

Dynamic life cycle investing

by

T.R.B. den Haan

to obtain the degree of Master of Science
at the Delft University of Technology,
to be defended publicly on Wednesday June 12, 2019 at 3:00 PM.

Student number: 4141237
Project duration: September 3, 2018 – June 12, 2019
Thesis committee: Prof. dr. ir. C.W. Oosterlee, TU Delft, supervisor
Dr. M. van der Schans, Ortec Finance, supervisor
Dr. J. van der Woude, TU Delft
K.W. Chau MPhil, TU Delft

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

Abstract

This thesis studies the asset allocation of a DC pension investor over a long time horizon. Investors allocate their portfolio wealth between two assets: a return portfolio and a matching portfolio. Investors can adjust their allocation once a year. Several dynamic investment strategies that improve investment results compared to fixed allocations or static life cycles are shown. The dynamic investment strategies have been constructed by using two different approaches. The first approach is rule-based and defines intermediate wealth targets for every year in the investment horizon. Investment decisions are taken based on performance compared to these targets. The second approach involves a dynamic programming algorithm. The asset allocation over time is not always stable when using dynamic programming. Methods to smooth the asset allocation over time and improve stability are discussed. Last, both approaches are combined in one strategy.

Preface

This thesis has been submitted for the degree of Master of Science in Applied Mathematics at Delft University of Technology. The academic supervisor of this research is Kees Oosterlee, professor at the Numerical Analysis group of Delft Institute of Applied Mathematics. The thesis project was carried out in collaboration with Ortec Finance under the daily supervision of Martin van der Schans.

I would like to thank Martin van der Schans, Kees Oosterlee and Ki Wai Chau for their valuable input and feedback during this project. Furthermore I would like to thank Jacob van der Woude for being part of the thesis committee.

Tim den Haan
Rotterdam, June 2019

Contents

1	Introduction	1
1.1	Goal	1
1.2	Setting.	2
1.3	Problem definition	3
1.4	Literature overview	3
1.4.1	Merton portfolio theory	3
1.4.2	Time consistent Mean Variance	4
1.4.3	Target range strategies.	5
1.5	Outline	5
2	Overview	7
2.1	Simulation.	7
2.1.1	Admissible control	9
2.2	Rule-based strategies	9
2.2.1	Bogle's rule	10
2.3	Dynamic programming strategy	11
2.3.1	Least squares Monte-Carlo method	11
2.3.2	Regression techniques.	12
2.A	Ortec Finance modelling approach	13
2.B	Salary table	15
2.C	Academic model	16
3	Rule based strategies	17
3.1	Cumulative target strategy	17
3.1.1	Wealth target construction	17
3.1.2	Target replacement ratio	19
3.1.3	Results	20
3.2	Individual target.	22
3.2.1	Results	23
3.3	Constant cash flow	24
3.3.1	Results	24
3.3.2	Academic model results	26
3.4	Conclusion	28
3.A	Strategy results.	29
4	Dynamic programming strategy	31
4.1	Algorithm	31
4.2	Fixed wealth target	32
4.2.1	Results	32
4.2.2	Convergence of the allocation decision	34
4.3	Stochastic wealth target	36
4.3.1	Results	36
4.4	Conclusion	38
4.A	Strategy results.	39
5	Control smoothing	41
5.1	Fictive transaction costs	41
5.1.1	Results	42
5.2	Bandwidth limitation	43
5.2.1	Results	44

5.3	Minimal improvement.	44
5.3.1	Results	45
5.4	Conclusion	48
5.A	Strategy results.	49
6	Combination strategy	51
6.1	Combination strategy.	51
6.2	Results	53
6.3	Academic model	53
6.4	Conclusion	55
7	Conclusion	57
7.1	Conclusion	57
7.2	Recommendations	58
	Bibliography	59

Introduction

This thesis studies pension investments over a long time horizon. Many people save for their pension income during their working career. The pension schemes they use can be divided into two categories: a DB scheme and a DC scheme. Most people in the Netherlands follow a DB (defined benefit) scheme. Such a pension plan is known as a collective pension plan because the pension fund takes on the investment risk. Values of pension incomes for participants are dependent on factors such as age and tenure of service. The pension fund may raise or lower contribution requirements for participants in a DB scheme.

The other well-known pension scheme is a DC (defined contribution) scheme. The investment risk lies with participants and pension incomes are dependent on investment returns. Contributions to the DC pension scheme are fixed, age-dependent percentages of salaries of participants. In the Netherlands, 15% of the pension entitlements were DC in 2015 and this share is increasing [14].

DC pension investors distribute their pension savings among various assets via funds offered by pension managers. Currently, many of such investment strategies are only dependent on time remaining until retirement [1] [5] [15]. The share of safer assets such as bonds increases as retirement gets closer, to avoid the risk of losing a considerable amount of savings when close to retirement.

Intuitively, it makes sense to take the current portfolio value into account when making an investment decision. When the current value of the portfolio is ignored, the investor may take more risk than required. Improvements can be expected when making use of more information in the decision making process. Investors are able to steer towards a target by using such information.

Investment strategies that make use of the current portfolio state are known as dynamic investment strategies. This thesis shows dynamic investment strategies for individual DC pension investors who manage their own pension savings.

1.1. Goal

This thesis can be split in two separate parts, each with their own goal. The goal of the first part is to construct dynamic investment strategies that outperform those used in practice. For this, two approaches are used: rule-based strategies and a dynamic programming strategy. For a strategy to outperform another, it is not enough to have simply higher average outcomes or more upward potential; it is assumed that pension investors prefer certainty above all else. Their main objective is not to gain as much wealth as possible, it is to reach a (realistic) pre-defined target with minimal risk.

Optimal strategies found using these approaches can be difficult to apply in practice. For example, the annual turnover is often too high. It is not practical to sell all bonds in the portfolio such that all wealth can be invested in stocks, only to switch completely towards bonds the year after. The goal of the second part of this thesis is therefore to smooth optimal solutions such that they can be used in real life.

1.2. Setting

This thesis considers long-term investors who put their retirement savings in a DC pension scheme. The investment period is modelled in discrete time: The investor's age is assumed to be 26 at time t_0 the beginning of the investment horizon and the investor is assumed to retire at the age of 67 at time $t_n = T$, the end of the investment horizon.

The investor manages a portfolio during the investment period. The portfolio consists of two assets: an equity-like asset and a bond-like asset. The equity-like asset has a higher expected return and higher volatility compared to the bond-like asset. Both assets carry risk: there is no risk-free rate in the modelling setting.

Investors manage their portfolio by adjusting the balance between these two assets. The asset weight of portfolio wealth W for the equity-like asset at time $t_i, i \in \{0, \dots, n\}$ is denoted by α_{t_i} . The asset weight for the bond-like asset at time t_i is equal to $1 - \alpha_{t_i}$. The investor is not allowed to short-sell assets or borrow money: $0 \leq \alpha_{t_i} \leq 1$.

No assumptions are made regarding the dynamics of the assets. Instead, real-world scenarios from Ortec Finance are used. Ortec Finance forecasts the economy, builds econometric models and advises on optimal decision making.

A risk-free rate is not available in the framework. Designed strategies do not rely on asset dynamics and can still be used in case the underlying asset return model is modified.

A more detailed description of the investment setting can be found in Section 2.1.

At terminal time T the portfolio is liquidated and an annual pension income for 20 years is bought. The value of the annual pension income is expressed in terms of the *replacement ratio*: the ratio between the pension income and the average salary of the investor.

Definition 1.1 (Replacement ratio). Let I_{avg} be the time-average labour income of the investor,

$$I_{\text{avg}} := \frac{1}{n} \sum_{i=0}^{n-1} \tilde{q}_{t_i},$$

where \tilde{q}_{t_i} is the salary of the investor earned in year t_i , corrected for inflation from the period $[t_i, T)$. I_{avg} contains a correction for inflation such that salaries for different years can be weighted equally.

Let P be the annual pension income bought at the end of the investment horizon. The replacement ratio RR at terminal time T is defined as

$$RR_T := \frac{P}{I_{\text{avg}}}$$

The investor compares investment results by using a *utility function*. This thesis uses the following definition of a utility function:

Definition 1.2 (Utility function). A function $f : D \rightarrow \mathbb{R}$ is called a utility function if it is concave and continuous on D .

The risk-averse nature of the investor is reflected in the concave property of the utility function definition. An example of a utility is the shortage:

Definition 1.3 (Shortage). Let S be a representative set of outcomes for an investment strategy. Let z^* be the target value of the investor. Let $|S|$ denote the number of elements in set S . The shortage μ of a set is equal to

$$\mu(S) = \frac{1}{|S|} \sum_{s \in S} \min[0, s - z^*],$$

Investors use the shortage measure to indicate the target-reachability of an investment strategy. Investment risk is compared by using the *conditional value at risk* measure:

Definition 1.4 (Conditional Value at Risk). Let X be a random variable accounting for the outcomes of an investment strategy. Let $0 < \gamma < 1$. The Conditional Value at Risk (CVaR) at confidence level γ is the expected value of the worst γ cases:

$$\text{CVaR}_\gamma = -\frac{1}{\gamma} \int_0^\gamma \text{VaR}_\eta(X) d\eta,$$

where VaR_η is the Value at Risk at confidence level η ,

$$\text{VaR}_\gamma = \inf [x \in \mathbb{R} : P(X \leq x) \geq \gamma]$$

Dynamic investment strategies use the state of the portfolio in the decision making process. The portfolio state at time t is denoted by Z_t . The portfolio state contains information about the investment progress. Examples of portfolio states are portfolio wealth W_t , expected terminal wealth $\mathbb{E}[W_T | \mathcal{F}_t]$ or the expected replacement ratio.

1.3. Problem definition

The objective of the investor is given in Problem 1.5.

Problem 1.5 (Optimal asset allocation). Let \mathcal{S} be a dynamic investment strategy that assigns the asset allocation α_t at time t depending on the portfolio state Z_t :

$$\begin{aligned} \mathcal{S} &: [t_0, T] \times D \rightarrow [0, 1], \\ \alpha_t &= \mathcal{S}(t, Z_t), \end{aligned}$$

where D is the domain of state variable Z . Let S be a representative set of outcomes of strategy \mathcal{S} . Let Z_T be the final state of the portfolio after executing strategy \mathcal{S} . Let γ be the confidence level of the conditional value at risk. The investor's objective is to optimize over a utility function U , e.g. the shortage, under the constraint of not increasing the investment risk:

$$\begin{aligned} &\max_S \mathbb{E}[U(Z_T)] \\ \text{s.t. } &\text{CVaR}_\gamma(S) \geq \text{CVaR}_\gamma(\tilde{S}), \end{aligned}$$

where \tilde{S} is a representative set of outcomes for a reference strategy such as an investment strategy used in practice.

Other constraints regarding investment risk can also be applied, such as the volatility of the expected pension income over time.

1.4. Literature overview

The problem setting described in Section 1.2 has been investigated using different approaches. The most relevant results for this thesis are summarized briefly below. The notation in this section differs from the rest of this thesis.

1.4.1. Merton portfolio theory

An asset allocation problem similar to the setting described in Section 1.2 was solved analytically in 1969 by Merton [13]. In Merton's setting, investors must allocate their assets between a risky and a risk-free asset. Investors also have the possibility to consume wealth during the investment period. The objective is to maximize

$$\max \mathbb{E} \left[\int_0^T e^{-\rho t} U(C_t) dt + \varepsilon^{1-\gamma} e^{-\rho T} U(W_T) \right],$$

where ρ is a subjective discount rate. U is a concave utility function and assumed to be of the form yielding constant relative risk aversion (CRRA),

$$-\frac{U''(x)}{U'(x)} = \beta,$$

where β is a constant. Wealth consumption at time t is denoted by $C_t \geq 0$ and $0 < \varepsilon \ll 1$. $\gamma < 1$ is a risk aversion parameter; the higher γ , the more risk the investor is willing to take. Wealth W_t is assumed to follow

$$dW_t = [(\alpha_t(\mu - r) + r)W_t - C_t]dt + \sigma\alpha_t W_t dZ_t,$$

where dZ_t is a Wiener increment. α_t is the wealth share invested in the risky asset, the remainder is invested in the risk-free asset with return r . The risky asset has an expected return of μ and a standard deviation σ .

A closed form solution exists. In this setting, the optimal allocation α_t^* is equal to

$$\alpha_t^* = \frac{\mu - r}{\sigma^2(1 - \gamma)}.$$

α_t^* is not dependent on W_t or t . The optimal allocation is constant.

A risk free rate is available in Merton's setting, while it does not exist in the setting described in Section 1.2. The bond-like asset in this thesis is not 'safe'- there is still some risk involved. Another difference is the utility function used in Chapter 4: it's not of the form of constant relative risk aversion.

1.4.2. Time consistent Mean Variance

In [11], Forsyth and Vetzal use a PDE solver to solve a time-consistent mean-variance problem. Mean-variance analysis optimizes the expected return (mean) for a given level of risk (variance) at every time t . Time-consistency is obtained by assuming that a similar mean-variance problem will be solved at all future times.

Forsyth and Vetzal also have a portfolio consisting of two assets, bonds B and stocks S . They assume the assets can be described by the following stochastic processes:

$$\begin{aligned} dS_t &= (\mu - \lambda \mathbb{E}[\xi - 1])S_t dt + \sigma S_t dZ + (\xi - 1)S_t dQ \\ dB_t &= rB_t dt, \end{aligned}$$

where μ is the drift rate, λ a positive intensity parameter, dZ a Wiener increment and ξ a random variable acting as a jump multiplier. dQ can be 0 or 1 and equals 1 with probability λdt . r is the risk-free rate. Bond asset B is risk-free. Value function V is defined as

$$V(z, t_i) = \sup_{\{\alpha_i, \dots, \alpha_n\}} \left\{ \mathbb{E}[W_T] - \rho(W_{t_i}) \text{Var}[W_T] \mid W_{t_i} = z \right\} \text{ s.t. } \{\alpha_i, \dots, \alpha_n\} = \{\alpha_i, \alpha_{i+1}^*, \dots, \alpha_n^*\},$$

where α_i^* is the optimal allocation at time t_i and $\rho > 0$ a scalarization parameter. W_{t_i} is the portfolio wealth at time t_i .

Expected quadratic shortfall

In the same paper, Forsyth and Vetzal define a different problem that uses a fixed terminal wealth target W^* . The corresponding value function V equals

$$V(z, t_i) = \inf_{\{\alpha_i, \dots, \alpha_n\}} \left\{ \mathbb{E} \left[\min(W_{t_i} - W^*, 0)^2 \mid Z_{t_i} = z \right] \right\}.$$

The introduction of a wealth target allows the investor to identify a surplus: wealth up to the target may be invested in stocks, any remainder is invested in the risk-free rate.

1.4.3. Target range strategies

In Zhang (2017) [23], a similar problem is solved using dynamic programming in combination with a least squares Monte-Carlo method. The objective function is defined by

$$\sup_{\{\alpha_i, \dots, \alpha_n\}} \mathbb{E} [f(W_T) | \mathcal{F}_{t_i}],$$

with

$$f(w) = (w - L_W) \mathbb{1}\{L_W \leq w \leq U_W\},$$

where L_W and U_W are lower and upper bounds for the final wealth W_T .

Zhang e.a. assume that upward potential is intertwined with downside risk. A difference compared to the setting of this thesis is the existence of a risk-free asset. Zhang allows the investor to invest wealth in a risk-free asset at any time. Final wealth W_T is steered towards a target range by investing the difference between the risk-free-discounted upper bound and wealth W_t in the risk-free asset.

1.5. Outline

Additional information on the simulation and underlying theoretical concepts of investment strategies can be found in Chapter 2. Several investment strategies are discussed in Chapters 3 to 6.

To reduce risk exposure, both Forsyth and Zhang suggest to invest excess wealth in a risk-free asset. Chapter 3 shows a rule-based investment strategy in which unnecessary risk is reduced in a different way: wealth is secured in the matching portfolio. The problem setting studied in this thesis does not have a risk-free rate. The matching portfolio is the most secure asset available in the investment setting. Results of Chapter 3 can be seen as a benchmark for subsequent chapters.

