and therefore that the lasso is the preferred estimations method for the conditional expectation of GBM returns.

The conditional expectation of the GBM returns is always μ , see equation (5.9), we therefore expect all the methods to have a constant term which equals the mean of the sample and zero coefficients otherwise. For the lasso, we indeed find a constant and zero coefficients. For OLS, however, we find non-zero coefficients, so OLS is effectively overfitting the model. Figure 5.1 shows the effect of overfitting on the MSE, as given by equation (5.8). In Figure 5.1, the MSE is plotted for number of previous time steps used in the regression, the number of features. As the number of features used increases, the MSE for OLS increases, for the lasso it is constant. Therefore, the lasso is preferred over OLS in the estimation of the conditional expectation of GBM returns.




Figure 5.1: Means squared error as given by equation (5.8). For OLS, the lasso and SVR, obtained with estimation of conditional expectation of GBM returns, distribution given by equation (5.9), for varying regression bases.

Figure 5.1 shows that the MSE of SVR grows faster and is larger than the errors for OLS and the lasso. Part of the reason for this is the setting of parameter ϵ , as used in equation (5.21). The parameter ϵ determines how much noise around the fitted function is allowed. When ϵ is small, part of the random noise will influence the fit, and hence the SVR overfits the model. Increasing ϵ does solve the overfitting problem, however, this makes estimation of the constant problematic. As proper estimation of the constant is required for GBM returns, SVR is not a suitable method.

For SSA, we need to choose a number of cells for which the conditional expectation is calculated. In this case the optimal number of cells is 1, this will yield the sample average which is the best approximation of the conditional expectation. The approximation will become worse when the number of cells is increased. Indeed, for one cell we find $MSE = 6.4 \times 10^{-6}$ which is in the same order as the MSE for the lasso. For 10 cells we find $MSE = 5.9 \times 10^{-5}$, which is worse than OLS and the lasso but better than SVR. Finally, for 100 cells we find $MSE = 4.5 \times 10^{-4}$, which is worse than SVR. We conclude, for the GBM, SSA does not add anything as the sample average is the best estimator.

Value function

In this subsection we evaluate our methods in estimating the value function, as given by equation (4.4), when the asset returns are given by the GBM. We find that OLS and the lasso both perform very well in estimation of the value function, though both are prone to overfitting. Furthermore, we find that the SVR is able to find the right functional form for the value function, however, it performs slightly worse than OLS and the lasso.

We estimate the value function by regressing on the wealth of the previous time step only versus on the wealth over the entire path. We compare the results for these two regression bases. Note that for the GBM, the conditional expectation at any time should have no dependence on any of the previous times. Although independence implies zero weights for all but the most recent wealth value, the wealth values are highly correlated and it is therefore expected that our numerical methods will give weight to all values in the state space.

Figure 5.2 confirms our expectations. On the top row of Figure 5.2 we show the OLS estimate of the value function, the green line. The blue line represents the sample averages from conditional sampling and the blue dots represent the samples obtained from conditional sampling. The bottom row shows the relative absolute difference between the estimate using OLS and the conditional sample average. On the left side we have the results for regression on previous wealth only, on the right side for regression on the entire state space. Regression on the entire state space improves the solution, though not by much. This improvement




Figure 5.2: Top panels show the value function, given by equation (4.4), estimated through OLS regression in green, the value function found through conditional sampling in blue and the samples from conditional sampling as blue dots. GBM asset returns are used for the wealth process, the optimal asset allocation is found using the forward algorithm, as described in Section 4.5. The bottom panels show the relative absolute difference between estimate of the value function through OLS and conditional sample average of the value function.

is due to overfitting, only the previous wealth values should influence the conditional expectation, but as wealth values are highly correlated over time (in the order of 0.95!), older wealth values also have non-zero coefficients.

The peak in relative error for a wealth value of 13000 in Figure 5.2 is caused by the small variance in final replacement ratio values. When the size of wealth is 13000 at the conditioning time, we are on 'the perfect path', we can reach the target wealth by only investing in the risk-free asset. This causes a low variance in the squared difference between replacement ratio and target, the dependent variable in the regression. Low variance of the dependent variable causes high variance in estimated parameters, see Example 5.1, and therefore a worse estimator.

The lasso is also not able to shrink the coefficients of irrelevant wealth values. So, we do not obtain a solution that only depends on the last wealth value. The correlation between wealth values over time is so high, that the lasso result barely differs from the OLS result. This is unfortunate, as the lasso was meant to prevent overfitting. Still, lasso does not perform worse than OLS and therefore both OLS and the lasso can be used in estimation of the value function for independent returns.

Finally, we test whether SVR can fit the (approximately) quadratic function. The parameters of the SVR found through 10-fold cross-validation are: $\epsilon = 0.1$, $C = 10^3$ and $\gamma = 0.01$. The result of the estimations can be found in Figure 5.3, which is build-up in the same way as Figure 5.2. The top panels of Figure 5.3 show that SVR is able to approximate the functional form of the value function. The relative error, shown in the bottom panels of Figure 5.3, is larger than for OLS and the lasso. This is not surprising, for OLS and the lasso we have specified the approximately correct functional form. For SVR however, we have specified no form at all, and




Figure 5.3: Top panels show the value function, given by equation (4.4), estimated through SVR in green, the value function found through conditional sampling in blue and the samples from conditional sampling as blue dots. GBM asset returns are used for the wealth process, the optimal asset allocation is found using the forward algorithm, as described in Section 4.5. The bottom panels show the relative absolute difference between estimate of the value function through SVR and conditional sample average of the value function.

hence for the same number of samples it logically has worse performance.

When we compare the left and right side of Figure 5.3, we notice that regression on the entire state space has larger relative errors than regression on current wealth only. This result is in contrast with the result for OLS and the lasso. The worse performance can be explained by the number of features: when we regress on the entire state space, we have many more features to take into account, while the number of samples stays the same. Hence, per feature, we have less data, which accounts for the worse performance when regressing on the entire state space. Despite being able to approximate the right functional form, we conclude that SVR is not preferred over OLS or the lasso for estimation of the value function.

5.4.3. Autoregressive returns

To test how the introduction of path-dependent returns influences the estimation of conditional expectation, we introduce the following autoregressive model with one lag for the excess returns:

$$r_t^e = \alpha r_{t-1}^e + \beta + \epsilon_t, \quad (5.10)$$

where $\alpha = 0.8$, $\beta = 0.0074$ (such that the expected stationary excess return is 3.7%) and $\epsilon_t \sim \mathcal{N}(0, 0.2^2)$. Any autoregressive model has a representation based on all previous time steps, given by:

$$r_t^e = \sum_{i=1}^t \alpha^{t-i} (\beta + \epsilon_i) + \alpha^t r_0^e. \quad (5.11)$$

Therefore, returns are correlated with all previous returns, although the correlation will be exponentially decreasing.

We present the results using the autoregressive return structure and 10^4 sample paths.

Conditional expectation of returns

In this subsection we evaluate our methods for estimation of the conditional expectation of autoregressive asset returns. We show that OLS is the preferred method for estimation when the regression basis is small, and the lasso is preferred when the regression basis is large. Furthermore, we show that both SVR and SSA are not able to properly estimate the conditional expectation of autoregressive return.

The conditional expectation of returns is given by $E[r_t^e | r_{t-1}^e] = \alpha r_{t-1}^e + \beta$. A linear function in one parameter, the previous return value. Despite dependence on the previous return only, we make the regression basis bigger, as for the GBM, adding more previous time steps, to see how the various regression methods cope with this.

In Figure 5.4, we show the out-of-sample MSE, given by equation (5.8), for OLS, the lasso and SVR. Figure 5.4 shows that the performance difference between OLS and the lasso is subtle. A look at the estimated parameters reveals that OLS overfits, it has non-zero coefficients for every parameter. Still, OLS performs better than the lasso, this is caused by shrinking of estimated coefficients done in the lasso. The parameter α , as in equation (5.11), is more correctly specified by OLS as the lasso tends to shrink this parameter below its actual value. Nevertheless, as we increase the size of the regression basis, the OLS results becomes worse, where the lasso result stays constant. Preference between OLS and the lasso therefore depends on the size of the regression basis: for a small basis we prefer OLS, for a large basis we prefer the lasso.

Figure 5.4 furthermore reveals that SVR performs the least and clearly overfits as we increase the regression basis. As for the GBM, SVR is not a suitable method for the estimation of conditional expectation of asset returns.

For SSA, it is not possible to predict the optimal number of cells, as the sample average is not the best estimator of conditional expectation. Still, for SSA, it turns out that the sample average, one cell, works better than any larger number of cells. The MSE increases as we increase the number of cells. For one cell, the MSE is 0.06, many times larger than for any of the regression methods. This result makes it doubtful whether SSA is a useful method for estimating conditional expectations.