This thesis uses a utility function in combination with dynamic programming in Chapter 4. Just like [23], dynamic programming in combination with the least squares Monte-Carlo method is used.

Investment decisions of dynamic programming strategies not always stable over the time horizon: the annual turnover is large. Chapter 5 shows methods to reduce the annual turnover at the cost of optimality. Chapter 6 combines results from the rule-based and dynamic programming approach.

2

Overview

As noted in Section 1.1, two types of investment strategies are studied in this thesis. This chapter offers the formal of problem definitions for both the rule-based approach and the dynamic programming approach. Details about the investment setting illustrated in Section 1.2 and relative theoretical background are also provided.

2.1. Simulation

The investment period is simulated in the following setting. Investors control their investments by managing a portfolio. Two assets are available in the market: a risky equity-like asset called 'return portfolio' and a more secure bond-like asset called 'matching portfolio'. The return portfolio has a higher expected return and higher volatility compared to the matching portfolio. The portfolio is managed over time: The investment period starts at time t_0 and at terminal time T the portfolio is liquidated. The amount of wealth invested in the return portfolio at time t is denoted by S_t . B_t indicates the matching portfolio wealth at time t . The total portfolio wealth W of the portfolio at time t thus equals:

$$W_t = S_t + B_t.$$

The portfolio wealth can be rebalanced annually. Just before rebalancing, a contribution c_t to the portfolio is made. c_t can also be referred to as the premium paid at time t . An age-dependent percentage p_t of the investor's salary q_t over the previous time period is invested in the two assets at time t . Salary q_t follows a so-called career path: it is assumed that the annual salary of the investor increases over the investment period. Premium percentage p_t is assumed to increase every 5 years.

The investor is not allowed to short-sell assets or borrow money. Therefore, the return portfolio share α_{t_i} at time t_i ,

$$\alpha_{t_i} = \frac{S_{t_i}}{W_{t_i}},$$

is constrained to $0 \leq \alpha_{t_i} \leq 1$. The matching portfolio share is equal to $1 - \alpha_{t_i}$. α_{t_i} will also be referred to as the asset allocation at time t_i . The subscript t is often dropped to avoid unclear notation such as $\alpha_{t_{i-1}}$. The notation α_{i-1} is used instead.

This thesis uses a density forecast generated by Ortec Finance, also known as an Ortec Finance scenario set. The modelling approach for such a scenario set is explained briefly in Section 2.A. The density forecast is obtained by Monte Carlo simulation and consists of 2000 sample paths that contain the following variables:

- Return portfolio return r_{t_i} is the return over the return portfolio in the period $[t_{i-1}, t_i)$.
- Matching portfolio return r_{m,t_i} is the return over the matching portfolio in the period $[t_{i-1}, t_i)$.

- Annual inflation ζ_{t_i} is the inflation over the period $[t_{i-1}, t_i)$.
- Annual wage inflation w_{t_i} is the wage inflation over the period $[t_{i-1}, t_i)$.
- Market value factor M_{t_i} is an interest rate variable determining the price at time t_i of the future pension income.

Annual salary q_{t_i} , the salary earned by the investor over the period $[t_{i-1}, t_i)$, is randomized by multiplying the career path found in Table 2.1 with the wage inflation path for each scenario.

The matching portfolio and market value factors are tailored to the investment horizon and the initial age of the investor. The market value factor is used to construct the price of a bond with 20 annual coupons, starting at the end of the investment horizon. Such bonds can be used to secure pension income during the investment period. The matching portfolio replicates the relative change of this bond price over time. Small deviations are caused by inflation.

Academic model

Results for Chapters 3 and 6 are also computed for an academic model using standard time series. This model gives a basic insight in the dynamics of the scenarios from the Ortec Finance scenario set. Parameters have been estimated on the Ortec Finance scenario set. The modelling approach of Ortec Finance is briefly explained in Section 2.A.

Return portfolio return at rebalancing times $t_i \in \mathcal{T}$ is assumed to be distributed normally and independent of past returns. The expected return for the return portfolio is 6% and the standard deviation is assumed to be 0.16:

$$r_{t_i} \sim N(0.06, 0.16^2) .$$

Inflation at time t_i , ζ_i , is assumed to be dependent on the inflation at t_{i-1} and the return portfolio return at time t_i , r_i :

$$\zeta_i = a_\zeta \zeta_{i-1} + b_\zeta r_i + \varepsilon_\zeta .$$

Wage inflation w_i follows similar dynamics:

$$w_i = a_w w_{i-1} + b_w r_i + \varepsilon_w .$$

The matching portfolio return and the relative change in market value factor are correlated strongly. The relative change in market value factor, r_{MVF} , is simulated first:

$$r_{MVF,i} = a_{MVF} r_{MVF,i-1} + b_{MVF} r_i + \varepsilon_{MVF,i} .$$

The market value factor at time t_i is equal to

$$M_i = M_{i-1} (1 + r_{MVF,i}) . \quad (2.1)$$

The market value factor is bound at every t_i between $M_{\min,i}$ and $M_{\max,i}$. If M_i exceeds these bounds, M_i can be updated as follows

$$M_i = \begin{cases} M_{\min,i} + |r_{MVF,i}|, & \text{if } M_i < M_{\min,i} \\ M_{\max,i} - |r_{MVF,i}|, & \text{if } M_i > M_{\max,i} \\ M_i, & \text{otherwise} \end{cases} .$$

The relative change in market value can now be updated accordingly:

$$r_{MVF,i} = \frac{M_i}{M_{i-1}} - 1,$$

such that $r_{MVF,i}$ gives a random change in market value factor such that the market value factor stays between bounds $[M_{\min,i}, M_{\max,i}]$.

The matching portfolio return at time t_i is then equal to

$$r_{m,i} = a_m r_{MVF,i} + b_m r_i + \varepsilon_m$$

Values for each variable can be found in Table 2.2.

2.1.1. Admissible control

The investor is able to adjust the asset allocation at adjusting times \mathcal{T} .

Definition 2.1 (Adjusting times). \mathcal{T} is the set of adjusting times,

$$\mathcal{T} = \{t_0 < t_1 < \dots < t_{n-1} < t_n = T\}.$$

Adjusting times are equidistant. In this thesis, adjusting times are assumed to be the first day of each investment year.

Final time T is included in \mathcal{T} for notational purposes. In practice, the portfolio is liquidated at time T and a pension product is bought. No wealth remains to adjust at time T .

The asset allocation over time serves as the input for the optimal asset allocation problem as defined in Problem 1.5. The investor can only change the output of the investment process by altering the asset allocation over time. α_i over time defines the investment strategy followed by the investor. The investor's input at adjusting times is referred to as an admissible control. The definition for an admissible control for this thesis is considered the following:

Definition 2.2 (Admissible control). Let \mathcal{A} be the sequence of return portfolio weights at adjusting times $t \in \mathcal{T}$,

$$\mathcal{A} := \{\alpha_0, \alpha_1, \dots, \alpha_{n-1}, \alpha_n\}.$$

\mathcal{A} is an admissible control if $0 \leq \alpha_i \leq 1$ for all $t_i \in \mathcal{T}$. The tail of the control starting at t_i is denoted by \mathcal{A}_i ,

$$\mathcal{A}_i := \{\alpha_i, \alpha_{i+1}, \dots, \alpha_{n-1}, \alpha_n\}.$$

α_n is included for simplified notation, but serves no purpose because the portfolio is liquidated at time T .

In practice, investment risk is reduced towards retirement by decreasing α_i over time. Control \mathcal{A} is therefore known as the glide path in the pension industry. A constant control \mathcal{A} will also be referred to as a static mix. A control or investment strategy that is only dependent on time, $\mathcal{A} = \mathcal{A}(t)$, will also be referred to as a static life cycle.

Dynamic investment strategies use the portfolio state at time t to determine the asset allocation at time t . The portfolio is adjusted after the results from the previous time period are known. To illustrate this, times t^- and t^+ are introduced:

Let $t_i^- = t_i - \varepsilon$, with $0 < \varepsilon \ll 1$, and $t_i^+ = t_i + \varepsilon$. A subscript i for variables from the investment setting, such as the asset allocation α_i , indicates the value of that variable at time t_i^+ , just after adjusting time $t_i \in \mathcal{T}$.

Variables that describe results of the previous period are assumed to be known at time t_i^- . These include the salary q_i , return portfolio return r_i , matching portfolio return $r_{m,i}$ and annual inflation ζ_i . The annual contribution c_i is assumed to be made at time t_i^- .

2.2. Rule-based strategies

Rule-based strategies construct a control \mathcal{A} by determining each α_i on a pre-defined set of rules. Rules can be dependent on time t and state Z . A strategy with a set of rules only dependent on time is known as a static life cycle. The goal is to find a set of rules that have a better performance compared to static life cycle strategies used in practice.

For a rule-based approach, Problem 1.5 is adjusted into Problem 2.3. A rule-based strategy also takes into account target state Z_t^* and the investor optimizes over the shortage measure:

Problem 2.3 (Optimal rule-based asset allocation). Let \mathcal{S} be a dynamic rule-based investment strategy that assigns the asset allocation α_t at time t depending on portfolio state Z_t and target state Z_t^* ,

$$\begin{aligned} \mathcal{S} &: [t_0, T] \times \mathbb{R}_{>0}^2 \rightarrow [0, 1], \\ \alpha_t &= \mathcal{S}(t, Z_t, Z_t^*), \end{aligned}$$

Let S be a representative set of outcomes of strategy \mathcal{S} . Let Z_T be the final state of the portfolio after executing strategy \mathcal{S} . Let γ be the confidence level of the conditional value at risk. The investor's objective is to optimize over the shortage under the constraint of not increasing the investment risk:

$$\begin{aligned} &\max_S \mu(Z_T) \\ \text{s.t.} \quad &\text{CVaR}_\gamma(S) \geq \text{CVaR}_\gamma(\tilde{S}), \end{aligned}$$

where \tilde{S} is a representative set of outcomes for a static life cycle used in practice.

It is assumed that upward potential is intertwined with downside risk. If investors make a decision with a higher expected return, they also increase their risk profile and the probability for losses increases. A second assumption for rule-based strategies in this thesis is that investors are only interested in reaching their goal. They feel indifferent about any two values above their target.

2.2.1. Bogle's rule

A well-known rule-based strategy is known as Bogle's rule [3]. One decision rule defines the strategy: at age τ , the investor allocates $(100-\tau)\%$ of the portfolio towards equity, while the remainder is invested in bonds. The rule states to allocate most of the portfolio wealth to assets with a high expected return at a younger age, but with an increased risk profile. Over the years, more and more wealth is assigned towards safer assets. The underlying idea is that the investor can afford the increased risk during the first years of the strategy, because enough time is left for recovery.

In practice, salaries and contributions to the pension scheme often increase over time. At the beginning of their career, investors earn a modest salary and the contributions to their pension are relatively small. At a later age, investors are able to direct a large share of their (increased) annual income towards their pension. Such large contributions are mainly invested in bonds when Bogle's rule is applied in practice.

An example of Bogle's rule can be found in Figure 2.1.




Figure 2.1: Example of an asset allocation strategy based on Bogle's rule.

Bogle's rule is only dependent on time and is therefore a static life cycle. At time $t = 0$, the glide path is known for all t . The strategy cannot be adjusted during the investment period.

Variants on Bogle's rule are used in practice by pension managers. The share of bonds is increased overtime. A static life cycle strategy by Achmea [1] is shown in Figure 2.2.




Figure 2.2: Example of a life cycle strategy applied in practice.

2.3. Dynamic programming strategy

Dynamic programming was introduced by Richard Bellman [2]. He constructed a mathematical optimization technique that can be applied to a broad range of problems. It can also be applied to the setting of Section 2.1. A complex optimization problem (optimizing the investment strategy) is solved by reducing the problem to smaller scale sub-problems (optimizing the investment decision at time t_i for every $t_i \in \mathcal{T}$). Sub-problems are nested - the optimal investment decision at time t_i is dependent on what will happen at time t_{i+1} . The sub-problem at time t_i is solved under the assumption that future sub-problems at time $t_{i+1}, \dots, t_n = T$ will also be optimized. Sub-problems are therefore solved recursively, backwards in time.

Problem 2.4 defines the sub-problem that is solved to obtain the optimal investment decision at time $t_i \in \mathcal{T}$.

Problem 2.4 (Optimal investment decision at time t_i). Let \mathcal{A} be an admissible control. Let $U : \mathbb{R}_{>0} \rightarrow \mathbb{R}$ be a utility function that measures final state Z_T . The optimal investment decision α_i^* at time t_i is the solution of

$$V(z, t_i) = \sup_{\mathcal{A}_i} \mathbb{E} [U(Z_T) | Z_{t_i} = z]$$

such that $\mathcal{A}_i = \{\alpha_i, \mathcal{A}_{i+1}^*\} := \{\alpha_i, \alpha_{i+1}^*, \dots, \alpha_n^*\}$
 $\alpha_i \in [0, 1]$,

where α_{i+1}^* indicates the optimal investment decision for the optimality problem at t_{i+1} .

2.3.1. Least squares Monte-Carlo method

Now consider the dynamic programming procedure from an algorithmic perspective. Solving the sub-problem at time t_i may alter state Z at t_{i+1} . A different solution may be optimal at time t_{i+1} for a new state $Z_{t_{i+1}}$. Optimizing the investment problem using dynamic programming requires a nested simulation approach. A least squares Monte-Carlo method is used to avoid nested simulations and exponential computational times.

The least squares Monte-Carlo method was introduced by Longstaff & Schwartz in 2001 [12]. They constructed a simple method for pricing American options by simulation. The conditional expectation of the pay-off under the assumption of not exercising the option is estimated by using cross-sectional information already available in the simulation. Realized pay-offs from continuation (or, in the pension investment setting, final utility $U(Z_T)$) are regressed on functions of state variables. The fitted value provides an estimate of the conditional expectation.

2.3.2. Regression techniques

Several regression techniques are used in this thesis. They are defined in this section.

Definition 2.5 (Regress now). Regress now estimates $\mathbb{E} [Y_{t_{i+1}} | X_{t_i}]$, $X_{t_i} \in \mathcal{F}_{t_i}$ by using a set of basis functions Φ with index set \mathcal{J} :

$$Y_{t_{i+1}} \approx \sum_{j \in \mathcal{J}} c_j \varphi_j(X_{t_i}),$$

with c_j coefficients found by using least squares regression and $\varphi_j \in \Phi$. Substituting gives

$$\begin{aligned} \mathbb{E} [Y_{t_{i+1}} | X_{t_i}] &\approx \mathbb{E} \left[\sum_{j \in \mathcal{J}} c_j \varphi_j(X_{t_i}) \middle| X_{t_i} \right] \\ &= \sum_{j \in \mathcal{J}} c_j \varphi_j(X_{t_i}). \end{aligned}$$

Definition 2.6 (Local regression). Local regression was introduced by Cleveland in 1979 [6]. Local regression, also known as LOESS, estimates a linear or quadratic polynomial fit at x by using weighted least squares regression. Weights for an observation (x_i, y_i) are dependent on the distance between x_i and x [7]. The smoothness of the fit is dependent on the percentage of observations taken into account when evaluating at x .

Let n be the amount of observations and let $0 < a \leq 1$ be the neighbourhood parameter: the share of observations used for the weighted least squares regression at the evaluation point. Let $k = an$ rounded up to an integer value. Let $\Delta_i(x)$ be the Euclidean distance of x to x_i . Let $\Delta_{(i)}(x)$ be the values of these distances ordered from smallest to largest.

The weight w_i for an observation (x_i, y_i) is then equal to

$$w_i(x) = T \left(\frac{\Delta_i(x)}{\Delta_{(k)}(x)} \right),$$

with

$$T(u) = \begin{cases} (1 - u^3)^3, & \text{for } 0 \leq u < 1, \\ 0, & \text{for } u \geq 1 \end{cases},$$

also known as the tricube weight function.