The underperformance of SSA might be due to the metric we use to evaluate performance: mean squared error. The MSE is based on the squared errors, exactly the measure the regression techniques try to minimize. When we use the *mean absolute error* (MAE), the absolute error instead of the squared error, SSA performs more as expected. We have found the MAE to be non-monotonic in the number of cells, in our sample we found a minimum at 50 cells. So, using MAE as metric, the sample average is not the best approximation for the conditional expectation. Still, performance is well below that of OLS: SSA has an MAE of 0.18 versus 0.003 for OLS.

Value function

In this subsection we evaluate our methods in estimating the value function, as given by equation (4.4), when the asset returns are given by the autoregressive process given by equation (5.11). As returns are no longer independent over time, we know that increasing the size of the regression basis can benefit the estimation of the value function. We saw for the GBM, with independent returns, however, that the correlation between wealth values of different time steps is high. High correlation between time steps makes it is doubtful that we can distinguish between whether a feature is relevant for the estimation or whether it is taken into account because it is highly correlated with a relevant feature. We show that both OLS and the lasso are able to approximate the value function well, which shows Algorithm 3 is suitable for asset returns with dependence over time. On top of that, we show that SVR is able to outperform OLS and the lasso in some cases, showing that removing the assumption that the value function is quadratic in wealth, can be beneficial.

Figure 5.5 shows the results for lasso regression. Figure 5.5 is build up in the same ways as Figure 5.2. We present the results for the lasso, as the difference between OLS and the lasso in performance is small. The lasso performs slightly better and has zero coefficients, where OLS only has non-zero coefficients. To obtain these results, the regularisation parameter of the lasso λ , as in equation (5.19), took the high value of 10^8 . Such a value for λ indicates a large penalty on regression coefficient's size. Comparing Figure 5.5 to Figure 5.2, we see that the performance for autoregressive returns is worse than for GBM returns. This result follows from the more complicated nature of the model for the returns. When comparing the left and right side of Figure 5.5, we see a considerable improvement when we regress on the entire state space. This is as expected, the correlation between the returns over time implicates there is information in older values of wealth and therefore more precise estimation is possible. Furthermore, with the relative error at less than 1 %, this example shows that our algorithm can indeed be used for a model with non-independent returns.




Figure 5.4: Means squared error as given by equation (5.8). For OLS, the lasso and SVR, obtained with estimation of conditional expectation of autoregressive returns, given by equation (5.11), for varying regression bases.




Figure 5.5: Top panels show the value function, given by equation (4.4), estimated through the lasso in green, the value function found through conditional sampling in blue and the samples from conditional sampling as blue dots. Autoregressive asset returns are used for the wealth process, the optimal asset allocation is found using the forward algorithm, as described in Section 4.5. The bottom panels show the relative absolute difference between estimate of the value function through SVR and conditional sample average of the value function.




Figure 5.6: Top panels show the value function, given by equation (4.4), estimated through SVR in green, the value function found through conditional sampling in blue and the samples from conditional sampling as blue dots. Autoregressive asset returns are used for the wealth process, the optimal asset allocation is found using the forward algorithm, as described in Section 4.5. The bottom panels show the relative absolute difference between estimate of the value function through SVR and conditional sample average of the value function.

For the SVR ($\epsilon = 0.1$, $C = 10^4$, $\gamma = 0.1$), the results are surprising. As for the GBM returns, the regression on the entire state space performs poorly, which can be seen on the right side of Figure 5.6. With regression on the previous wealth only, however, the SVR performs better than OLS and the lasso. The reason for this outperformance is that the value function with autoregressive returns may not be quadratic. Assuming a quadratic function in wealth for the value function works well for OLS and lasso estimates, it is still an imposed assumption. As this assumption is not present in the SVR estimation, this can explain the out performance of SVR for regression on the current wealth only. Nevertheless, OLS and lasso performance is better than SVR when using the entire state space and are therefore still preferred methods.

5.4.4. MGARCH returns

To evaluate the performance of our methods on more complicated returns structures, we introduce a GARCH-in-mean (MGARCH) model for the asset returns. The model we use was proposed by Bollerslev et al. [8], we use the parameters estimated by de Goeij and Marquering [16]. The model is multivariate and models long term zero-coupon bond ($i = 1$), short term zero-coupon bond ($i = 2$) and equity ($i = 3$) returns. The model is given by:

$$r_{i,t+1}^e = \mu_i + \lambda \sum_{j=1}^3 w_{j,t+1} \sigma_{ij,t+1} + \epsilon_{i,t+1}, \quad (5.12)$$

$$\sigma_{ij,t+1} = \gamma_{i,j} + \alpha_{ij} \epsilon_{i,t} \epsilon_{j,t} + \beta_{ij} \sigma_{ij,t}, \quad (5.13)$$

where $\epsilon_{i,t+1}$ is multivariate normally distributed with mean vector zero and covariance matrix with elements $\sigma_{ij,t}$. The coefficients can be found in Appendix 5.B, Table 5.1.

The model specified by equations (5.12) and (5.13) is a model where the volatility is dependent on previous volatility values and the realization of the random component. In such a model, volatility is persistent,

high volatility causes high volatility in the next period etc. The volatility component is also present in the equation for the return (5.12), which means high volatility also mean higher expected returns. Because the volatility influences the conditional expectation, we expect all methods to have difficulty with estimating the conditional expectation of this model. The only dependence the volatility has on the observed values r_t^e is via the square of random part ϵ_t and it will therefore be hard to model the conditional expectation using a quadratic regression basis of previous returns.

We simulate 10^4 paths, with 150 time steps, of which we only use the last 50. We only use the last 50 to ensure we are in the stationary state, the start values should have no influence on the return dynamics.

Conditional expectation of returns

In this subsection we show how our methods perform in estimation of the conditional expectation of MGARCH asset returns, given by equation (5.12). The results are disappointing for all methods. Both OLS and SVR have relatively good in-sample performance, but both clearly overfit: their out-of-sample performance is poor. The lasso has only zero coefficients, meaning the lasso estimate is given by the sample average. SSA yields the same result: for one cell, the sample average, the performance is best. When the MAE is applied, 2 cells is optimal for SSA, and performance matches the performance of OLS. All methods for estimating conditional expectation fail for the MGARCH returns, they do not perform better than using the sample average. As the lasso estimator is the sample average, the lasso is the preferred method when asset returns are given by an involved model.

Even when we append the regression basis with the previous values of the volatility ($\sigma_{ij,t}$), the regression methods are not able to capture the model structure. Note however that the parameters to be estimated are small, order 10^{-3} , which could explain the underwhelming results.

5.5. Conclusion

In this chapter we presented and tested multiple methods to estimate conditional expectation in a simulation setting. We tested all methods on estimation of the conditional expectation of asset returns, generated through simulation. For OLS and the lasso the results look promising, in the case of more basic return dynamics, OLS and the lasso are able to capture the model structure. OLS tends to overfit when the regression basis consists of many features and therefore the lasso is generally preferred.

SVR performance was worse than that of OLS and the lasso, SVR tends to overfit and the sample sizes used in the test were too small to sufficiently counter this. As we will not use larger samples in the simulation-based algorithm presented in Chapter 4, SVR is not suitable for the estimation of conditional expectations of returns.

SSA had a poor performance in all of the returns models we used, under the MSE metric. When the MAE metric was applied, SSA still underperformed in the GBM and AR case, but performed on par with OLS in the MGARCH case.

We tested all regression-based methods for estimation of the value function, the main function used in algorithm presented in Chapter 4. For all regression techniques we tested whether only the last time step or all previous time steps (entire state space) should be incorporated in the regression basis.

For independent returns, OLS performed slightly better when regressed on the entire state space than when regressed only on the last time step. This is entirely due to overfitting, as we know the value function to be dependent on the last time step only (for independent returns). The results for OLS and the lasso were almost identical. SVR was able to fit the quadratic functional form, but performed slightly worse than OLS and the lasso.

For autoregressive returns, regressing on the entire state space performed better than on the last time step only, when we used OLS and the lasso. As for independent returns, OLS and the lasso has similar results. Nevertheless, as the lasso set some coefficients to zero, this method is preferred. The performance of the lasso when regressing on the entire state space has us convinced that this method is suitable for our optimization algorithm when returns are not independent.

A final, surprising result, is the performance of SVR for autoregressive returns. SVR was able to outperform the lasso when regressed on the last time step only. We expect this is due to the functional form of the value function when returns are not independent. For the lasso we used a quadratic regression basis, which may not be the correct functional form for the value function. Therefore SVR, in which we do not impose any

functional, could have outperformed the lasso. Nevertheless, when regressed on the entire state space, the lasso performs better than SVR and is therefore the preferred method.

In the next chapter we investigate how our optimization algorithm performs with some of the methods tested in this chapter incorporated.