2.A. Ortec Finance modelling approach

This section briefly describes the proprietary model developed by Ortec Finance. Results of Chapters 3, 4, 5 and 6 have been computed using scenarios generated by this model. This section is a summary of [20] and [10], more detail can be found in these documents.

Ortec Finance uses time series models combined with Monte Carlo simulation to create sample paths (scenarios) for several hundreds of economic and financial variables such as inflation or asset returns.

The model is set up using the following principles:

1. Realistic scenarios based on robust stylized facts,
2. Out-of-sample testing of risk and return,
3. Views and expert opinion.

Simple random walk models can give good out-of-sample performance in terms of expectations, but such models do not always capture important stylized facts. The following stylized facts are incorporated in the model:

- Term structure of risk and return: risk and return vary with the investment horizon. For example: volatilities increase with the horizon, but they do not always follow the square root pattern that is implied by a simple random walk model. Some variables may show lower volatility due to mean reversion, other variables may show higher volatility due to trends. The sample paths should show an appropriate auto-correlation pattern incorporating this stylized fact, see also [4].
- Business cycles: medium term fluctuations are inherent to economies and markets. Business cycles span a time period of 1 to 8 years. For example: stock prices tend to go down before GDP goes down. Another example could be that when GDP is recovering, unemployment is still be increasing for some time.
- Time varying and state dependent volatility: are relevant for generating realistic short-term scenarios that are used for risk management purposes. For example: low volatility in years preceding the financial crisis. Another example is that volatility of interest rates is lower when interest rates themselves are low.
- Tail risk: correlations between asset returns can increase in bad economic and financial market conditions. Benefits of diversification of assets are reduced in such cases.
- Non-normal distributions: normal distributions are not always realistic. For example: return distributions on short horizons can have fat tails. Many examples can be found in literature, for example [16].
- Yield curves: it is more complicated to generate scenarios for variables that have a maturity dimension such as bond yields. For example, yield curves move over time in specific ways, including parallel shifts, tilts and flex movements, see also [9].

In order to incorporate these features, the model stacks several modelling approaches. The core of the model consists of three factor models that are used to construct scenarios:

- Trend model: A long term trend model drives long term (decades) returns,
- Business cycle model: A medium term business cycle model drives medium term (annual) returns,
- Monthly model: A short term monthly model drives short term (monthly) return.

The underlying data for each model is constructed by using a Fourier-type filter combined with a principal component analysis applied to long term historical data. Next, a Principal Components Analysis is used to estimate common factors for each frequency component. The next step is to estimate a Vector Autoregressive Model of order 1 (VAR(1)) for each common factor. Relevant economic and financial variables are regressed on the factors. [18].

This way, the first two stylized facts have been covered. Additional modelling is required incorporate the other stylized facts in a consistent way. The modelling approaches often deviate from standard academic models, due to the required consistency between variables. For example, Dutch and German interest rates should not show completely different behaviour. Additional information can be found in [10].

In the second step of the model the scenarios are calibrated by using out-of-sample testing. In practice, out-of-sample tests can only be performed by constructing scenarios and then let time pass to collect new data that can be used to perform the tests. More immediate testing can be done by pseudo-out-of-sample testing: consecutive model parameters are estimated without forward looking data. The downside of this approach is that it is unavoidable that the model also uses information of the full data sample.

In the third step information for which historical data is scarce is added to the scenarios. Experts from different fields are consulted to improve a small part of variables of the generated scenarios. Scenarios are also calibrated to economic views in this step. For example, the impact of Brexit on scenarios for the future is not represented in historical data.

2.B. Salary table

The salary path and annual contribution percentages are shown in Table 2.1 [21]. At age x , the corresponding premium percentage at age x of the salary earned at age $x - 1$ is contributed to the portfolio. Note that under Dutch law, not all of the salary can be used to contribute to the pension before taxes. No premium is contributed over the first part of the salary: this part of the salary is called the franchise. The franchise value for this thesis is assumed to be 13123 at initial time. For example, the premium paid at age 26 is equal to $0.089 \cdot (42090 - 13123) = 2667.06$.

Age	Salary (EUR)	Premium (%)
25	42090	0
26	44195	8.9
27	46404	8.9
28	50466	8.9
29	52989	8.9
30	54579	10.4
31	59300	10.4
32	62265	10.4
33	64133	10.4
34	69616	10.4
35	73097	12.0
36	75290	12.0
37	81662	12.0
38	84111	12.0
39	86635	12.0
40	93967	14.0
41	98665	14.0
42	103598	14.0
43	108778	14.0
44	117480	14.0
45	127440	16.3
46	133812	16.3
47	138495	16.3
48	141265	16.3
49	144090	16.3
50	146251	19.0
51	148445	19.0
52	149929	19.0
53	149929	19.0
54	149929	19.0
55	149929	22.3
56	149929	22.3
57	149929	22.3
58	149929	22.3
59	149929	22.3
60	149929	26.5
61	149929	26.5
62	149929	26.5
63	149929	26.5
64	149929	26.5
65	149929	30.6
66	149929	30.6
67	0	0

Table 2.1: Salary path and contribution percentages during the investment horizon.

2.C. Academic model

The following values have been used to construct the academic-model-based scenario set:

Variable	Value
a_ζ	0.7
b_ζ	-0.015
ε_ζ	$N(0.001, 0.018^2)$
ζ_0	0.01
a_w	0.65
b_w	-0.02
ε_w	$N(0.0065, 0.022^2)$
w_0	0.015
a_{MVF}	-0.2
b_{MVF}	0.1
$\varepsilon_{MVF,i}$	$N(0, \sigma_i^2)$
$r_{MVF,1}$	Lognormal (shape = 0.333, scale = 0.9267, loc = -0.8822)
M_0	4.93
a_m	0.9
b_m	0.1
ε_m	$N(0, 0.03^2)$
$M_{\min,i}$	$\begin{cases} 0.5, & \text{if } 2 \leq i \leq 17 \\ 1, & \text{if } 18 \leq i \leq 22 \\ 1 + (i - 22) \frac{7}{19}, & \text{if } i \geq 23 \end{cases}$
$M_{\max,i}$	$\min\{i + 11, 24\}$
σ_i	$\begin{cases} 0.2, & \text{if } i \leq 20 \\ 0.2 - (i - 20) \frac{0.135}{21}, & \text{if } i > 20 \end{cases}$

Table 2.2: Variables used to construct the academic-model-based scenario set.

3

Rule based strategies

This chapter shows rule-based strategies that are feasible solutions for Problem 2.3. In particular, specific rule-based strategies are introduced that perform well for an individual DC pension investor. Such strategies have, to our knowledge, not been studied anywhere else in literature. The rule-based strategies in this chapter have been designed from scratch.

The state variable for the rule-based strategies shown is the portfolio wealth, $Z_t = W_t$. A wealth target $W^*(t)$ is defined for all adjusting times $t \in \mathcal{T}$ and corresponds to a replacement ratio target over the time horizon. Wealth targets are stochastic: they are dependent on annual inflation i_t , salary q_t and a minimal required return r .

The investor's target for this chapter is to reach a replacement ratio of 70%. The replacement ratio has been defined in Definition 1.1. The risk measure is chosen to be the 10% CVaR.

3.1. Cumulative target strategy

The cumulative target strategy is a dynamic, risk-averse investment strategy. The goal is to reach a replacement ratio target with minimal risk exposure. The strategy aims for a replacement ratio target and decreases risk exposure because excess pay-offs are no longer rewarded. The cumulative target strategy is similar to strategies studied in [23] and [11]: a wealth target is introduced such that excess wealth can be defined and risk can be reduced once sufficient progress has been made. The cumulative target strategy does not invest excess wealth in a risk-free asset. Instead, part of the pension income is bought and secured until the end of the investment horizon. At each adjusting time, the price of a bond securing a target pension income is compared to the portfolio value of the investor.

Strategy 3.1 (Cumulative target strategy). The cumulative target strategy compares portfolio wealth W_t to a wealth target $W^*(t)$ for $t \in \mathcal{T}$. The wealth target $W^*(t)$ represents the price of a bond with 20 annual coupons of value $\tilde{W}(t)$ starting at the end of the investment horizon. Once $W_t \geq W^*(t)$, the investor buys the bond and keeps it until retirement: the portfolio wealth at t is secured in the matching portfolio. W_t and all future returns gained by W_t remain invested in the matching portfolio for the remainder of the investment horizon, $[t, T]$.

If the target value has not been reached, the investor buys return portfolio shares instead.

Outcomes are compared to static mixes and static life cycles. Improvements in goal-reachability while reducing the risk exposure are pursued, as stated in problem 2.3.

3.1.1. Wealth target construction

Wealth target $W^*(t)$ is constructed as follows. Future annual pension incomes are discounted and translated into a value at time t . It is assumed the investor receives an annual pension income for twenty years at retirement. Pension incomes can be discounted using a post-retirement discount rate r_p . Let

$$R(T) = \sum_{k=0}^{19} \frac{1}{(1+r_p)^k}, \quad (3.1)$$

such that $R(T)$ represents the value of an annual income of €1 for 20 years starting at the end of the investment horizon T .

This value can be discounted further using a discount rate r . Increasing r , increases the wealth target and thus the risk the investor has to take in order to reach the target:

$$R(t_i) = \frac{R(T)}{[1 + r + I(T; t_i)]^{T-t_i}},$$

the value of an annual income of €1 for 20 years starting at time T , at time $t_i \in \mathcal{T}$. Here, I is equal to

$$I(T; t_i) = \mathbb{E} \left[\zeta_\tau \mid \mathcal{F}_{t_i} \right], \quad (3.2)$$

the expected annual inflation for the remainder of the time horizon given filtration \mathcal{F}_{t_i} , for $\tau \in \{t_{i+1}, \dots, T\}$. For simplicity, the expected annual inflation is assumed to be equal for years in the remainder of the time horizon, $[t_{i+1}, T]$.

The discount rate can be used to impose a real annual return requirement, independent of inflation, when annual inflation is included in the wealth target definition.

Now consider

$$\widetilde{W}(t_i) = \sum_{k=0}^i \frac{c_k}{\widetilde{R}(k; t_i)}. \quad (3.3)$$

with c_k the annual contribution to the pension at time t_k , also called premium or cash flow. The annual inflation ζ_i up to time t_i is known, thus

$$\widetilde{R}(k; t_i) = \frac{R(t)}{\prod_{n=k+1}^i (1 + r + \zeta_n)}.$$

At time t_i , inflation rates $\zeta_1, \zeta_2, \dots, \zeta_i$ are known. $\widetilde{W}(t_i)$ is known at time t_i and represents the expected annual pension income based on the contributions made up to time t_i and an expected annual return of r for the remainder of the investment horizon. The market price at time t_i of an annual income in the future can be easily calculated using the market value factor M_i , which has been introduced in Section 2.1. Thus, the target wealth W^* at time t_i equals

$$W^*(t_i) = M_i \widetilde{W}(t_i). \quad (3.4)$$

The value of $W^*(t)$ is dependent on all contributions to the portfolio up to time t . $W^*(t)$ is therefore known as the cumulative target. Investors are able to build up a buffer for the future when returns exceed the return requirement of r . The opposite also holds: when returns are lower than required, returns in the future must compensate for losses experienced in the past.

It is assumed that inflation during the retirement period is accounted for in the market value factor M_i . Therefore, inflation is not included in Equation (3.1).

The expected annual inflation from Equation (3.2) is estimated by using regress now as defined in Definition 2.5 with $\Phi = \{1, x\}$. The future cumulative inflation $Y_{t_{i+1}}$ is regressed on the past cumulative inflation X_{t_i} to estimate the future annual inflation at time t_i :

$$x_j = \prod_{k=1}^i (1 + \zeta_k)$$

$$y_j = \prod_{k=i+1}^n (1 + \zeta_k),$$

with $x_j \in X_{t_i}$ and $y_j \in Y_{t_{i+1}}$ for each j in the scenario set.

The resulting regression function is of the form $f = ax + b$. The expected annual inflation for each scenario is equal to:

$$I(T; t_i | x = x_j) = \sqrt[n-(i+1)]{ax_j + b} - 1.$$

3.1.2. Target replacement ratio

The variable r used in the construction of the wealth target can be interpreted in multiple ways. First, it serves as a discount rate, used to compute the present value of contributions made in the future. Second, it can also be viewed as an annual return requirement: each contribution is required to have an average annual return of r . A third interpretation of r is that of a future expected annual return. The computation of the expected replacement ratio requires a future annual return assumption:

Definition 3.2 (Expected replacement ratio). Let $t \in \mathcal{T}$ and let \mathcal{F}_t be the filtration at time t . The expected replacement ratio RR_t is defined as

$$RR_t := \frac{\mathbb{E}[P|\mathcal{F}_t]}{\mathbb{E}[I_{\text{avg}}|\mathcal{F}_t]},$$

where

$$\mathbb{E}[P|\mathcal{F}_t] = \frac{\mathbb{E}[W_T|\mathcal{F}_t]}{\mathbb{E}[M_T|\mathcal{F}_t]},$$

with

$$\mathbb{E}[W_T|\mathcal{F}_t] = [1 + r + I(T; t)]^{T-t} W_t + \sum_{k=t+1}^{T-1} [1 + r + I(T; t)]^{T-k} \mathbb{E}[c_k|\mathcal{F}_t].$$

Computation of the expected replacement ratio requires four estimators. The discount rate r is used as estimator for future expected annual return. The estimator for the future inflation $I(T; t)$ is the same as used for Equation (3.2). The estimator for the future annual contributions is similar to the inflation estimator. Future salaries are estimated by estimating the future wage inflation. The age-dependent percentages of salaries that will be contributed in the future are assumed to be given. Last, the estimator for the market value factor at the end of the investment horizon, $\mathbb{E}[M_T|\mathcal{F}_t]$, uses regress now between M_t and M_T with $\Phi = \{1, x\}$.

The market value factor is considered independent of discount rate r , inflation and wage inflation. The division operator can therefore be taken out of the expected value operator.

The computation of the target replacement ratio at time t is similar to the computation of the expected replacement ratio. The only difference is that the portfolio wealth W_t is replaced with the target wealth $W^*(t)$. The target wealth definition causes the target replacement ratio $RR^*(t)$ to be independent of the market value factor:

$$RR^*(t) = \frac{\mathbb{E}[P^*|\mathcal{F}_t]}{\mathbb{E}[I_{\text{avg}}|\mathcal{F}_t]},$$

and by using Equation (3.4),

$$\begin{aligned} \mathbb{E}[P^*|\mathcal{F}_t] &= \frac{\mathbb{E}[W^*(T)|\mathcal{F}_t]}{\mathbb{E}[M_T|\mathcal{F}_t]} \\ &= \mathbb{E}[\widetilde{W}(T)|\mathcal{F}_t], \end{aligned}$$

because M_T is independent of $\widetilde{W}(T)$.

The investor's target for this chapter is to reach a replacement ratio of 70%. Wealth target function $W^*(t)$ defines a wealth target in terms of portfolio wealth instead:

$$\mathbb{E}[\widetilde{W}(T)|\mathcal{F}_t] = RR^*(t) \mathbb{E}[I_{\text{avg}}|\mathcal{F}_t].$$

To steer towards a fixed replacement ratio target, $\widetilde{W}(t)$ would have to be altered for each scenario. It is easier to allow $RR^*(t)$ to differ slightly between scenarios from a computational point of view. $\widetilde{W}(t)$ remains as defined in Equation (3.3). Numerically, the target replacement ratio within a scenario is almost constant throughout time, as can be seen in the bottom-right plot of Figure 3.1. Small alterations are caused by the estimators for inflation and wage inflation. Alterations of up to 0.01 within a scenario are observed for a discount rate of 2.5%. Target replacement ratios are between 0.6847 and 0.7033 for a discount rate of 2.5%.

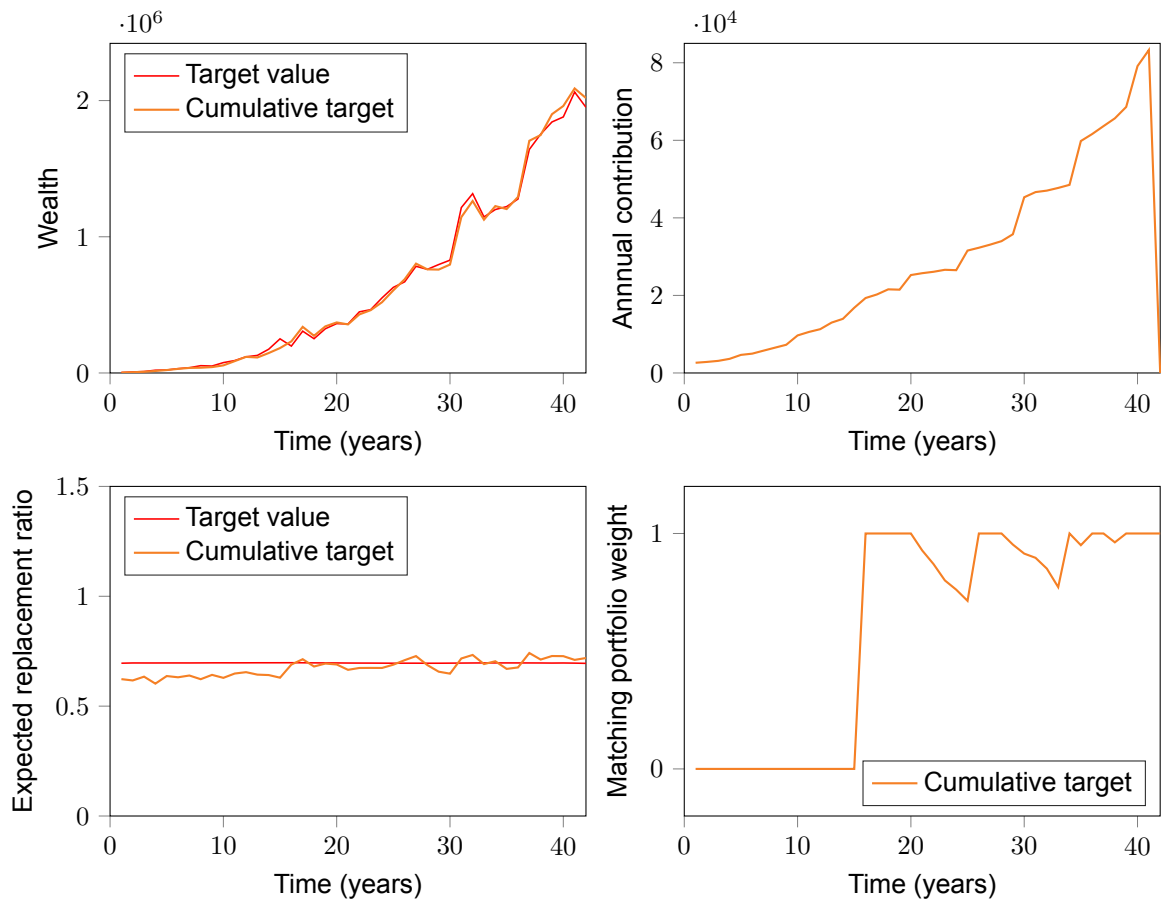


Figure 3.1: Example of a scenario simulation using the cumulative target strategy.

3.1.3. Results

Figure 3.1 shows an example of a single scenario simulation. The top-left plot compares the portfolio wealth to the wealth target over the time horizon. A decrease in market value factor may also cause a decrease in target value. The top-right plot shows the annual contribution for each adjusting time. No contribution is made in the final year. The bottom-left plot shows the expected replacement ratio and the target replacement ratio. The bottom-right plot shows the decision making process of the cumulative target strategy. The matching portfolio weight is equal to 1 at time t when the wealth target at time t has been reached. The first time this happens in Figure 3.1 is in year 16. The portfolio value exceeds the target value for the next four years as well, thus the matching portfolio weight remains at 1. In year 21, the portfolio wealth drops below the target value: the contribution of year 21 is invested in the return portfolio. The remaining portfolio wealth stays in the matching portfolio, because it was secured in the matching portfolio at the previous adjusting time.

Discount rate r relates to the risk nature of the investor. A higher discount rate implies a higher annual return requirement. Investors must allocate more wealth to the return portfolio to increase their returns, increasing their risk profile. Figure 3.2 shows a comparison between the cumulative target strategy for different discount rates, static mixes and static life cycles. The shortage measure and 10% CVaR are shown. The shortage measure has a target replacement ratio of 70%. Post retirement rate r_p is set at the fixed value of $r_p = 2.5\%$ for all results.

The static life cycles shown are variants on Bogle's rule as shown in Section 2.2.1. The offensive life cycle starts at 70% return portfolio and decreases linearly to 0 starting 15 years before retirement. The other two life cycles start at 60% return portfolio and they start to decrease linearly 20 (neutral) and 25 (defensive) years before retirement.




Figure 3.2: Risk measure comparison for the cumulative target strategy.

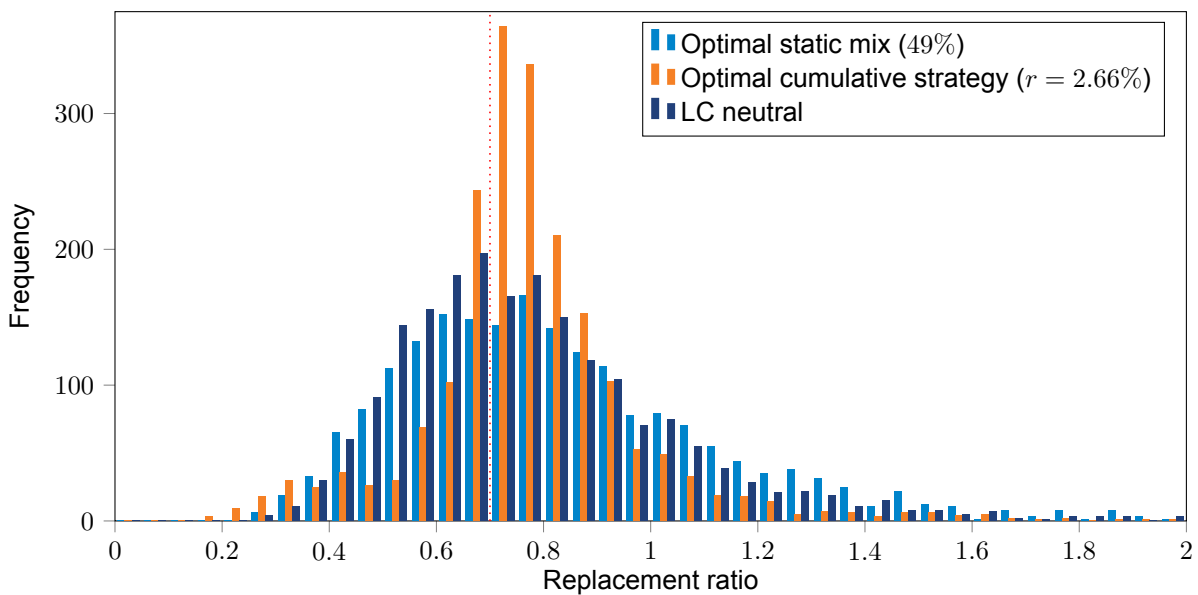


Figure 3.3: Replacement ratio distribution for different investment strategies. The optimal instance of a strategy is the instance with the lowest shortage measure.

The investor wishes to improve on goal-reachability compared to a static mix or static life cycle, without taking more risk. Such improvements can be found in the top-left corner of Figure 3.2. The cumulative target strategy with discount rates between 2% and 3% offers the most improvement in goal-reachability compared to static mixes or static life cycles. The risk level is slightly lower for the discount rate that gives the optimal shortage measure ($r = 2.66\%$, see Table 3.1).

Table 3.1 also shows the 5% CVaR risk measure. Compared to the 5% CVaR of static life cycles, this value is lower for all discount rates above 1.5%. Figure 3.3 explains this observation: it shows the

distribution of replacement ratios for different strategies, including the optimal static mix and optimal cumulative target strategy in terms of shortage.

In worst-case scenarios, the cumulative target strategy performs worse compared to static mixes or static life cycles. The defined wealth target is not reached, causing the investor to be invested for 100% in the return portfolio over the time horizon. The return portfolio is performing below expectation because the target is out of reach. The investor is stuck in a downwards spiral and it becomes increasingly more difficult to get out of it. Future investments must make up for the difference for every year the return portfolio performs below expectation.

Other statistics such as the median replacement ratio can also be found in Table 3.1.

Expected replacement ratio

One of the main advantages of a static life cycle strategy is the reduced risk exposure in the years before retirement. Investors create more certainty during these final years by increasing their matching portfolio weight. An accurate estimation for the pension income can be made during the final years of a static life cycle strategy.

The expected replacement ratio as defined in Definition 3.2 can be used as an estimator for the replacement ratio. Comparing the estimation error for different strategies gives insight in predictability of the cumulative target strategy.

Five years before terminal time T an estimation for the replacement ratio is made. Table 3.1 in the appendix of this chapter compares the estimation error between the cumulative target strategy and static life cycles. A positive prediction error implies the predicted value turned out to be too conservative. The volatility of such a prediction for the cumulative target strategy is comparable to the volatility of a static life cycle. The main cause for this is that for most scenarios, the return portfolio weights are similar for static life cycles and cumulative target strategies during the final investment years.

3.2. Individual target

The target value of the cumulative target strategy explained in Section 3.1 at time t is based on the cumulative contributions up to time t . The individual target strategy shown in this section defines a separate target for each premium.

Strategy 3.3 (Individual target strategy). The individual target strategy compares the progress of each annual contribution c_k to a wealth target $W_k^*(t_i)$ for all $t_k \leq t_i \in T$. The wealth target $W_k^*(t_i)$ represents the price of a bond with 20 annual coupons of value \widetilde{W}_k^* starting at the end of the investment horizon. Once $W_{k,t_i} \geq W_k^*(t_i)$, the investor buys the bond and keeps it until retirement: the portfolio wealth gained by contribution c_k , W_{k,t_i} , remains invested in the matching portfolio for the remainder of the investment horizon, $[t_i, T]$.

If the target value has not been reached, the investor invests W_{k,t_i} in the return portfolio instead. The investor repeats the strategy for each contribution made.

The individual target strategy tracks the growth of each individual contribution to the portfolio. Premiums paid at time t are required to undo any potential misfortune encountered during the period $[t_0, t]$ for the cumulative target strategy shown in Section 3.1. This is no longer the case for the individual target strategy: setbacks are only attempted to be compensated using premiums that encountered setbacks in the first place. The opposite phenomenon is also observed: new premiums no longer benefit from any potential surplus from previous years.

For the individual target strategy, the cumulative problem is split into n individual sub-problems, one for each annual contribution. Sub-problem k , $k \in \{0, \dots, n-1\}$ has an initial contribution of c_k and an investment horizon of $[t_k, T]$. Within sub-problem k , all other contributions are equal to 0.

The cumulative wealth target $W^*(t)$, as defined in Equation (3.4), is based on the cumulative premiums up to time t . A similar target can be defined for each individual premium. Let

$$\widetilde{W}_k^*(t_i) = \frac{c_k}{\widetilde{R}(k; t_i)},$$

such that

$$W_k^*(t_i) = M_i \widetilde{W}_k^*(t_i), \quad (3.5)$$

the individual target at time t_i for the premium paid at time $t_k \leq t_i$. Note that the sum of all individual targets is equal to the cumulative target from the first section of this chapter.

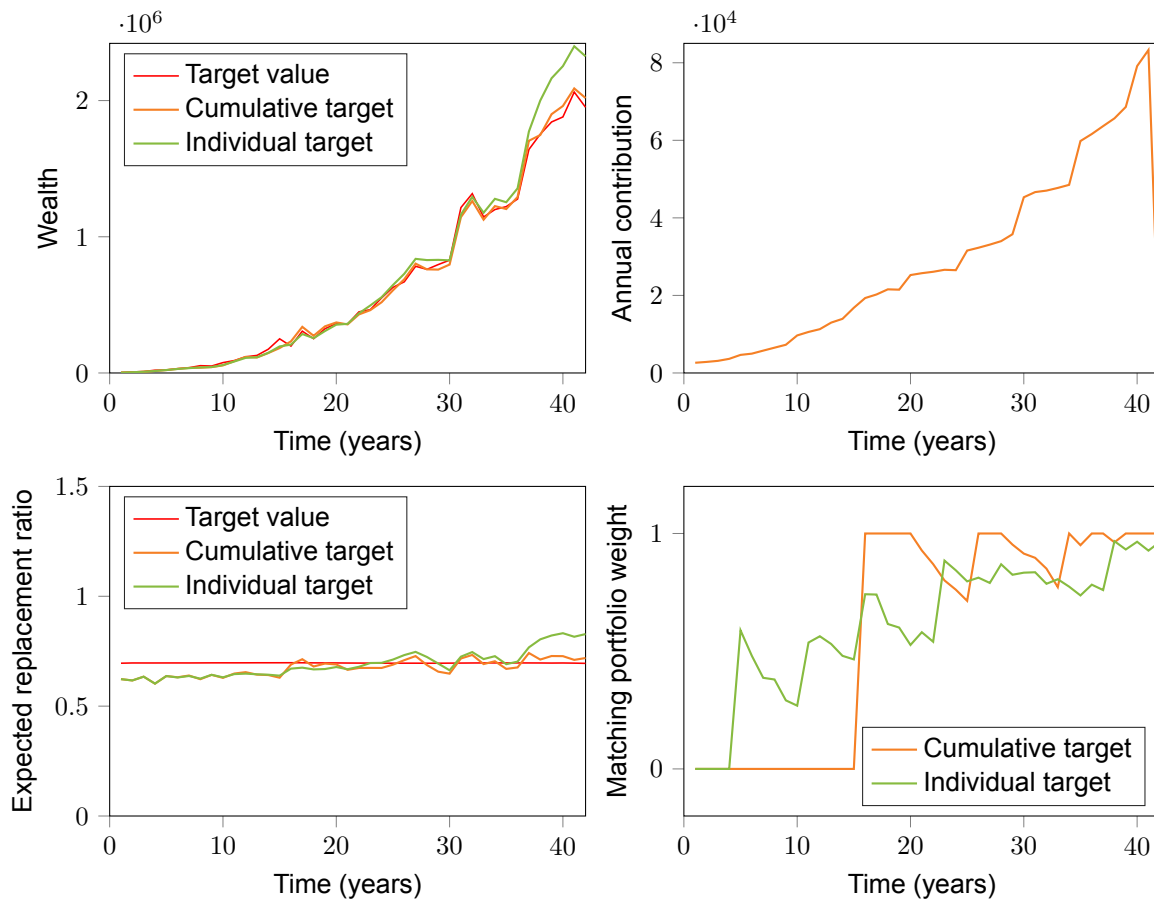


Figure 3.4: Example of a scenario simulation using the individual and cumulative target strategy.

Investment decisions within sub-problem k remain the same as for the cumulative target strategy. Once a premium reaches the target, this premium and all wealth gained by this premium are secured in the matching portfolio until final time T . If the individual target has not yet been reached, the premium remains invested in the return portfolio.

The investment decision of the individual target strategy is obtained by combining the investment decisions of all sub-problems.

3.2.1. Results

Figure 3.4 adds the individual target strategy to the previously shown scenario simulation. A matching portfolio weight of 1 at time t no longer implies reaching a target, it indicates all premiums up until t have reached their target. In some scenarios, the contribution at time t already exceeds the individual wealth target at time t . Individual targets can be relatively low due to a low market value factor M_t .