5.A. Regression techniques

In this appendix we elaborate on the regression techniques used in Chapter 5: ordinary least squares (OLS) regression, the lasso and support vector regression (SVR).

Throughout this section, y is the dependent variable, the variable of which we wish to estimate the conditional expectation, and x is the independent variable, the variable on which we condition. We can have multiple independent variables, in which case we use the vector $\vec{x}$. A subscript i is used to indicate an observation of y or x .

5.A.1. Ordinary Least Squares regression

In OLS we wish to estimate the parameter β in the following equation:

$$y = x\beta + \epsilon, \quad (5.14)$$

or in vector notation:

$$y = \vec{x}^\top \vec{\beta} + \epsilon. \quad (5.15)$$

Equations (5.14) and (5.15) are linear in the coefficient β ; we are considering linear OLS.

Estimation of β is done through minimization of the sum of squared residuals:

$$\min_{\beta} \sum_{i=1}^N (y_i - x_i\beta)^2. \quad (5.16)$$

An explicit solution to this optimization problem exists. Let $\vec{y}$ be the column vector with observations of y and X the matrix where each row is an observation of $\vec{x}$. Then the optimal coefficients $\hat{\beta}$ are given by:

$$\hat{\beta} = (X^\top X)^{-1} X^\top \vec{y}.$$

Lemma 5.2. When $\mathbb{E}[\epsilon_i|X] = 0$ for all i , $\hat{\beta}$ is unbiased

Proof. Substituting $\vec{y} = X\beta + \vec{\epsilon}$, where $\vec{\epsilon}$ is the vector of all residuals ϵ_i , into $\hat{\beta} = (X^\top X)^{-1} X^\top \vec{y}$ and taking expectation yields:

$$\begin{aligned} \mathbb{E}[\hat{\beta}] &= \mathbb{E} \left[(X^\top X)^{-1} X^\top (\vec{\epsilon} + X\beta) \right] \\ &= \mathbb{E} \left[(X^\top X)^{-1} X^\top \vec{\epsilon} + \beta \right] \\ &= \beta + \mathbb{E} \left[\mathbb{E} \left[(X^\top X)^{-1} X^\top \vec{\epsilon} | X \right] \right] \\ &= \beta + \mathbb{E} \left[(X^\top X)^{-1} X^\top \mathbb{E}[\vec{\epsilon} | X] \right] \\ &= \beta, \end{aligned}$$

which shows that $\hat{\beta}$ is unbiased. □

Furthermore, when all ϵ_i have the same variance σ^2 and they are uncorrelated with each other, the Gauss-Markov theorem states that OLS yields the minimum variance unbiased estimator for β [24]. The following example gives an idea of the size of the variance of the estimator.

Example 5.1 (Variance of the least squares estimator). Suppose we try to fit $y = \beta_0 + \beta_1 x + \epsilon$, where ϵ has variance σ^2 . Then $\text{Var}[\beta_1] = \frac{\sigma^2}{\sum_{i=1}^N (x_i - \frac{1}{n} \sum_{i=1}^n x_i)^2}$. Adapted from example 4.2 from [24].

Example 5.1 shows that the variance of the coefficients decreases as the sample size grows and as the residuals are larger. The example holds for the one-dimensional case, there is only one independent variable. When there are more independent variables, the variance of the coefficients depends on the correlation between variables. When two variables are more positively correlated, variance of the regression coefficients will be higher. In small samples, this will lead to inaccurate estimation of the coefficients [24].

Linear OLS has desirable properties, it is unbiased and it is the minimum variance unbiased estimator. In small samples with many features, however, the OLS estimator can be unstable. It is possible to mitigate the unstable behaviour of OLS by introducing some bias, by using the lasso estimator.

5.A.2. Lasso regression

Tibshirani [42] introduced the lasso, the least absolute shrinkage and selection operator, to improve two aspects of OLS: prediction accuracy and interpretation. Prediction accuracy refers to the problem mentioned in the previous section: in small samples with many features, the variance of the OLS estimate is high. Interpretation refers to determining which independent variables have the strongest effects on the dependent variable. By shrinkage of coefficients, often setting them to zero, the lasso yields more interpretable results.

Definition 5.4 (Lasso estimator). When the functional form is $y = \vec{x}^\top \vec{\beta} + \epsilon$, the lasso estimator $\hat{\beta}$ for $\vec{\beta}$ is given by:

$$\hat{\beta} = \arg \min_{\vec{\beta}} \sum_{i=1}^N (y_i - \vec{x}_i^\top \vec{\beta})^2, \quad (5.17)$$

$$\text{s.t.} \quad \|\vec{\beta}\|_1 \leq t. \quad (5.18)$$

The definition of the lasso is as that of OLS, but with a penalty term. The $L1$ -norm of the coefficient vector $\vec{\beta}$ has to be smaller than or equal to t . An alternative definition of the lasso is given by:

$$\min_{\vec{\beta}} \sum_{i=1}^N (y_i - \vec{x}_i^\top \vec{\beta})^2 + \lambda \|\vec{\beta}\|_1. \quad (5.19)$$

λ is called the regularisation parameter and it determines the trade-off between bias and variance of the coefficients. When $\lambda = 0$ we have OLS and hence an unbiased estimator. For $\lambda > 0$ we have a penalty on the size of the coefficients, shrinking the coefficients and inducing sparsity: many coefficients will be zero.

No explicit solutions exist for the lasso problem, however, equation (5.19) represents a convex optimization problem, which can efficiently be solved using numerical methods, see for example [17].

Choosing the regularisation parameter

The optimal value of the regularisation parameter λ is not known. To find a suitable value for λ we can split the data in a training set and a validation set. We fit the model for different values of λ on the training set, and test the fitted model on the validation set. The value of λ for which $\sum_{i=1}^N (y_i - x_i \beta_\lambda)^2$ of data in the validation set is smallest, is the λ we pick.

When we only have a small dataset, we wish to use all the data, and not use a part as validation set. In this case we can use k -fold cross validation. In cross validation, we split the data into k sets. For different values of λ , we fit the model on all the data except for the j -th set. We use the j -th set as validation set and record $\sum_{i \in j\text{-th set}} (y_i - x_i \beta_\lambda)^2$ for every λ . We do this for every set from 1 to k and take the average over the sums of squared errors, we then pick λ for which this mean is smallest.

The lasso is a useful technique when dealing with high-dimensional data, with automatic feature selection and shrinking, reducing variance in the estimated coefficients. Computationally, the method is slightly less efficient than OLS, especially when cross-validation is used for finding the optimal value of the regularisation parameter.

As with OLS, the functional form fitted by the lasso is determined by the user. In some cases, especially in the high-dimensional case, it can be hard to guess a proper functional form. For these cases, we introduce support vector regression.




Figure 5.7: Geometric interpretation of support vector regression. Left graph shows a linear function with the corresponding ϵ -bounds. The right graph indicates how errors are penalized. Figure 1 from [40].

5.A.3. Support Vector Regression

Support vector regression is an application of the support vector machine, developed by Vapnik and Chervonenskis in 1963, however the current standard was described by Cortes and Vapnik [14]. We define SVR following Smola and Schölkopf [40]. We start with the problem where the function we want to estimate is linear.

Definition 5.5 (Support vector regression for linear functions). We estimate the coefficients β in the function $y = x^\top \beta + b$ by solving:

$$\min_{\beta} \frac{1}{2} \|\beta\|_2^2 + C \sum_{i=1}^N (\xi_i + \xi_i^*), \quad (5.20)$$

$$\text{s.t.} \quad \begin{cases} y_i - x_i^\top \beta - b \leq \epsilon + \xi_i \\ x_i^\top \beta + b - y_i \leq \epsilon + \xi_i^* \\ \xi_i, \xi_i^* \geq 0. \end{cases} \quad (5.21)$$

In SVR, we seek small β , hence the minimization of the L_2 -norm of β . Furthermore, we accept a small error in the estimated function: when $|y_i - x_i^\top \beta - b| < \epsilon$, the estimation is good enough. When $|y_i - x_i^\top \beta - b| > \epsilon$, we introduce the slack variables ξ_i, ξ_i^* . The slack variables are given by:

$$\begin{aligned} \xi_i &= y_i - x_i^\top \beta - b - \epsilon, & y_i - x_i^\top \beta - b > \epsilon \\ \xi_i^* &= x_i^\top \beta + b - y_i - \epsilon, & y_i - x_i^\top \beta - b < -\epsilon. \end{aligned}$$

The parameter C determines the trade-off between simplicity, i.e., minimizing beta, and the sum of residuals, i.e., minimizing the slack variables.

A geometrical interpretation of the SVR problem can be found in Figure 5.7. Figure 5.7 shows a linear function with the ϵ -bounds and the slack variable ξ . Points outside of the ϵ -band, the slack variables, are penalized according to the function on the right of the figure.