Figure 3.5 adds the individual target strategy to the results from figure 3.2: both the 10% CVaR and shortage measures are shown for different discount rates for the individual target strategy. Just like the cumulative target strategy, discount rates between 2% and 3% provide the most improvement regarding goal-reachability. The risk level of the individual target strategy is an improvement compared to static life cycles: the 10% CVaR measure is higher. Table 3.1 also shows the 5% CVaR for the individual target strategy. This value is still lower compared to static life cycles, but an improvement compared to the cumulative target strategy has been made.

Figure 3.6 shows the distribution of replacement ratios for different strategies, including the optimal static mix and optimal cumulative and individual target strategy in terms of shortage. The individual target strategy performs better in the left-sided tail compared to the cumulative target strategy. This is because some contributions are now able to reach their individual target, while the collective target remains out of reach.

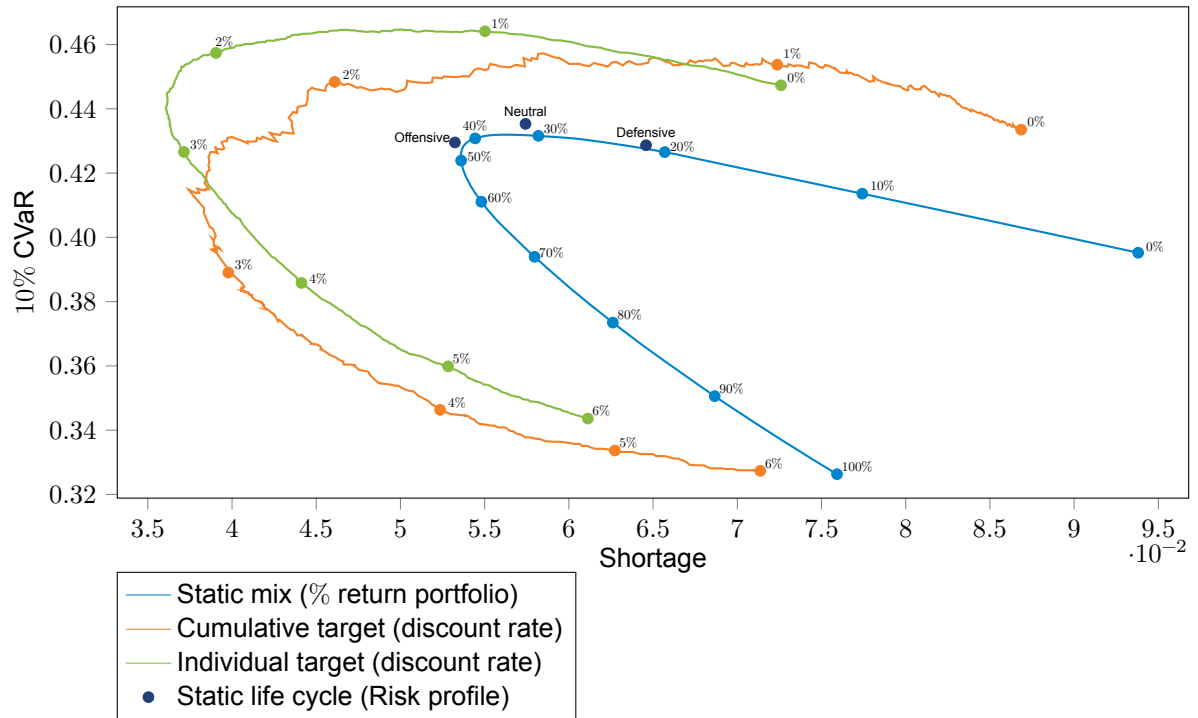


Figure 3.5: Risk measure comparison for the individual target strategy.

In worst case scenarios for the individual target strategy no contributions reach their individual target and the investor remains invested in the return portfolio for 100%. The discount rate is too ambitious for these scenarios: the required return cannot be reached.

Estimations of the replacement ratio 5 years before retirement have also been made for the individual target strategy. The estimation errors of static life cycles, the cumulative target strategy and the individual target strategy are comparable.

Other statistics for the individual target strategy such as the median replacement ratio can be also found in Table 3.1.

3.3. Constant cash flow

The cash flow seen in figures 3.1 and 3.4 increases over time. In Chapter 1, the assumptions for cash flows have been explained. Annual contributions increase over time due to two factors: an increase in salary over time and an increase in salary share p_i contributed. The goal of the Dutch government is to have a contribution to the pension savings that is independent of age in the future: investors have to contribute a fixed share π of their annual salary to their pension savings [17]. This section explores if the designed strategies are still of value for such a constant cash flow. The discounted value of the original and new cash flow should be equal for a fair comparison. The cumulative wealth target definition already includes a discounted cash flow. To find the fair percentage π , the following objective function is minimized:

$$\min [W^*(T; p_0, \dots, p_n) - W^*(T; \pi)]^2$$

where discount rate $r = 2.5\%$ and post-retirement rate $r_p = 2.5\%$.

3.3.1. Results

Figure 3.7 shows a comparison between the individual and cumulative target strategy for different discount rates, static mixes and static life cycles. Results are very similar to those found for an increasing cash flow.

Figure 3.8 shows the distribution of replacement ratios for different strategies, including the optimal static mix and optimal cumulative and individual target strategy in terms of shortage.

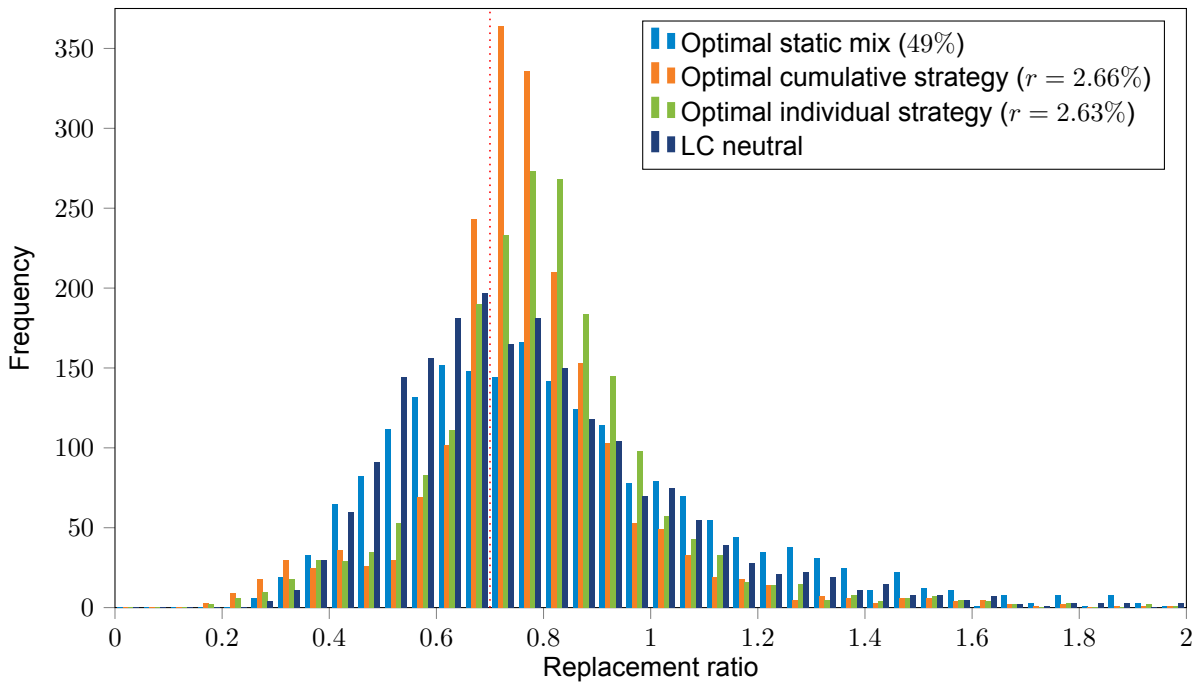


Figure 3.6: Replacement ratio distribution for different investment strategies. The optimal instance of a strategy is the instance with the lowest shortage measure.

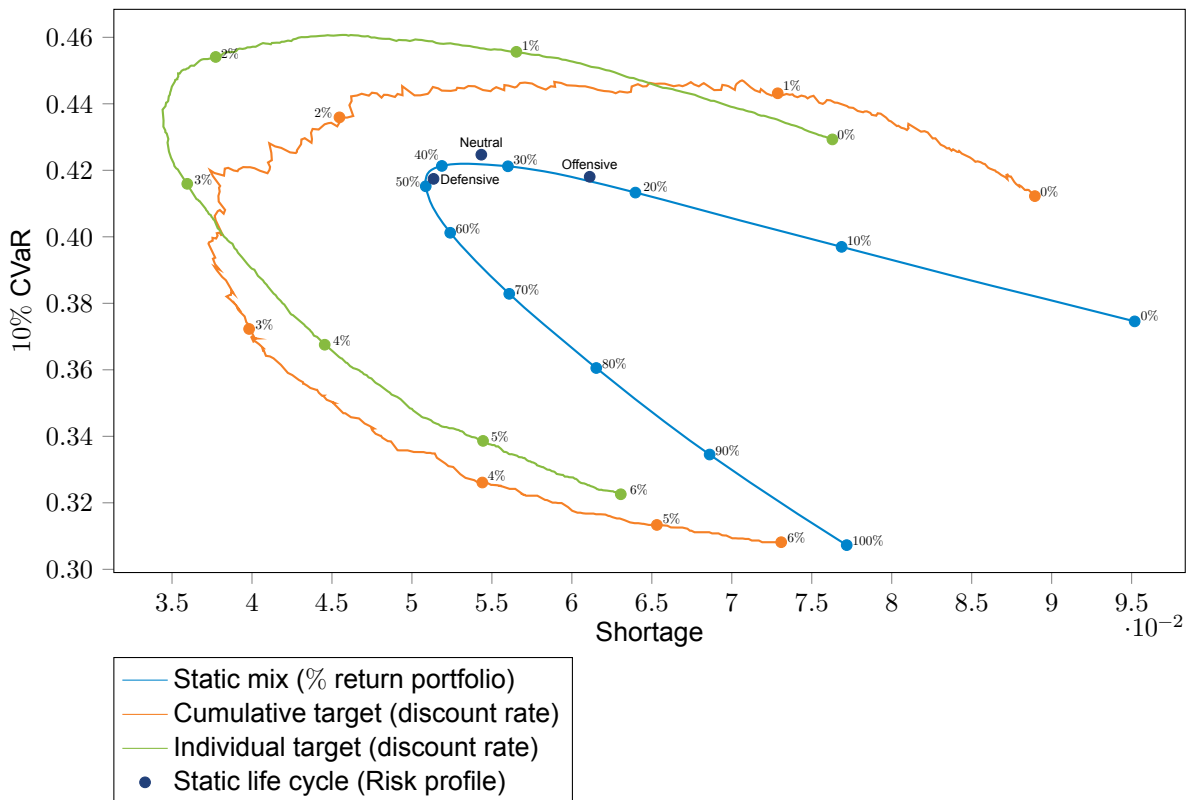


Figure 3.7: Risk measure comparison for the individual and cumulative target strategy for a constant cash flow.




Figure 3.8: Replacement ratio distribution for different investment strategies using a constant cash flow. The optimal instance of a strategy is the instance with the lowest shortage measure.

Table 3.2 provides all statistics that have also been discussed for the results of an increasing cash flow. Results of the two different cash flow scenarios are very similar and therefore the results of the constant cash flow scenario are not discussed in detail. Both the cumulative target strategy and individual target strategy can still be used in a constant cash flow scenario.

3.3.2. Academic model results

The results in this section have been computed by using scenarios from the academic model described in Section 2.1.

Results for the academic model can be found in Figures 3.10 and 3.9. Similar results as in the previous sections are found. The individual target strategy still provides the most improvement compared to a static mix or static life cycle. The difference between the cumulative and individual target strategy is smaller when using scenarios from the academic model. This can also be seen in the statistics found in Table 3.3.

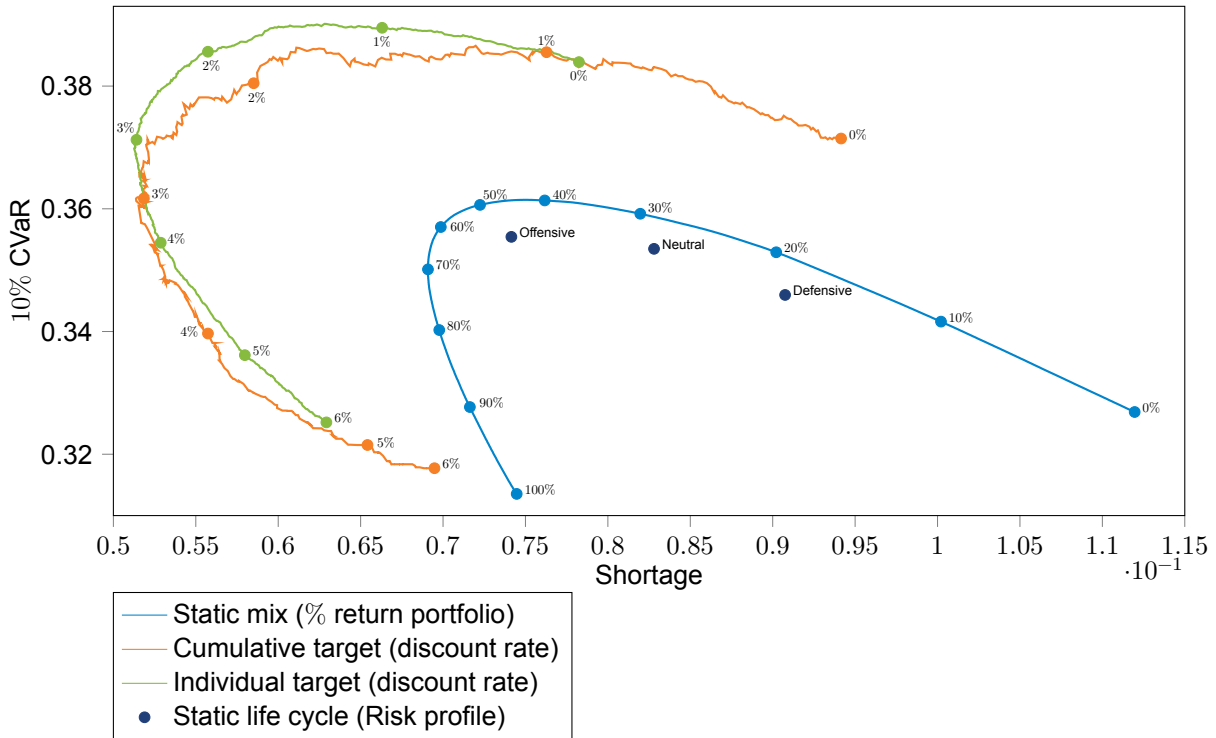


Figure 3.9: Risk measure comparison for the individual and cumulative target strategy using scenarios generated by the academic model.




Figure 3.10: Replacement ratio distribution for different investment strategies using scenarios generated by the academic model. The optimal instance of a strategy is the instance with the smallest shortage measure for that strategy.

3.4. Conclusion

Two rule-based strategies have been compared to static mixes and static life cycles by using simulation of representative scenarios. Both strategies make their decisions based on wealth targets depending on an annual return requirement r .

Cumulative target strategies shown in Section 3.1 have a constant cumulative target. Required performances from individual contributions may vary. In case of setbacks early on, the remaining premiums are required to also make up for under performances, on top of their own required performances. The opposite also holds: In case of a surplus, remaining premiums have effectively lower performance requirements.

Individual target strategies shown in Section 3.2 have a variable cumulative target. Required performances from individual contributions are constant. Each premium is expected to gain an annual return of r . The performance of a single premium does not influence the requirements for other premiums. The cumulative target can be viewed as variable: no attempts are made to fix or compensate under performances of premiums. This could be seen as lowering the cumulative wealth target.

Individual target strategies suit the risk profile of the investor better compared to cumulative target strategies. Investors wish to reach their replacement ratio target. Goal-reachability is displayed in the measure that sums the shortage, the CVaR indicates the shortfall in case of under performance. The shortage of the individual target strategy is lower compared to the cumulative strategy for all discount rates. The 10% CVaR is higher for individual target strategies compared to cumulative strategies for discount rates that correspond to a replacement ratio target of 70%.

A disadvantage of both rule-based strategies is the all-or-nothing investment decision of securing wealth in the matching portfolio. The portfolio remains 100% invested in the more risky return portfolio when targets are not reached. In such worst case scenarios the investor does not create any stability when following the cumulative or individual wealth target strategy.

The main advantage of a static life cycle strategy is high certainty regarding the expected pension income during the final years of the investment horizon. The expected replacement ratio at time t_{n-5} has been used to investigate this property for the two rule-based strategies. The standard deviation for the estimated replacement ratios are similar to those from static life cycle strategies.

Last, the cash flow shape has been adjusted to investigate the influence of relative large contributions made during the final years of the investment horizon. Results are comparable to results of increasing cash flows.

To summarize,

- Cumulative and individual target strategies improve goal-reachability compared to static mixes and static life cycles, while also reducing risk exposure.
- The individual target strategy is superior to the cumulative target strategy, because required performances for annual contributions are stable over the investment horizon.
- During the final years of the investment horizon, predictability of static life cycles and target strategies is similar.
- There is a possibility that the cumulative and individual target strategy remain 100% invested in the risky asset because no targets are met.

3.A. Strategy results

	Static mix			Static life cycle			Cumulative target			Individual target		
	0%	100%	49%	Def.	Neut.	Off.	2%	3%	2.66%	2%	3%	2.63%
Averages (<i>RR</i>)												
Mean	0.686	0.994	0.838	0.755	0.784	0.834	0.744	0.796	0.776	0.772	0.819	0.801
10% CVaR	0.395	0.326	0.425	0.429	0.435	0.430	0.448	0.389	0.415	0.457	0.427	0.440
5% CVaR	0.363	0.278	0.381	0.389	0.395	0.387	0.364	0.315	0.328	0.384	0.352	0.363
Percentiles (<i>RR</i>)												
Median	0.640	0.829	0.778	0.711	0.736	0.771	0.704	0.787	0.756	0.743	0.812	0.787
10% VaR	0.445	0.410	0.498	0.491	0.501	0.505	0.566	0.530	0.568	0.566	0.548	0.564
5% VaR	0.404	0.336	0.438	0.443	0.449	0.441	0.483	0.400	0.423	0.485	0.442	0.459
Goal (70% <i>RR</i>)												
Shortage	0.094	0.076	0.054	0.065	0.057	0.053	0.046	0.040	0.037	0.039	0.037	0.036
Goal reached	0.374	0.626	0.626	0.516	0.563	0.614	0.516	0.766	0.705	0.631	0.751	0.717
Estim. error (<i>RR</i>)												
Mean	0.003	0.141	0.063	0.009	0.022	0.043	0.030	0.000	0.021	0.040	0.001	0.021
Std dev.	0.091	0.396	0.182	0.100	0.111	0.145	0.107	0.126	0.118	0.107	0.122	0.115


Table 3.1: Simulation statistics for different investment strategies with an annual contribution percentage as in Table 2.1

	Static mix			Static life cycle			Cumulative target			Individual target		
	0%	100%	49%	Def.	Neut.	Off.	2%	3%	2.68%	2%	3%	2.51%
Averages (<i>RR</i>)												
Mean	0.709	1.090	0.900	0.806	0.841	0.902	0.775	0.840	0.816	0.805	0.865	0.843
10% CVaR	0.375	0.307	0.416	0.418	0.425	0.417	0.436	0.372	0.399	0.454	0.416	0.433
5% CVaR	0.341	0.257	0.368	0.374	0.380	0.369	0.339	0.293	0.308	0.374	0.335	0.347
Percentiles (<i>RR</i>)												
Median	0.651	0.872	0.822	0.746	0.776	0.811	0.722	0.824	0.787	0.763	0.849	0.818
10% VaR	0.429	0.392	0.499	0.484	0.497	0.498	0.569	0.534	0.569	0.574	0.551	0.566
5% VaR	0.382	0.317	0.426	0.429	0.436	0.429	0.476	0.381	0.403	0.471	0.429	0.447
Goal (70% <i>RR</i>)												
Shortage	0.095	0.077	0.051	0.061	0.054	0.051	0.045	0.040	0.038	0.038	0.036	0.034
Goal reached	0.416	0.642	0.664	0.571	0.610	0.654	0.572	0.794	0.747	0.660	0.781	0.752
Estim. error (<i>RR</i>)												
Mean	0.022	0.185	0.093	0.033	0.047	0.072	0.029	-0.004	0.007	0.037	0.003	0.020
Std dev.	0.100	0.458	0.206	0.114	0.127	0.166	0.115	0.137	0.129	0.115	0.131	0.121

Table 3.2: Simulation statistics for different investment strategies with an annual contribution of 17.3% of the salary.

	Static mix			Static life cycle			Cumulative target			Individual target		
	0%	100%	70%	Def.	Neut.	Off.	2%	3%	3.11%	2%	3%	3.12%
Averages (<i>RR</i>)												
Mean	0.840	1.205	1.088	0.951	0.981	1.035	0.914	0.966	0.972	0.929	0.975	0.981
10% CVaR	0.327	0.314	0.350	0.346	0.353	0.355	0.380	0.362	0.362	0.386	0.371	0.370
5% CVaR	0.303	0.265	0.300	0.314	0.319	0.316	0.318	0.294	0.293	0.326	0.306	0.304
Percentiles (<i>RR</i>)												
Median	0.707	0.934	0.904	0.806	0.830	0.875	0.787	0.873	0.885	0.823	0.897	0.905
10% VaR	0.370	0.409	0.433	0.397	0.409	0.422	0.481	0.470	0.472	0.483	0.479	0.478
5% VaR	0.330	0.323	0.358	0.354	0.366	0.367	0.379	0.377	0.378	0.399	0.386	0.386
Goal (70% <i>RR</i>)												
Shortage	0.112	0.074	0.069	0.091	0.083	0.074	0.059	0.052	0.051	0.056	0.051	0.051
Goal reached	0.506	0.657	0.652	0.573	0.597	0.623	0.619	0.728	0.738	0.657	0.720	0.728
Estim. error (<i>RR</i>)												
Mean	0.061	0.185	0.141	0.076	0.086	0.103	0.100	0.065	0.061	0.105	0.068	0.063
Std dev.	0.167	0.513	0.349	0.196	0.212	0.241	0.188	0.196	0.197	0.190	0.193	0.194

Table 3.3: Simulation statistics for different investment strategies using scenarios generated by the academic model.



Dynamic programming strategy

This chapter applies dynamic programming to Problem 1.5 by solving Problem 2.4 for all adjusting times in the investment horizon. The used algorithm and utility function are explained in the first section.

The second section studies a simple setting with a fixed wealth target. The investor starts with an initial wealth of 10 and has a target wealth of 50. The state variable used is the portfolio wealth, $Z_t = W_t$. Similar settings that use fixed target-based dynamic programming optimization have been researched before in literature, for example in [23].

The third section goes back to the DC pension investor problem. Dynamic programming is applied to the pension case setting studied in Chapter 3. The investor now has a stochastic target: a 70% replacement ratio. The replacement ratio has been defined in Definition 1.1. Such a stochastic target requires a different state variable: $Z_t = RR_t$ as defined in Definition 3.2. Similar settings have been studied before in literature, for example in the context of mean-variance optimization [22]. A dynamic programming approach with a stochastic replacement ratio target has, to our knowledge, not been studied anywhere else in literature.

The risk measure is chosen to be the 10% CVaR.

4.1. Algorithm

Both Problem 1.5 and Problem 2.4 use a utility function to measure final state Z_T . The utility function in this chapter uses a minimal target state z_{\min}^* and an upper bound target state z_{\max}^* . Let $U : \mathbb{R}_{>0} \rightarrow \mathbb{R}$ such that

$$U(z) = \frac{-(z - \beta)^2 - (z - z_{\min}^*)^2}{z}, \quad (4.1)$$

where

$$\beta = \sqrt{2(z_{\max}^*)^2 - (z_{\min}^*)^2},$$

with $z_{\min}^* < z_{\max}^*$. The utility function can be found in Figure 4.1. U is clearly concave and continuous on the domain $\mathbb{R}_{>0}$.

A utility function like the one in Figure 4.1 corresponds with the assumption that upward potential is intertwined with downward risk. The investor avoids unnecessary risk by decreasing the utility above the upper bound target state.

Asset allocations for a dynamic programming strategy follow from Algorithm 4.1. The algorithm runs backwards in time, similar to [8] and [19]. The solution space is discretized such that an algorithm can be used to solve Problem 2.4: it is assumed that

$$S : \mathcal{T} \times \mathbb{R}_{>0} \rightarrow \{a_1, a_2, \dots, a_k\}, \quad (4.2)$$

where $a_j \in [0, 1]$ for $j \in \{1, \dots, k\}$. The investor can choose between at most k asset allocations at every time t .




Figure 4.1: Utility function for Chapter 4.

Algorithm 4.1 uses regression between state Z_t and expected utility $\mathbb{E}[U(Z_T)|\mathcal{F}_t]$. The local regression technique that is used has been defined in Section 2.3.2. The use of local regression is similar to the use of bundling in [8]: neighbourhood points are used in regressions for each step of the algorithm.

Solving the sub-problem at time t_i changes states $Z_{t_{i+1}}, \dots, Z_{t_{n-1}}$. These states have previously been used in local regressions to find the optimal solution for the sub-problems at times t_{i+1}, \dots, t_{n-1} . Thus, the optimal solutions for sub-problems at t_{i+1}, \dots, t_{n-1} may be different after solving the sub-problem at time t_i . This is why Algorithm 4.1 follows a snake-like pattern through time: after the sub-problem at time t_i is solved, future sub-problems are first updated forward in time. Sub-problems are updated backward in time at the next step until the sub-problem at time t_{i-1} is solved for the first time.

Once all sub-problems have been solved, the solution can be further improved by repeating the procedure. Algorithm 4.1 restarts at the beginning of the snake-like pattern through time. Each iteration of Algorithm 4.1 follows the snake-like pattern from T to t_0 once.

A dynamic programming strategy has many different inputs and assumptions such as utility function U , choice of regression techniques or choice of state variable Z_t . The pension investor's optimality problem is also complex due to the annual contribution to the pension and the stochastic replacement ratio target. A simplified investment setting is studied first to get a better understanding of the dynamic programming setting.

4.2. Fixed wealth target

First, the setting with an initial wealth and fixed wealth target is used. The state variable Z is chosen to be the portfolio wealth, $Z_t = W_t$. A fixed wealth target G is defined at the end of the investment horizon, $W_T^* = G$. The investor starts off with an initial wealth $K < G$. No contributions are made to the portfolio during the investment period.

The objective of the investor is to reach the target wealth G at time T with minimal risk. Utility function U is as defined in (4.1) with $z_{\min}^* = G$ and $z_{\max}^* = 3G$.

4.2.1. Results

The discretization used for the results in this section is $\{a_1, \dots, a_m\} = \{0, 0.1, \dots, 1\}$. The investor can choose between 11 different asset allocations at every rebalancing time. The neighbourhood parameter of the local regression is chosen to be 0.75. Decreasing this value causes regressions $f_{j,i}$ to oscillate: a smaller neighbourhood parameter results in over-fitting and optimization on the data set itself.

Figure 4.2 compares the distribution of terminal wealth of the dynamic programming strategy and the optimal static mix in terms of shortage measure. The dynamic programming strategy reduces the area of the probability distribution on the left-hand side of the target: the target-reachability increases by using a dynamic programming strategy compared to a static mix. This is reflected by the values of the shortage measure found in Table 4.1.

Algorithm 4.1: Dynamic programming investment strategy**Data:** Ortec Finance Scenario Set, asset allocations a_1, \dots, a_k **Result:** Optimal admissible asset control \mathcal{A}^* Initialization: create initial solution by choosing $\alpha_i = a_1$ for all $t_i \in \mathcal{T} = \{t_0, \dots, t_n = T\}$ for all scenarios.

```

1 for  $m = 1, \dots$ , number of iterations do
2   for  $j = 1, \dots, k$  do
3     Set  $\alpha_{n-1} = a_j$  for all scenarios
4     Determine the expected utility function  $f_{j,n-1}$ ,
        $f_{j,n-1}(z) = \mathbb{E} [U(Z_T) | Z_{n-1} = z, \alpha_{n-1} = a_j]$ ,
       by using local regression for  $Z_{n-1}$  and  $U(Z_T)$  (Least squares Monte-Carlo method)
5   Determine optimal asset allocation decision at time  $t_{n-1}$  for all scenarios:
        $m = \arg \max_{j \in \{1, \dots, m\}} f_{j,n-1}(Z_{n-1})$ 
        $\alpha_{n-1}^* = a_m$ 
6   for  $t_i = t_{n-2}, t_{n-3}, \dots, t_0$  do
7     Update optimal allocation decisions (backward update):
       for  $\tau = n, \dots, i + 1$  do
8       Update expected utility function  $f_{j,\tau}$  for  $j = 1, \dots, k$ 
9       Update optimal asset allocation decision  $\alpha_\tau^*$  for all scenarios
10    Determine optimal asset allocation decision at time  $t_i$  for the first time
       for  $j = 1, \dots, k$  do
11      Set  $\alpha_i = a_j$  for all scenarios
12      Determine the expected utility function  $f_{j,i}$ ,
          $f_{j,i}(z) = \mathbb{E} [U(Z_T) | Z_i = z, \alpha_i = a_j, \mathcal{A}_{i+1} = \mathcal{A}_{i+1}^*]$ ,
         by using local regression for  $Z_i$  and  $U(Z_T)$  (Least squares Monte-Carlo method)
13      Determine optimal asset allocation decision at time  $t$  for all scenarios:
          $m = \arg \max_{j \in \{1, \dots, k\}} f_{j,i}(Z_i)$ 
          $\alpha_i^* = a_m$ 
14      Update optimal allocation decisions (forward update):
       for  $\tau = i + 1, \dots, n - 1$  do
15        Update expected utility function  $f_{j,\tau}$  for  $j = 1, \dots, k$ 
16        Update optimal asset allocation decision  $\alpha_\tau^*$  for all scenarios

```

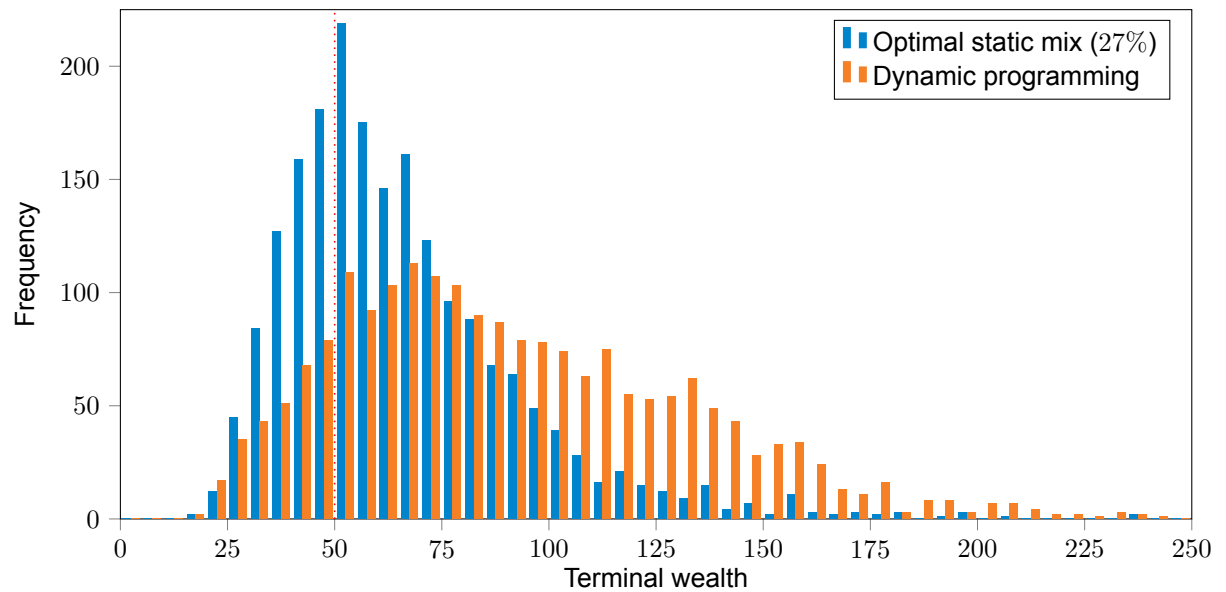


Figure 4.2: Terminal wealth distribution for the optimal static mix and dynamic programming strategy instances with initial wealth $K = 10$ and fixed wealth target goal $G = 50$. The optimal instance of a strategy is the instance with the lowest shortage measure.