The problem defined by equations (5.20) and (5.21) has a dual formulation. Skipping details, we can present the solution for β given by the dual problem: $\beta = \sum_{i=1}^N (\alpha_i - \alpha_i^*) x_i$, and therefore the estimated function is given by:

$$y = \sum_{i=1}^N (\alpha_i - \alpha_i^*) x_i^\top x_i + b, \quad (5.22)$$

where the coefficients α_i, α_i^* are subject to $\sum_{i=1}^N (\alpha_i - \alpha_i^*) = 0$ and $\alpha_i, \alpha_i^* \in [0, C]$.

The model coefficients β can be completely described by the data points x_i . Furthermore, α_i, α_i^* are zero when $|y_i - x_i^\top \beta - b| < \epsilon$, which means the model is completely determined by the data points that are outside of the ϵ -tube: the support vectors.

Under certain conditions, we can show that SVR is equivalent to L_1 -norm minimization, which we show in the following lemma.

Lemma 5.3. Assume we take $\epsilon = \frac{1}{C}$, then for the limit $C \rightarrow \infty$, the SVR for linear function given by Definition 5.5, is equivalent to:

$$\min_{\beta} \sum_{i=1}^N |y_i - x_i^\top \beta - b|. \quad (5.23)$$

Proof. As $\epsilon = \frac{1}{C}$, the first constraint from (5.21) is given by:

$$y_i - x_i^\top \beta - b \leq \frac{1}{C} + \xi_i.$$

Taking limits yields:

$$\begin{aligned} y_i - x_i^\top \beta - b &\leq \lim_{C \rightarrow \infty} \left(\frac{1}{C} + \xi_i \right), \\ y_i - x_i^\top \beta - b &\leq \xi_i. \end{aligned}$$

As we try to minimize ξ_i , we can combine the first and third constraint and replace the inequality by equality:

$$y_i - x_i^\top \beta - b = \xi_i, \quad \xi_i \geq 0 \quad (5.24)$$

Using the same reasoning for the second constraint, we obtain:

$$x_i^\top \beta + b - y_i = \xi_i^*, \quad \xi_i^* \geq 0. \quad (5.25)$$

Plugging equations (5.24) and (5.25) into equation (5.20) yields:

$$\min_{\beta} \frac{1}{2} \|\beta\|_2^2 + C \sum_{i=1}^N |y_i - x_i^\top \beta - b|.$$

As C is a constant, we can divide by C to obtain:

$$\min_{\beta} \frac{1}{2C} \|\beta\|_2^2 + \sum_{i=1}^N |y_i - x_i^\top \beta - b|.$$

Taking the limit yields the final result:

$$\begin{aligned} \lim_{C \rightarrow \infty} \min_{\beta} \frac{1}{2C} \|\beta\|_2^2 + \sum_{i=1}^N |y_i - x_i^\top \beta - b|, \\ \min_{\beta} \sum_{i=1}^N |y_i - x_i^\top \beta - b| \end{aligned}$$

□

Non-linear functions

Until now, we have considered the linear version of SVR. For the non-linear function of the SVR, we need to map the data into a (high-dimensional) feature space. For example, if we have one-dimensional data x , we can use the map $\Phi(x) = [x, x^2]$ to go from a linear to a quadratic function.

The power of SVR is that we do not need an explicit map $\Phi(\cdot)$: we can replace the inner product $x^\top x_i$ in equation (5.22) by the kernel function $k(x, x_i) := \Phi(x)^\top \Phi(x_i)$. We do not need to know $\Phi(\cdot)$ as long as we know the kernel $k(\cdot, \cdot)$.

Examples of kernels are:

- The linear kernel: $k(x, x') = x^\top x'$. Trivial kernel, corresponds to linear functions.
- The polynomial kernel: $k(x, x') = (x^\top x' + r)^n$, where $r \in \mathbb{R}$ and $n \in \mathbb{N}$ are variables to be specified. The functional form is a polynomial of order n .
- The radial basis function (rbf) kernel: $k(x, x') = \exp(-\gamma |x - x'|^2)$ where $\gamma > 0$. This kernel can be interpreted as we interpret statistical kernels, where this one corresponds to the normal density kernel. It is a weighting function that gives large weights to points x_i close to x and small weights to points far away from x . This function is useful when one want to avoid imposing a functional form on the regression. The radial basis function described here and used in this thesis is known as the Gaussian radial basis function.

- The sigmoid kernel: $\tanh(\langle x, x' \rangle + r)$ where $r \in \mathbb{R}$. Like the rbf kernel, the sigmoid kernel weights the input, but not symmetric like the rbf kernel. When using the sigmoid kernel, it is unclear whether the link between the primal and dual optimization problem always exists. The reason this kernel is used is its link with neural networks, however it does not seem to perform better than the rbf kernel.

The rbf kernel is the current standard in support vector regression and it has universal approximation properties, as shown by Park and Sandberg [36]. We were not able to generalize Lemma 5.3 to general kernel functions and therefore show we can approximate any function in the $L1$ -norm. Still, the results by Park and Sandberg have us convinced that the rbf kernel is suitable when we do not wish to impose any functional form on our regression. We will therefore focus on using the Gaussian rbf kernel in our research.

As with the lasso, there are parameters for which we have to chose values. These are ϵ and C in equations (5.20) and (5.21), and γ in the rbf kernel. We will use cross-validation to determine these values.

5.B. Appendix to Chapter 5

μ_1	-0.1047
μ_2	0
μ_3	0.1262
λ	0.0036
$w_{1,t}$, for all t	0.8
$w_{2,t}, w_{3,t}$ for all t	0.1
γ_{11}	0.0101
γ_{12}	0.0045
γ_{13}	0.0106
γ_{22}	1.7062
γ_{23}	0.5716
γ_{33}	0.0602
α_{11}	0.1406
α_{12}	0.0448
α_{13}	-0.0463
α_{22}	0.1342
α_{23}	-0.0339
α_{33}	0.0645
β_{11}	0.8434
β_{12}	0.8649
β_{13}	0.3286
β_{22}	0.1016
β_{23}	0.2941
β_{33}	0.9205

Table 5.1: Values for coefficients used in equations (5.12) and (5.13). Note that coefficients with a double subscript, e.g. γ_{ij} , are symmetrical: $\gamma_{ij} = \gamma_{ji}$.

6

Numerical experiments

In this chapter, we apply the life-cycle construction algorithm, that is, the forward algorithm given by equation (4.16) and the backward algorithm, Algorithm 3, to several density forecasts generated by models of increasing complexity.

This chapter consist of three parts: a description of numerical issues encountered with the implementation of the algorithm and presentation of results for two different asset return models; normal, independent asset returns and asset returns as modelled by the Ortec Finance Scenarios.

6.1. Numerical issues in implementation

In this section, we discuss the implementation of the mean-variance optimization of the replacement ratio. The optimization is done through the algorithm described in Section 4.5. The algorithm consist of two parts: the forward, given by equation (4.16), and backward algorithm, Algorithm 3. For both parts it is essential we can estimate the conditional expectation of asset returns, of which we will discuss the details first. Next, we discuss the optimization method used in both algorithms. Finally, we discuss the estimation of value functions necessary in the backward algorithm.

6.1.1. Estimation of conditional expectation of asset returns

In the forward and backward algorithm we need the values of the conditional expectation of asset returns: $\mathbb{E}[\tilde{R}_t | Z_t]$ and $\mathbb{E}[\tilde{R}_t \tilde{R}_t^\top | Z_t]$, see for example equation (4.24). In Chapter 5 we investigated methods of how to estimate conditional expectation in a simulation framework. We will use two of the methods investigated in Chapter 5, the lasso and SSA.

In Section 5.4, the lasso has shown to be the most promising regression method for the estimation of conditional expectation and is therefore chosen for the implementation. As all regression methods, however, the lasso has a drawback. The drawback is in the estimation of $\mathbb{E}[\tilde{R}_t \tilde{R}_t^\top | Z_t]$. In the case that there is more than one risky asset, $\mathbb{E}[\tilde{R}_t \tilde{R}_t^\top | Z_t]$ is a semi-positive definite matrix. Semi-positive definiteness (SPD) is a property essential for the functioning of the optimization algorithm. When we numerically estimate $\mathbb{E}[\tilde{R}_t \tilde{R}_t^\top | Z_t]$ through the lasso, the SPD property can be lost, which makes the algorithm fail.

There exist two solutions for the lack of the SPD property: use SSA or transform the estimated matrix $\mathbb{E}[\tilde{R}_t \tilde{R}_t^\top | Z_t]$ such that it has only non-negative eigenvalues. SSA uses sample averages to estimate the conditional expectation and is therefore guaranteed to preserve the SPD property. Transforming $\mathbb{E}[\tilde{R}_t \tilde{R}_t^\top | Z_t]$ to have only positive eigenvalues can be done efficiently calculating the eigenvalue decomposition and setting the negative eigenvalues to zero, see [37], Section 3.