The left-hand tail of the terminal wealth distribution also improves when using a dynamic programming strategy, although the performance in worst-case scenarios is slightly worse. The values for both the 5% and 10% CVaR are higher compared to a static mix.

Results are sensitive to local regressions used in Algorithm 4.1. Functions $f_{j,i}$ for different j become similar to each other as i increases: less time remains available to improve the portfolio wealth. The following section discusses the convergence of allocation decision and the dependence on regressions in more detail.

4.2.2. Convergence of the allocation decision

Figure 4.3 shows the allocation decision over time after 1 iteration of Algorithm 4.1.

Increasing the number of iterations causes small differences between Figure 4.3 and similar plots. This implies that the allocation decision has not converged after 1 iteration of Algorithm 4.1.

Differences are caused by the dependence of the allocation decision on functions $f_{j,i}$ found by using local regression. Running multiple iterations of Algorithm 4.1 changes the two sets of observations which are used when determining functions f . This causes intersections between $f_{j,i}$ for different j to change slightly, which in turn determine the interval in which strategy j is used for state variable Z_i .

These dynamics make it difficult to reach absolute convergence. It is assumed that one iteration of Algorithm 4.1 is sufficient to approach the optimal allocation decision. Multiple iterations of Algorithm 4.1 do not improve investment results: they are similar to results that have been computed using a single iteration of the algorithm. Allocation decisions over time change slightly when using more iterations. The difference in expected utility between two similar asset allocation decisions is small. The impact of differences in allocation decisions caused by using multiple iterations is therefore small.

The allocation intervals shown in Figure 4.3 are quite stable for the first 20 years of the investment horizon. In later years some instability can be observed. The main cause for instabilities are tiny improvements in expected utility at time t_i for different allocation decisions a_j . Functions $f_{j,i}$ for different j are very similar for $i \geq 20$. The sensitivity towards the regression increases and the asset allocation decision becomes increasingly unstable.

Statistics for the asset allocation over time can be found in Table 4.1. The asset allocation over time for individual scenarios can be unstable. Large jumps in asset allocations over a short time period are observed. The main reason these happen are tiny differences in expected utility. Requiring a minimal improvement regarding expected utility could provide improvements.

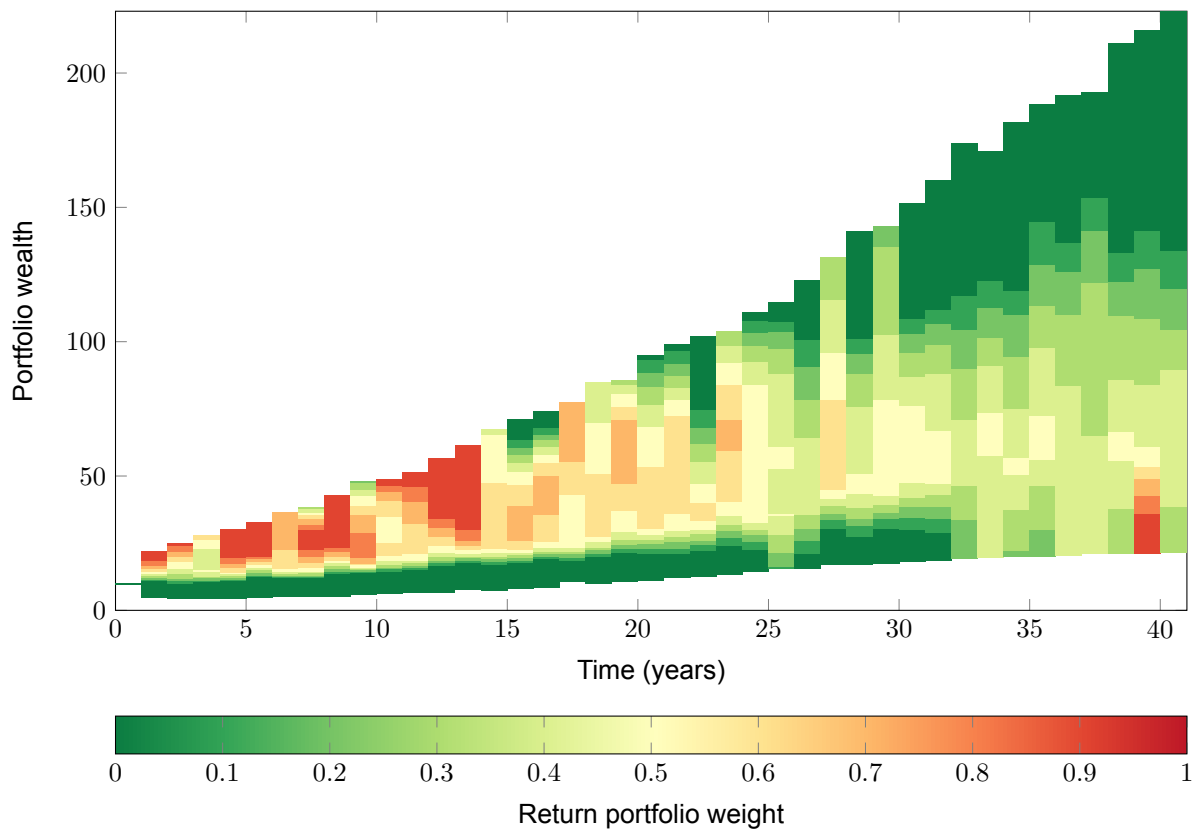


Figure 4.3: Asset allocation decisions over the investment horizon for a dynamic programming strategy with a fixed wealth target. The return portfolio weight is shown from 0 (green) to 1 (red) at each adjusting time.

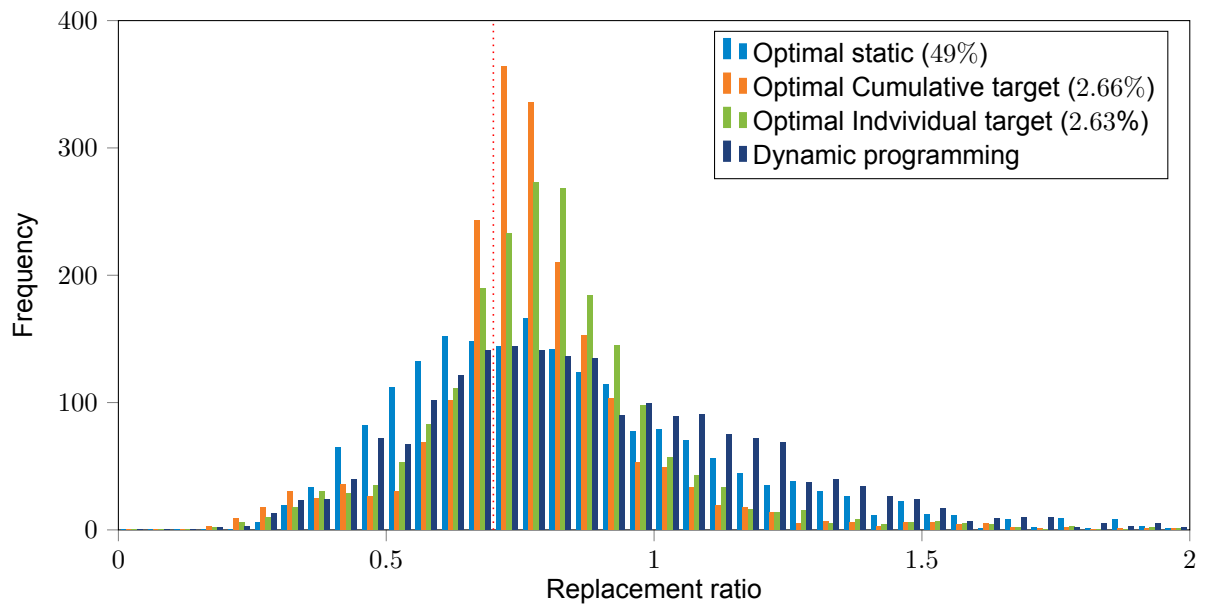


Figure 4.4: Replacement ratio distribution for the optimal static mix, optimal results from Chapter 3 and the dynamic programming strategy. The optimal instance of a strategy is the instance with the lowest shortage measure.

4.3. Stochastic wealth target

The previous section showed that dynamic programming combined with the utility function defined in Equation (4.1) provides improvements in reachability of the target and reduces downside risk for a setting with an initial wealth and a fixed wealth target.

This section goes back to the pension investment problem previously explained in Section 2.1 and studied in Chapter 3. This problem does not have an initial wealth. Instead, annual contributions are made to the portfolio. The pension investment problem also does have a fixed wealth target. It has a stochastic wealth target: a 70% replacement ratio.

Portfolio wealth W_t is not a sufficient for such a target. The same value for W_T may result in different replacement ratios between scenarios. A replacement ratio target requires a state variable that is expressed in terms of the target. A sufficient state variable is the expected replacement ratio RR_t , which has been defined in Definition 3.2. The utility function U used is as defined in (4.1) with $z_{\min}^* = 0.7$ and $z_{\max}^* = 2.1$.

4.3.1. Results

The discretization used for the results in this section is $\{a_1, \dots, a_m\} = \{0, 0.1, \dots, 1\}$. The investor can choose between 11 different asset allocations at every rebalancing time. The neighbourhood parameter of the local regression is chosen to be 0.6. Decreasing this value causes regressions $f_{j,i}$ to oscillate: a smaller neighbourhood parameter results in over-fitting and optimization on the data set itself.

Figure 4.4 compares the distribution of the replacement ratio between the dynamic programming investment strategy and the optimal static mix in terms of shortage. Note that this is the same optimal static mix as the one found in Figures 3.3 and 3.6. Risk measures and other statistics can be found in Table 4.2. The optimal instances from the cumulative target strategy and the individual target strategy have been added to Table 4.2 to compare between dynamic programming and rule-based strategies.

Dynamic programming strategy results are worse compared to those of rule-based strategies. The shortage measure is higher for the dynamic programming strategy, while the CVaR values are slightly lower. Rule-based strategies make use of market value factor M_t to construct the price of a bond that secures a target pension income at each adjusting time. The dynamic programming strategy does not use the information, because no target pension incomes have been defined at any adjusting times.

Another cause for this difference is the instability of the allocation decision over time as seen in




Figure 4.5: Asset allocation decisions over the investment horizon for a dynamic programming strategy with a replacement ratio target. The return portfolio weight is shown from 0 (green) to 1 (red) at each adjusting time.

Figure 4.5. Compared to Figure 4.3, the dynamic programming strategy for a replacement ratio target is unstable. The main cause for this is the difference in state variable used. For a fixed wealth target, a change in state variable (the portfolio wealth) is caused by the portfolio return. The stability of the portfolio return and the asset allocation decision are closely related: A low value of α_t decreases the volatility of the portfolio return.

For a stochastic wealth target (a replacement ratio of 70%), a change in state variable (the expected replacement ratio) is not only caused by portfolio return. A difference in expected wage inflation and the real wage inflation causes differences in (expected) salaries and annual contributions, thus also changing the expected replacement ratio. Differences in price inflation or expected market value factor at time T also cause changes in expected replacement ratio. The state variable for a stochastic wealth target is not only dependent on the portfolio return, increasing the difficulty of finding the relationship between the allocation decision and the change in state variable.

The sensitivity to regressions increases for a replacement ratio target: the annual contribution mechanism causes the investment decision at a time t_i to be of less influence compared to the fixed wealth target. Regressions $f_{j,i}$ for different j are almost identical: the difference in expected utility is tiny.

Large contributions are made during the final years. The investment decisions in these years carry more weight because the portfolio wealth has increased significantly compared to years at the beginning of the investment horizon.

4.4. Conclusion

Dynamic programming has been applied to a simple setting with a fixed wealth target in Section 4.2 and a complex setting with a stochastic wealth target in Section 4.3. Investors can use dynamic programming to steer towards a pre-defined target. Distributions of outcomes are more centred around the target value and the area below the target value is smaller compared to static mixes.

The dynamic programming investment strategy has some drawbacks. The impact of these drawbacks grows as the problem setting becomes more complex. First, restrictions encountered for the fixed wealth target are discussed.

The simple setting with a fixed wealth target shown in Section 4.2 is sensitive to the regressions used. The optimal asset allocation is chosen based on the maximum expected utility. Regressions are used to estimate the expected utility for each allocation decision at time t . These regressions are often similar: there is only a small difference in expected utility for different allocation decisions.

Small differences between regressions also cause the asset allocation over time to be volatile for some scenarios. The annual turnover for these scenarios is large. Investment decisions of the dynamic programming strategy are therefore not yet applicable in practice. Chapter 5 discusses techniques to decrease the annual turnover.

In the more complex pension case setting shown in Section 4.3, the impact of these drawbacks increases. Regressions used to estimate the expected utility become even more similar. This is largely due to the switch in state variable. For the simple setting, the portfolio wealth is sufficient as state variable because the target value is a fixed wealth level. Changes in portfolio wealth are only caused by portfolio return.

The pension case setting has a target that is stochastic: the replacement ratio. The state variable used is changed to the expected replacement ratio such that it can be compared to the target value. The downside of this approach is that changes in the state variable are not only caused by portfolio, but also by changes in estimation of random variables such as salary, inflation and market value factor. This reduces the dependency of the state variable on the allocation decision and causes expected outcome to be more similar for different allocation decisions.

To summarize,

- Dynamic programming investment strategies can be used to steer towards a fixed wealth target in a simplified setting and a more complex setting.
- Asset allocations over time have a large annual turnover in some scenarios and are not yet applicable in practice.
- The instability of asset allocations over time increases when the problem setting becomes more complex.
- The expected replacement ratio state variable makes it more difficult to find the relationship between allocation decision and change in state variable.

4.A. Strategy results

The tables in this section contain statistics for the dynamic programming strategy instances as seen in Figures 4.2 and 4.4. The asset allocation statistics give an insight of the behaviour of the asset allocation over time. They are to be interpreted as follows:

- **Average weight change:** the average asset weight change per adjusting time among all scenarios. For example, the asset weight for a dynamic programming strategy changes on average 0.119 per adjusting time.
- **Average jump:** the average asset weight change given that the the asset allocation is altered. For example, if the asset weight for a dynamic programming strategy changes, it jumps by an average value of 0.183.
- **Asset weight changes:** the share of adjusting time in which the asset allocation is altered. For example, the asset allocation for a dynamic programming strategy is altered in 64.9% of the adjusting times.
- **Maximum asset weight changes:** the share of adjusting time in which the asset allocation change is maximum. For example, the asset allocation for a dynamic programming strategy is altered from 0 to 1 or from 1 to 0 in 0.1% of the adjusting times.

Example values have been taken from Table 4.1.

	Static mix 27%	Dynamic Programming $\alpha_t \in [0, 0.1, \dots, 1]$
Averages (Wealth)		
Mean	66.00	91.64
10% CVaR	31.95	34.44
5% CVaR	28.58	29.03
Percentiles (Wealth)		
Median	59.87	84.25
10% VaR	37.49	44.32
5% VaR	32.84	35.25
Goal (Wealth of 50)		
Shortage	3.004	1.698
Goal reached (%)	0.695	0.853
Asset allocation ($\Delta\alpha_t$)		
Average weight change	0	0.119
Average jump	0	0.183
Asset allocation(Ratio of adjusting times)		
Asset weight changes	0	0.649
Maximum asset weight changes	0	0.001

Table 4.1: Simulation statistics for the dynamic programming strategy with a fixed wealth target of 50 and an initial wealth of 10.