6.1.2. Convex optimization algorithm

Using the lasso with negative eigenvalue correction or SSA for the estimation of the conditional expectation of asset returns means the optimization problem in both the forward and backward algorithm is a constrained convex optimization problem. The optimization problem is a quadratic problem given by equation (4.16) for the forward algorithm and equation (4.26) for the backward algorithm, constrained by the no-leverage and

Table 6.1: Statistics of replacement ratio and the squared difference between the replacement ratio and the target, $(RR - \gamma)^2$, for dynamic asset allocation. Standard error, defined as sample standard deviation divided by the square root of the sample size, displayed in brackets when applicable.

	Mean of $(RR - \gamma)^2$	Mean of RR	Median of RR	Minimum RR	Maximum RR
Forward	$2.40 * 10^{-4}$ ($4.5 * 10^{-5}$)	69.6% ($3.4 * 10^{-2}$)	70.0%	49.8%	71.1%
Backward 1	$2.19 * 10^{-4}$ ($2.8 * 10^{-5}$)	69.6% ($3.2 * 10^{-2}$)	69.7%	52.5%	80.4%
Backward 2	$1.25 * 10^{-4}$ ($1.2 * 10^{-5}$)	69.6% ($2.4 * 10^{-2}$)	69.8%	56.8%	77.6%
Backward 3	$0.85 * 10^{-4}$ ($0.7 * 10^{-5}$)	69.6% ($2.0 * 10^{-2}$)	69.8%	61.1%	75.2%

no-shorting constraints. The no-leverage constraint is an equality constraint, the no-shorting constraint is an inequality constraint. The constrained quadratic problem is convex and can therefore be efficiently solved by the Python package *cvxopt*, which we will use in our implementation.

6.1.3. Estimation of the value function

The value function as defined in Problem 4.4, equation (4.15a), is estimated four times per time step in the backward algorithm. In Chapter 5 we showed how different regression methods perform in the estimation of the value function. The results indicate that the lasso is the preferred method for estimation of the value function. When implementing the backward algorithm, however, OLS is preferred for estimation of the function that maps wealth values to the value function, equation (4.25), as non-zero coefficients are required for the backward algorithm to converge.

Furthermore, for stability reasons, a bundled regression approach is necessary for the estimation of the value function. A bundled approach means we sort our samples in to a number of bundles of equal size, on which we separately estimate the value function. Bundling is used in [12] and is shown to be beneficial, especially in the presence of constraints. We bundle by sorting on the values of W_t , the wealth at the time on which is conditioned. By bundling we increase the accuracy of the backward algorithm and we therefore use it in all our tests.

Finally, we use the entire state-space, so all previous values of wealth and interest rates, for estimation of the value function. We choose to do so because the results from Chapter 5 indicate that using the entire state-space is beneficial when asset returns are not independent over time.

6.2. Normal, independent asset returns and constant risk-free rate

With the details of implementation completed, we can present the results of our numerical tests. The first test is done using risky asset returns that are normally distributed, independent over time. The risk-free rate is constant for all tenors and for all times. We use two risky assets, equity and bonds, and a risk-free, of which the distribution is given by:

$$\vec{R}_t \sim \mathcal{N}(\vec{\mu}, \Sigma) \text{ with } \vec{\mu} = [1.08, 1.05, 1.043]^\top \text{ and } \Sigma = \begin{bmatrix} 0.2^2 & -0.015 & 0 \\ -0.015 & 0.1^2 & 0 \\ 0 & 0 & 0 \end{bmatrix}. \text{ Here the third return asset is risk-}$$

free cash, uncorrelated with the other variables.

We generate $M = 2000$ sample paths, and run 3 backward updates. For the estimation of the conditional expectation of returns we use SSA with 1 cell, i.e., the sample average. We use 10 bundles in the estimation of the value function. The investment horizon is 40 years, there are 40 rebalancing moments and the target replacement ratio is 70%. This target matches what is expected of the ingoing cash flows in the current Dutch pension system. For the incoming cash flows and calculation of the replacement ratio we use the assumptions presented in Section 3.4.2.

We present the asset allocation from the forward run and the one after three backward updates, we'll also provide a table with statistic of the replacement ratio and the value function.

Figure 6.1 shows the mean asset allocation over a 2000 paths. The mean of the asset allocation creates a life-cycle with equity at the start and just cash at the end, the classic life-cycle. Figure 6.2 shows the asset allocation after three backward updates, which looks very similar to Figure 6.1. From Table 6.1 it becomes clear that the backward algorithm improves the forward solution. This improvement is the decrease in the mean of $(RR - \gamma)^2$, while the summary statistics of the replacement ratio have relatively little change. The decrease in the mean of $(RR - \gamma)^2$ shows we have decreased variance of the replacement ratio while keeping the average stable.




Figure 6.1: Mean of asset allocation over time after forward algorithm. Independent returns and constant risk-free.




Figure 6.2: Mean of asset allocation over time after 3 backward updates. Independent returns and constant risk-free.

Table 6.2: Statistics of replacement ratio and the squared difference between the replacement ratio and the target, $(RR - \gamma)^2$, for deterministic asset allocation. Standard error, defined as sample standard deviation divided by the square root of the sample size, displayed in brackets when applicable.

	Mean of $(RR - \gamma)^2$	Mean of RR	Median of RR	Minimum RR	Maximum RR
Forward	$2.29 * 10^{-3}$ ($8.4 * 10^{-5}$)	70.1% (0.1)	69.7%	58.0%	94.2%
Backward 1	$2.08 * 10^{-3}$ ($7.4 * 10^{-5}$)	70.0% (0.1)	69.5%	58.2%	91.7%
Backward 2	$1.96 * 10^{-3}$ ($6.9 * 10^{-5}$)	70.0% (0.1)	69.5%	58.5%	90.4%
Backward 3	$1.92 * 10^{-3}$ ($6.7 * 10^{-5}$)	69.9% (0.1)	69.5%	58.6%	89.8%

The results presented in Table 6.1 are from the dynamic optimization, every path has their own allocation. We also present the result when we use the mean asset allocation of all paths, as a deterministic life-cycle. The results can be found in Table 6.2. In this table we see that when we use the mean asset allocation for all paths, the backward solutions perform better than the forward one. Comparing Table 6.1 and Table 6.2, we find that the mean of the squared difference between replacement ratio and target has increased by a factor 10 for a deterministic asset allocation, but the mean of the replacement ratio has come up. This means that variance of the replacement ratio will have increased by using a deterministic asset allocation. This is also visible in the minimum and maximum RR: their difference is approximately 35 percent-point for the deterministic asset allocation, compared to 20 percent-point for the dynamic asset allocation. The minimum and maximum both have increased, however, which in the eyes of the investor is good. For this increase, he has to take on some extra variance.

Besides the mean of the asset allocation over all paths, we are also interested in behaviour of the asset allocation for a single path. A comparison of Table 6.1 and 6.2 shows that using the deterministic asset allocation, the mean over all paths, yields slightly higher average replacement ratio, but it is necessary to take on more variance. A dynamic approach can therefore be a better alternative. A dynamic asset allocation, however, may have undesirable properties, such as large instability in the asset allocations or extremely risky allocation near the end of the investment horizon. To check whether dynamic asset allocations have desirable properties, we compare the asset allocation for the minimum, median and maximum realization of the replacement ratio.

The asset allocations of minimum, median and maximum replacement ratio can be found in Appendix 6.A, Figure 6.8, 6.9 and 6.10. Three things stand out: firstly, asset allocation is volatile, secondly, for the minimum replacement ratio the final asset allocation is risky and finally, for both the median and maximum allocation the entire wealth is put into the riskless asset in the last periods. Volatility of the asset allocation is caused by slight variation in the sample mean of the asset returns, which varies over time. The variation of the mean of asset return causes small but visible changes in the asset allocations, which shows up as volatility. This volatility can be damped by using analytical calculations of conditional expectation. The riskiness of the final asset allocation for the minimum replacement ratio is a consequence of the replacement ratio being far from the target, and hence much risk is necessary to get close to the target. Finally, in the median and maximum case, the target replacement ratio can be obtained by investing all wealth in the riskless asset and as the riskless asset is non-stochastic, this is always preferred.

The decrease in risk-taking as soon as reaching the target is possible with only non-risky assets is one of the main strengths of the target-based dynamic asset allocation. By decreasing risk-taking when we are close to target, we can reach the target with little variance. With a deterministic asset allocation, we would not automatically decrease risk when we are close to target and hence not have this decrease in variance.

We think that when the volatility in the asset allocation is smoothed out, the dynamic asset allocations are suitable for real investment purposes. A method to reduce volatility in the asset allocation is using regress-later, which we investigate in the next section.

6.2.1. Regress-later

In Chapter 5, we introduced two types of regression techniques: regress-now and regress-later. With regress-now we calculate conditional expectation from data using the values of variables on which we condition. This technique is used in the previous tests. In this section, we use regress-later, where we calculate conditional expectation from data using future values of variables on which we condition.

In Chapter 5 we showed that, when the conditional expectation of asset returns is not analytically tractable, regress-now is necessary for the estimation of conditional expectations in Algorithm 3. For normal indepen-




Figure 6.3: Mean of asset allocation over time from forward algorithm. Independent returns and constant risk-free. Conditional expectations of asset returns are calculated analytically.

dent asset returns, however, the conditional expectation of the returns is analytically tractable. Hence we can compare regress-now and regress-later.

For our test of regress-later, we use the analytical conditional expectation of asset returns in both the forward and the backward algorithm. Regress-later, as defined by Definition 5.2, is applied to the estimation of the conditional expectation of the value function, given by $V_t = \mathbb{E}[(RR - \gamma)^2 | Z_t]$, in step 2, 3 and 4 of the backward algorithm, Algorithm 3. We compare our findings to those of the previous section.

The result of using analytical conditional expectations of asset returns in the forward algorithm are unsurprising: the mean asset allocation is more smooth than for estimated conditional expectations, compare Figure 6.1 to Figure 6.3.

The use of regress-later in the backward algorithm does not improve the quality of the algorithm, there is no convergence any more. This contradicts results from literature, where regress-later is more stable than regress-now. Unlike existing literature, however, we do not optimize wealth, we optimize the replacement ratio. The influence of interest rates, despite being constant and non-stochastic, on the calculation of the replacement ratio could diminish the positive effect of regress-later on the convergence. This result strengthens our believe that regress-now is a suitable method for the estimation of conditional expectation in our algorithm.

Having investigated the performance of the algorithm with normal, independent asset returns, we investigate a more complicated model with non-normal asset returns which are dependent on previous values and, on top of that we add stochastic interest rates correlated with asset returns. These asset returns and interest rates are obtained from the Ortec Finance Scenarios and the results are presented in the next section.

6.3. Ortec Finance Scenario returns

The Ortec Finance Scenarios (OFS) returns are generated with model stacking of several approaches. The model facilitates: fat tails, time-varying volatility and is suitable for high-dimensional forecasting on both short and long horizons [26]. Describing the complete model is beyond the scope of the thesis.

The OFS generates monthly returns on a wide range of assets and monthly levels of interest rates, inflation and economic growth for different parts of the world. For our test asset returns and interest rate levels are the quantities of interest. Our minimum asset holding period is one year, and hence we convert the monthly values to yearly values.

In our test we use the yearly return on European equity and the nominal interest rates of zero-coupon German government bonds as generated by the OFS. European equity is the most general choice for a Euro-denominated risky asset and German government bonds are considered as the most save (from default) Euro-denominated asset [19]. We use the rates on German government bonds to calculate the replacement ratio but also the calculate the return on German government bonds of different tenors. We do not choose specific tenors to invest in, but we create a basket which covers the liabilities generated by the replacement ratio, the liability covering portfolio (lcp).

The lcp consists of zero-coupon bonds which mature exactly at a pension payment date. By combining bonds that mature at different pension payout dates we obtain a portfolio that exactly covers the payments needed by a retiree. By choosing the lcp as the less risky asset we do not have to chose which tenor bond we invest in, we automatically invest in the bonds that perfectly correlate with the pension payments. Note that as time goes by the average tenor of bonds in the lcp decreases, because retirement approaches, and therefore the volatility of the lcp decreases as time goes by. Volatility decreases because bonds with shorter tenors are less sensitive to interest rate changes, as described in Section 3.3.1. Furthermore, when all wealth is invested in the lcp, the volatility of the replacement ratio is zero: any change in interest rates which changes the calculation of the replacement ratio is offset by value changes of the lcp.

In the tests with OFS asset returns and interest rates, we use 2000 paths, which we divide into 4 bundles for estimation of conditional expectations. Both conditional expectations, of returns and the value function, are calculated using the lasso.

Figure 6.4 and 6.5 show the mean asset allocation of 2000 paths from the forward algorithm and after 3 backward updates. The algorithm converges, the mean of the quadratic difference between replacement ration and target shows a sharp decrease between de forward and the third backward iteration. Furthermore, after the first 15 years, the mean asset allocation shows the classic life-cycle property: equity allocation decreases. Nevertheless, the asset allocation is erratic. The forward solution seems reasonable, the improved backward solution, however, is non-monotonic and volatile. Nonetheless, we can explain this behaviour and how to resolve it.

The first reason for the volatile behaviour of the asset allocation is the modelling of interest rates: the OFS predicts low rates in the short-term and then an increase towards a higher long-term expectation. An increase in rates means a negative return on bonds and, as our lcp consist of long-term bonds, the change in interest rates is accompanied by high volatility in the lcp. As a result we might expect that in the first years we mainly invest in equity, but we clearly do not, as shown by Figure 6.5. This can be the result of the negative correlation between lcp and replacement ratio, which makes the lcp a good hedge against interest rate changes. Given the contradictory effects of rising rates on the attractiveness of the lcp as an asset, we cannot be certain whether this causes the erratic asset allocation in the first 10 years.

A second reason for the volatile behaviour of the asset allocation is the estimation of the value function in the first periods. The value function is the conditional expectation of the quadratic difference between the replacement ration and the target; in the first periods the estimate of this value function is challenging: the wealth values are not predictive of final wealth and interest interest rate values have insignificant correlation with the interest rate values used to calculate the replacement ratio. This can cause our estimation of the value function to be unstable and therefore cause the algorithm to produce insufficient results. Bundling should be able to reduce this issue and when that is not enough, the number of samples should be increased.

Both causes of the erratic asset allocation can be solved by reducing the investment horizon. By optimizing only the last 15 years we have steady-state asset returns and we increase the correlation between first period wealth values and the final value function.

Besides decreasing the investment horizon, we also decrease the number of rebalancing moments. This should smooth out variability in asset returns and therefore cause a less volatile asset allocation. We decrease




Figure 6.4: Mean of asset allocation over time from forward algorithm. OFS returns and interest rates. Lasso was used for all estimations of conditional expectation.




Figure 6.5: Mean of asset allocation over time after three backward updates. OFS returns and interest rates. Lasso was used for all estimations of conditional expectation.

Table 6.3: Statistics of replacement ratio and the squared difference between the replacement ratio and the target, $(RR - \gamma)^2$, for OFS returns, using a dynamic asset allocation. Standard error, defined as sample standard deviation divided by the square root of the sample size, displayed in brackets when applicable.

	Mean of $(RR - \gamma)^2$	Mean of RR
5 rebalancing moments, forward solution	3.11×10^{-2} ($8.0 * 10^{-4}$)	37.0% (0.27)
5 rebalancing moments, 3 backward updates	3.07×10^{-2} ($7.8 * 10^{-4}$)	36.9% (0.26)
15 rebalancing moments, forward solution	3.15×10^{-2} ($9.7 * 10^{-4}$)	38.8% (0.31)
15 rebalancing moments, 3 backward updates	2.54×10^{-2} ($6.9 * 10^{-4}$)	39.2% (0.26)

the rebalancing frequency to three years, so, over 15 years we have five rebalancing moments. We furthermore assume the investor starts building up wealth 15 years before retirement, hence the accumulated wealth will be lower and we chose a lower target for the replacement ratio: 50%.

The resulting mean asset allocations are visible in Figure 6.6 and 6.7. The mean allocation following from the forward algorithm is less volatile than previously and shows classic life-cycle properties. The same holds for the mean asset allocation after three backward updates, it is smooth and monotonic. We conclude that reducing the number of rebalancing moments increases the quality of appearance of the asset allocation.

Although appearance improved by using less rebalancing moments, performance in terms of mean squared difference between replacement ratio and target has decreased. Table 6.3 shows that using 15 rebalancing moments performs better than 5 rebalancing moments, i.e., the mean of $(RR - \gamma)^2$ is smaller for 15 rebalancing moments than for 5 rebalancing moments. Note that this only holds after backward updating, the solutions following from the forward algorithm have approximately equal results. Equality in results from the forward algorithm is caused by the assumption under which the forward algorithm operates: all wealth is invested in the risk-free after the current period. With 15 rebalancing moments, the assumption that all wealth is invested in the risk-free has more influence, as there are more rebalancing moments which require use of the assumption.

All in all we can conclude that using less rebalancing moments yields a more attractive looking asset allocation, but in terms of performance it is worse. In the end, it is up to the investor to decide the number of rebalancing moments.

Finally, we investigate how does the mean of the asset allocation perform as a deterministic life-cycle. We test the mean of the asset allocation with 5 rebalancing moments as a deterministic life-cycle, the resulting statistics are given in Table 6.4. Comparing Table 6.4 to Table 6.3 shows that using a deterministic asset allocation performs only slightly worse than using the dynamic asset allocation, as measured by the mean of the squared difference between the replacement ratio and the target. This shows that using the mean dynamic asset allocation as a deterministic life-cycle is a reasonable solution to the life-cycle problem.

Despite the small difference between the dynamic and deterministic final result, there can be big differences on individual paths. Every path has its own allocation when the asset allocation is dynamic, paths where returns are higher than expected in the first periods, will have less risky asset allocations in following periods. This less risky asset allocation will reduce variance of wealth and therefore the replacement ratio over the path. When the asset allocation is deterministic, such a reduction of variance over the path would not take place.

We compare the deterministic asset allocation to Bogle's rule, presented in Section 3.4.1. Results for Bogle's rule can be found in Table 6.4. Bogle's rule performs worse than the mean asset allocation in terms of average quadratic difference between target and replacement ratio. Also, the mean of the replacement ratio is higher for the mean dynamic asset allocation than for Bogle's rule. Moreover, comparing Table 6.3 with Table 6.4, we find that the dynamic asset allocation performs better than Bogle's rule, implying that using a dynamic asset allocation is better than the current standard in life-cycle investing.

6.4. Conclusion

In this chapter we have tested the forward algorithm, given by equation (4.16) and the backward algorithm, Algorithm 3, for two types of asset return models. The first conclusion is that the algorithm does indeed converge, for every backward update we perform, the solution gets better.

Besides convergence, we found that asset allocations can be volatile, caused by varying estimations of the conditional expectation of asset returns. Although variation in estimates of conditional expectations should




Figure 6.6: Mean of asset allocation over time from forward algorithm. Investment horizon of 15 years with 5 rebalancing moments. OFS returns and interest rates. Lasso was used for all estimations of conditional expectation.




Figure 6.7: Mean of asset allocation over time after three backward updates. Investment horizon of 15 years with 5 rebalancing moments. OFS returns and interest rates. Lasso was used for all estimations of conditional expectation.

Table 6.4: Statistics of replacement ratio and the squared difference between the replacement ratio and the target, $(RR - \gamma)^2$, for OFS returns, using a deterministic asset allocation. Standard error, defined as sample standard deviation divided by the square root of the sample size, displayed in brackets when applicable.

	Mean of $(RR - \gamma)^2$	Mean of RR
5 rebalancing moments, forward solution	3.2×10^{-2} ($7.5 * 10^{-4}$)	37.6% (0.29)
5 rebalancing moments, 1 backward update	3.2×10^{-2} ($7.4 * 10^{-4}$)	37.4% (0.29)
5 rebalancing moments, 2 backward updates	3.2×10^{-2} ($7.4 * 10^{-4}$)	37.4% (0.28)
5 rebalancing moments, 3 backward updates	3.2×10^{-2} ($7.4 * 10^{-4}$)	37.4% (0.29)
Bogle's rule	3.3×10^{-2} ($8.6 * 10^{-4}$)	36.3% (0.27)

decrease when using regress-later, regress-later did only improve the asset allocation found using the forward algorithm, not the solutions after backward updates. We conclude that using regress-now suffices for our algorithm.

Volatility in the asset allocation can be suppressed by using less rebalancing moments, this does however worsen the performance of the solution.

Finally, we investigated how the average of the dynamic asset allocations obtained using the algorithm performed as a deterministic asset allocation. We found that the deterministic allocation improves on an example of the current standard in life-cycle asset allocation: Bogle's rule. We therefore believe that using the mean of the dynamic asset allocation is a good approach to finding a deterministic asset allocation for life-cycle investing. Nevertheless, the deterministic asset allocation performs worse than the dynamic asset allocation and therefore a dynamic approach to life-cycle investing is preferred.

6.A. Appendix to Chapter 6




Figure 6.8: Asset allocation over time for the minimum replacement ratio. Independent returns and constant risk-free.




Figure 6.9: Asset allocation over time for the median replacement ratio. Independent returns and constant risk-free.




Figure 6.10: Asset allocation over time for the maximum replacement ratio. Independent returns and constant risk-free.

Conclusion

In this chapter, we present our conclusions and give recommendations for further research.

7.1. Conclusions

In this thesis, we have investigated how mean-variance optimization can be used for life-cycle investing. More specifically, how we can use mean-variance optimization to construct a life-cycle with an optimal replacement ratio, i.e., the fraction of labour income that can be replaced by a pension.

Contrary to common investing wisdom, mean-variance optimization does not lead to deterministic life-cycles with a decreasing risk over time. When we introduced mean-variance optimization for the replacement ratio, in Chapter 3, we found the same result: risk-taking increases as time goes by. This result is surprising as the replacement ratio is negatively correlated with, less risky, bonds, which should make investing in bonds attractive as time passes. Despite the hedging property of bonds, mean-variance optimal asset allocations for the replacement ratio can be increasing in risk over time. We concluded that using mean-variance optimization for deterministic life-cycles is not a suitable approach to the life-cycle problem.

As a solution, we introduced dynamic mean-variance optimization such that it can be used to optimize the replacement ratio in Chapter 4. We adjusted the dynamic mean-variance optimization algorithm by Cong and Oosterlee [12] to be able to optimize the replacement ratio. We also generalized the algorithm to handle problems where asset returns can be driven by any process.

We compared several methods for estimation of conditional expectation in a simulation setting in Chapter 5. We concluded that lasso regression is the preferred method for estimating conditional expectations when the asset returns are generated by involved models, for example autoregressive models. Using these results we improved the developed algorithm for dynamic mean-variance optimization of the replacement ratio.

Numerical experiments in Chapter 6 showed that our dynamic mean-variance optimization algorithm for the replacement ratio is convergent and that resulting asset allocations perform better than Bogle's rule, a famous asset allocation rule [7]. Furthermore, we showed that dynamic allocations perform better than deterministic allocations, the current standard in life-cycle investing.

7.2. Recommendations for further research

Our research has focused on using dynamic mean-variance optimization for the replacement ratio. Although we applied dynamic mean-variance optimization to the replacement ratio, our extension of dynamic mean-variance optimization for involved models of asset returns is more generally applicable: other metrics that are a continuous function of wealth can be mean-variance optimized using our algorithm.

We believe that the methods in sections 4.5 to 5.2 can be improved further with parallel computing. The most straightforward improvement is using more samples in the algorithm. By using more samples, machine-learning methods like support vector regression could improve on the estimation of conditional expectation.

Furthermore, we found volatility of the asset allocation can be caused by fluctuations in the conditional expectation of asset returns. When there is a suspicion that fluctuations in conditional expectations are caused by estimation methods or noise, more samples should be used to obtain more stable estimations. Increasing the number of samples requires a more efficient implementation of the algorithm to keep computation time down, which is why we recommend research into parallelization.

A final recommendation is further research into risk-aversion. In Problem 4.4, there is no explicit mention of risk-aversion. Implicitly, risk-aversion is determined by choosing the target, a high target implying low risk-aversion and a low target implying high risk-aversion. Despite this implicit link, choosing a low target does not guarantee an asset allocation that always corresponds to a risk-averse investor profile, as risky equity allocation can be high at the start of the investment period. Therefore, further research into incorporating risk-aversion into the optimization, e.g., constraining the asset allocation so that all demands of the investor are satisfied, is recommended.

7.3. Recommendations for Ortec Finance

Section 6.3 shows that the constructed dynamic life-cycles have several advantages over deterministic life-cycles. Using a dynamic life-cycle increases the probability of hitting the target for the replacement ration and decreases the variance around this target. We therefore advise to further develop the algorithm for optimization of the replacement ratio described in this thesis.

In further development of the algorithm, the following issues have to be solved. The first is volatility of the asset allocation. Volatility of the asset allocation is partly caused by fluctuation of the conditional expectation of asset returns. When an investor is convinced that these fluctuations in conditional expectation are actual predictions of future asset returns, then they should be taken into account and volatility of the asset allocation is a logical consequence. A decision on whether it is desirable to take into account predictability of asset returns into life-cycle design should be made.

Another issue is the risk-aversion of the investor. A risk-averse investor should choose a small target, so that less risky assets are sufficient to reach this target. Despite choosing a small target, however, when asset returns are lower than expected, the resulting asset allocation might be too risky for the investors appetite. As recommended in the previous section, research into constraining the asset allocation so that all demands of the investor are satisfied is necessary.

Finally, investors are worried about the downside of their investments. The dynamic mean-variance optimization is a trade-off between expectation and variance, variance to the up- and downside. We therefore advise Ortec Finance to do research into optimization methods that only measure the downside of investments, e.g., mean-CVaR optimization.

Bibliography

- [1] I. Bajeux-Besnainou and R. Portait. Dynamic asset allocation in a mean-variance framework. *Management Science*, 44(11):S79–S95, November 1998.
- [2] D. P. Baron. On the utility theoretical foundations of mean-variance analysis. *The Journal of Finance*, 32(5):1683–1697, December 1997.
- [3] J. Barraquand and D. Martineau. Numerical valuation of high dimensional multivariate american securities. *The Journal of Financial and Quantitative Analysis*, 30(3):383–405, September 1995.
- [4] S. Basak and G. Chabakauri. Dynamic mean-variance asset allocation. *The Review of Financial Studies*, 23(8):2970–3016, August 2010.
- [5] T. R. Bielecki, H. Jin, S. R. Pliska, and X. Y. Zhou. Continuous-time mean-variance portfolio selection with bankruptcy prohibition. *Mathematical Finance*, 15(2):213–244, April 2005.
- [6] Z. Bodie, R. C. Merton, and W. F. Samuelson. Labor supply flexibility and portfolio choice in a life cycle model. *Journal of Economic Dynamics and Control*, 16(3):427–449, 1992.
- [7] J. C. Bogle. *Bogle on Mutual Funds: New Perspectives for the Intelligent Investor*. McGraw-Hill, 1993.
- [8] T. Bollerslev, R. F. Engle, and J. M. Wooldridge. A capital asset pricing model with time-varying covariances. *Journal of Political Economy*, 96:116–131, 1988.
- [9] M. W. Brandt, A. Goyal, P. Santa-Clara, and J. R. Stroud. A simulation approach to dynamic portfolio choice with an application to learning about return predictability. *The Review of Financial Studies*, 18(3):831–873, 2005.
- [10] M.J. Brennan, E. S. Schwartz, and R. Lagnado. Strategic asset allocation. *Journal of Economic Dynamics and Control*, 21:1377–1403, 1997.
- [11] F. Cong and C. W. Oosterlee. Accurate and robust numerical methods for the dynamic portfolio management problem. *Computational Economics*, pages 1–26, 2016.
- [12] F. Cong and C. W. Oosterlee. Multi-period mean-variance portfolio optimization based on Monte-Carlo simulation. *Journal of Economic Dynamics and Control*, 64:32–28, 2016.
- [13] F. Cong and C. W. Oosterlee. On pre-commitment aspects of a time-consistent strategy for a mean-variance investor. *Journal of Economic Dynamics and Control*, 70:178 – 193, 2016.
- [14] C. Cortes and V. Vapnik. Support-vector networks. *Machine Learning*, 20:273–297, 1995.
- [15] L. N. Coyle and J. J. Yang. Analysis of the ssap method for the numerical valuation of high-dimensional multivariate american securities. *Algorithmica*, 25(1):75–98, 1999.
- [16] P. de Goeij and W. Marquering. Modeling the conditional covariance between stock and bond returns: A multivariate garch approach. *Journal of Financial Econometrics*, 2(4):531–564, 2004.
- [17] B. Efron, T. Hastie, I. Johnstone, and R. Tibshirani. Least angle regression. *The Annals of Statistics*, 32(2): 407–499, 2004.
- [18] J. Estrada. The glidepath illusion: An international perspective. *The Journal of Portfolio Management*, 40(4):54–62, 2014.
- [19] European Central Bank. Recent developments in government bond yield spreads. https://www.ecb.europa.eu/pub/pdf/other/mb200805_focus03.en.pdf. Accessed: 21-04-2017.

- [20] P. A. Forsyth, Y. Li, and K. R. Vetzal. Are target date funds dinosaurs? Failure to adapt can lead to extinction. *Global Risk Institute [Internet]*, October 2016.
- [21] M. Friedman and L. J. Savage. The expected-utility hypothesis and the measurability of utility. *Journal of Political Economy*, 60(6):463–474, December 1952.
- [22] L. Garlappi and G. Skoulakis. Numerical solutions to dynamic portfolio problems: The case for value function iteration using taylor approximation. *Computational Economics*, 33(2):193–207, 2009.
- [23] P. Glasserman and B. Yu. *Simulation for American Options: Regression Now or Regression Later?*, pages 213–226. Springer Berlin Heidelberg, 2004.
- [24] W. H. Greene. *Econometric Analysis*. Prentice Hall, 1993.
- [25] J. C. Hull. *Options, futures, and other derivatives*. Prentice Hall, 1988.
- [26] B. Kramer. Projection methods: Simulating economic scenarios. <http://www.ortec-finance.nl/~media/Files/Research/Working%20Papers/Applied/Projection%20methods%20Simulating%20economic%20scenarios.ashx>, March 2014. Accessed: 28-04-2017.
- [27] D. Li and W. Ng. Optimal dynamic portfolio selection: Multiperiod mean-variance formulation. *Mathematical Finance*, 10(3):387–406, July 2000.
- [28] D. Li, X. Y. Zhou, and A. E. B. Lim. Dynamic mean-variance portfolio selection with no-shorting constraints. *SIAM Journal on Control and Optimization*, 40(5):1540–1555, 2002.
- [29] A. E. B. Lim and X. Y. Zhou. Mean-variance portfolio selection with random parameters in a complete market. *Mathematics of Operations Research*, 27(1):101–120, February 2002.
- [30] F. A. Longstaff and E. S. Schwartz. Valuing american options by simulation: A simple least-squares approach. *The Review of Financial Studies*, 14(1):113–147, 2001.
- [31] L. C. MacLean, Y. Zhao, and W. T. Ziemba. Mean-variance versus expected utility in dynamic investment analysis. *Computational Management Science*, 8(1):3–22, 2011.
- [32] H. Markowitz. Portfolio selection. *The Journal of Finance*, 7(1):77–91, March 1952.
- [33] R. C. Merton. Lifetime portfolio selection under uncertainty: The continuous-time case. *The Review of Economics and Statistics*, 51(3):247–257, 1969.
- [34] J. von Neumann and O. Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, Princeton, NJ, 1953.
- [35] OECD. Glossary of statistical terms. <https://stats.oecd.org/glossary/detail.asp?ID=5900>. Accessed: 01-02-2017.
- [36] J. Park and I. W. Sandberg. Universal approximation using radial-basis-function networks. *Neural Computation*, 3(2):246–257, 1991.
- [37] R. Rebonato and P. Jäckel. The most general methodology to create a valid correlation matrix for risk management and option pricing purposes. *Journal of Risk*, 2(2):17–27, 1999.
- [38] P. A. Samuelson. Lifetime portfolio selection by dynamic stochastic programming. *The Review of Economics and Statistics*, 51(3):239–246, August 1969.
- [39] R. J. Shiller. Life-cycle personal accounts proposal for social security: An evaluation of president bush's proposal. *Journal of Policy Modeling*, 28(4):427–444, May 2006.
- [40] A. J. Smola and B. Schölkopf. A tutorial on support vector regression. *Statistics and Computing*, 14: 199–222, 2004.
- [41] TIAA-CREF Asset Management. Tiaa-cref lifecycle funds: Methodology and design, 2014.

- [42] R. Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society, Series B (Methodological)*, 58(1):267–288, 1996.
- [43] J. H. van Binsbergen and M. W. Brandt. Solving dynamic portfolio choice problems by recursing on optimized portfolio weights or on the value function? *Computational Economics*, 29(3):355–367, 2007.
- [44] Vanguard Research. Vanguard life-cycle investing model: A framework for building target-date portfolios. <https://personal.vanguard.com/pdf/ISGLCIM.pdf>, September 2016. Accessed: 12-01-2017.
- [45] L. M. Viceira. Optimal portfolio choice for long-horizon investors with nontradable labor income. *The Journal of Finance*, 56(2):433–470, April 2001.
- [46] L. M. Viceira. Life-cycle funds. Available at SSRN: <https://ssrn.com/abstract=988362> or <http://dx.doi.org/10.2139/ssrn.988362>, May 2007.
- [47] J. Wang and P. A. Forsyth. Numerical solution of the hamilton-jacobi-bellman formulation for continuous time mean variance asset allocation. *Journal of Economic Dynamics and Control*, 34(2):207–230, February 2010.
- [48] J. Wang and P. A. Forsyth. Continuous time mean variance asset allocation: A time-consistent strategy. *European Journal of Operational Research*, 209(2):184–201, March 2011.
- [49] D. Williams. *Probability with Martingales*. Cambridge University Press, 1991.
- [50] X. Y. Zhou and D. Li. Continuous-time mean-variance portfolio selection: A stochastic LQ framework. *Applied Mathematics and Optimization*, 42(1):19–33, 2000.