	Static mix	Rule Based		Dynamic Programming
	49%	Cum. Target	Ind. Target	$\alpha_t \in [0, 0.1, \dots, 1]$
Averages (<i>RR</i>)		2.66%	2.63%	
Mean	0.838	0.776	0.801	0.894
10% CVaR	0.425	0.415	0.440	0.423
5% CVaR	0.381	0.328	0.363	0.363
Percentiles (<i>RR</i>)				
Median	0.778	0.756	0.787	0.836
10% VaR	0.498	0.568	0.564	0.518
5% VaR	0.438	0.423	0.459	0.447
Goal (70% <i>RR</i>)				
Shortage	0.054	0.037	0.036	0.044
Goal reached (%)	0.626	0.705	0.717	0.696
Estim. error (<i>RR</i>)				
Mean	0.063	0.011	0.021	0.073
Std dev.	0.182	0.118	0.115	0.189
Asset allocation ($\Delta\alpha_t$)				
Average weight change	0	0.034	0.046	0.175
Average jump	0	0.181	0.071	0.283
Asset allocation(% of adjusting times)				
Asset weight changes	0	0.188	0.653	0.618
Maximum asset weight changes	0	0.022	0.004	0.029

Table 4.2: Simulation statistics for the dynamic programming strategy with a replacement ratio target of 70%.

5

Control smoothing

This chapter shows methods to decrease the annual turnover of asset allocations found in chapter 4. The stability of the allocation decision over time is improved. Three approaches are investigated: the introduction of fictive transaction costs, the addition of a bandwidth limitation and a requirement of minimal improvement. All three approaches add a smoothing term to the value function of a dynamic programming sub-problem. The smoothing terms can be compared to a regularization term used in Lasso regression to regularize ordinary least squares regression. The smoothing terms have, to our knowledge, not been combined with a dynamic programming optimization problem before in existing literature.

5.1. Fictive transaction costs

Fictive transaction costs discourage investors to make large changes in asset allocations between consecutive adjusting times. The goal is to increase stability of the allocation decision over time. The four statistics describing the asset allocation found in Section 4.A are expected to improve when fictive transaction costs have been added to the model.

Fictive transaction costs are implemented in the model by expanding the value function from sub-problem 2.4. Expected utility at time T decreases depending on the annual turnover. The value function is modified into

$$V(z, t_i) = \sup_{\mathcal{A}_i} \mathbb{E} \left[U(Z_T) - \lambda(z) |\alpha_i - \alpha_{i-1}| \mid Z_{t_i} = z \right], \quad (5.1)$$

for $t_0 \neq t_i \in \mathcal{T}$, with λ a function representing fictive transaction costs. Fictive transaction costs are dependent on the asset allocation at the previous adjusting time. Algorithm 4.1 determines the optimal asset allocation backwards in time. Fictive transaction costs at time t_i are dependent on the optimal asset allocation at time t_{i-1} . Thus, α_i^* is dependent on α_{i-1}^* when adding fictive transaction costs as in Equation (5.1), while α_{i-1}^* is computed after α_i^* when optimizing backwards in time.

Fictive transaction costs can be implemented as in Equation (5.1) because of the snake-like pattern through time used in Algorithm 4.1. A single iteration of Algorithm 4.1 is summarized in Algorithm 5.1. At step t_i , the optimal allocation at time t_i is determined for the first time. Next, the optimal allocation at times t_{i+1}, \dots, T are updated forward in time. This can also be seen as solving a sub-problem with investment horizon $[t_i, T]$ forward in time.

Solving the sub problem with investment horizon $[t_i, T]$ forward in time means that $\alpha_i^*, \alpha_{i+1}^*, \dots, \alpha_n^*$ are computed consecutively. Fictive transaction costs can be added to the sub-problem at times t_{i+1}, \dots, T , because α_{i+1}^* is computed after α_i^* .

The restriction applies that the fictive transaction costs at time t_i for a sub-problem with investment horizon $[t_i, T]$ cannot yet be computed, because α_{i-1}^* is computed in the next step of Algorithm 4.1. Therefore, they are set to 0 when solving the sub-problem with investment horizon $[t_i, T]$.

Algorithm 4.1 can therefore also be interpreted as solving a sequence of sub-problems. Each sub-problem is solved forward in time, as seen in Algorithm 5.1. The investment horizon of these sub-problems increases backwards in time with every step: First, the sub-problem with investment horizon




Figure 5.1: Replacement ratio distribution of dynamic programming strategy instances with different levels of fictive transaction costs. The orange histogram for $\rho = 0$ is the same as the dark-blue dynamic programming histogram seen in figure 4.4.

$[t_{n-1}, T]$ is solved. The next step solves the sub-problem with investment horizon $[t_{n-2}, T]$ forward in time. An assumption regarding the optimal asset allocation at t_{n-1} is required: the solution of the sub-problem solved in the previous step is used.

Algorithm 5.1: Summary of a single iteration of Algorithm 4.1

- | |
|---|
| <ol style="list-style-type: none"> 1 for $t_i = t_{n-2}, t_{n-3}, \dots, t_0$ do 2 Update optimal allocation decisions (backward update) for t_{n-1}, \dots, t_{i+1} 3 Solve sub-problem with investment horizon $[t_i, T]$ forward in time |
|---|

Fictive transaction costs are not only dependent on the difference in asset allocation between two adjusting times, they are also dependent on the portfolio wealth. Investors pay more fictive transaction costs if they wish to buy more stocks or bonds. More stocks are required to adjust the portfolio balance if the portfolio wealth is larger. The expected utility variable has already been scaled towards portfolio size. Therefore, let

$$\lambda(z) = \rho \cdot \mathbb{E} [U(Z_T) | Z_{t_i} = z], \quad (5.2)$$

where $\rho > 0$ is a constant.

5.1.1. Results

Figure 5.1 shows results of some dynamic programming strategy instances with fictive transaction costs applied, for different values of smoothing parameter ρ .

The statistics in Table 5.1 show that the stability of the asset allocation is dependent on the level of fictive transaction costs. A higher value of smoothing parameter ρ increases the stability of the asset allocation over time. The same behaviour can also be seen in Figure 5.2. It shows the asset allocation over time for different levels of fictive transaction costs.

Fictive transaction costs can be too high, as can be seen for a value of $\rho = 0.1$. In such a case, it is too expensive to change the asset allocation. The investor is stuck with the initial asset allocation.

In more realistic cases of fictive transaction costs, for smoothing parameters $\rho = 0.001$ and $\rho = 0.01$, the stability of the asset allocation over time improves. The number of asset weight changes over

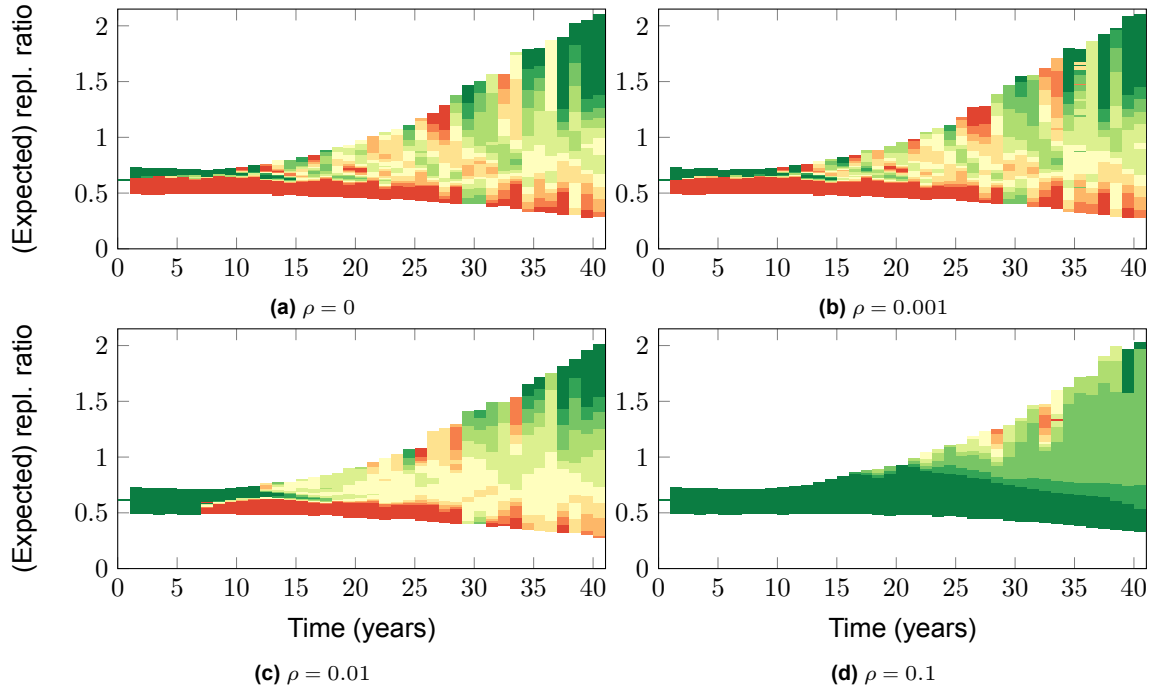


Figure 5.2: Asset allocation over time for different levels of Fictive transaction costs. Colours indicate the value of the return portfolio weight between 0 (green) and 1 (red), identical as in Figure 4.5. Figure 5.2a is the same plot as Figure 4.5.

time decreases, as well as the difference in asset allocation between two consecutive adjusting times (annual turnover). This comes at the cost of optimality: the investor can no longer choose any asset allocation without consequences as is the case with $\rho = 0$.

Results in terms of shortage measure and various risk measures are therefore worse compared to results found in Chapter 4. This can also be seen in Figure 5.1. Figure 5.7 compares the results of the various methods of control smoothing to the results of rule-based strategies and the unrestricted dynamic programming strategy from Chapter 4.

5.2. Bandwidth limitation

Alternatively, a bandwidth can be introduced to prevent big jumps in the asset allocation. It is used to decrease the annual turnover and improve the stability of the allocation decision over time.

Given asset allocation $\alpha_i = a$ at time t_i , the allocation α_{i+1} at the next rebalancing time is limited to the interval $B = [a - b, a + b] \cap [0, 1]$, where $b \in [0, 1]$. Parameter b indicates the maximum change in asset allocation allowed between rebalancing times.

A bandwidth limitation can be applied by using the same mechanism as in the previous section. A large penalty is subtracted from the expected utility at time t_{i+1} for allocation decisions outside interval B . No penalty is applied as long as $\alpha_{i+1} \in B$. Equation 5.1 is modified such that

$$V(z, t_i) = \sup_{\mathcal{A}_i} \mathbb{E} \left[U(Z_T) - \lambda(z) \mathbb{1}_{\{|\alpha_i - \alpha_{i+1}| > b\}} | Z_{t_i} = z \right], \quad (5.3)$$

with $\lambda(z)$ as defined in (5.2) and ρ large, such as $\rho = 10^4$. Allocation decisions that lie outside interval B become unattractive due to the large penalty and are no longer chosen. A bandwidth limitation has been applied.

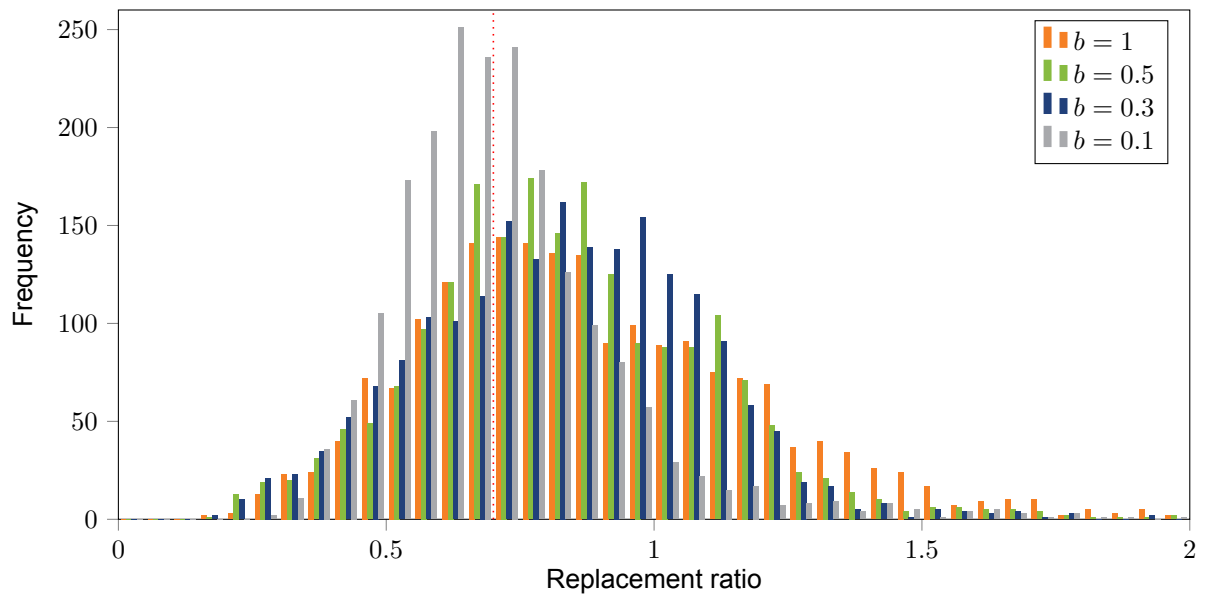


Figure 5.3: Replacement ratio distribution of dynamic programming strategy instances with different bandwidth limitations applied. The orange histogram for $b = 1$ is the same as the dark-blue dynamic programming histogram seen in figure 4.4.

5.2.1. Results

Figure 5.3 shows results of some dynamic programming strategy instances with different bandwidth limitations applied. Performance of the dynamic programming strategy lowers when the bandwidth is decreased. Statistics in Table 5.2 show that the asset allocation changes over time reduce when the bandwidth is decreased. Figure 5.4 shows that the allocation decision over time also becomes more stable when a bandwidth limitation is applied.

The allocation decision gets stuck around the initial allocation decision for a small bandwidth. This can be seen when $b = 0.1$. Larger bandwidths, for $b = 0.3$ or $b = 0.5$, improve the stability of the asset allocation over time and decrease the annual turnover. Just like with fictive transaction costs, a bandwidth restricts the investor. The solution space is now limited, causing the results in terms of shortage measure or various risk measures to deteriorate compared to the results found in Chapter 4. Figure 5.7 compares the results of the various methods of control smoothing to the results of rule-based strategies and the unrestricted dynamic programming strategy from Chapter 4.

5.3. Minimal improvement

Another option is to introduce a minimal improvement requirement to reduce annual turnover. It is used to encourage investors to remain with their current asset allocation. The goal is again to decrease annual turnover and improve the stability of the allocation decision over time.

If the expected utility improvement is below a threshold, investors remain at their current asset allocation.

This can also be implemented by using the same mechanism as used for the fictive transaction costs. To imply a minimal improvement, a reward is added. An allocation decision for which the asset allocation does not change, $|\alpha_i - \alpha_{i-1}| = 0$, receives an improvement of p percentage points. No penalties are applied to all other allocation decisions. Then, if the expected utility of a different asset allocation is higher, the improvement is at least $p\%$ compared to the current allocation.

Equation (5.1) is modified such that

$$V(z, t_i) = \sup_{\mathcal{A}_i} \mathbb{E} \left[U(Z_T) - \lambda(z) \mathbb{1}_{\{|\alpha_i - \alpha_{i-1}|=0\}} | Z_{t_i} = z \right], \quad (5.4)$$

with $\lambda(z)$ as defined in (5.2) and $\rho = \frac{p}{100}$, such that the expected utility of remaining at the current asset allocation improves by a percentage of p .

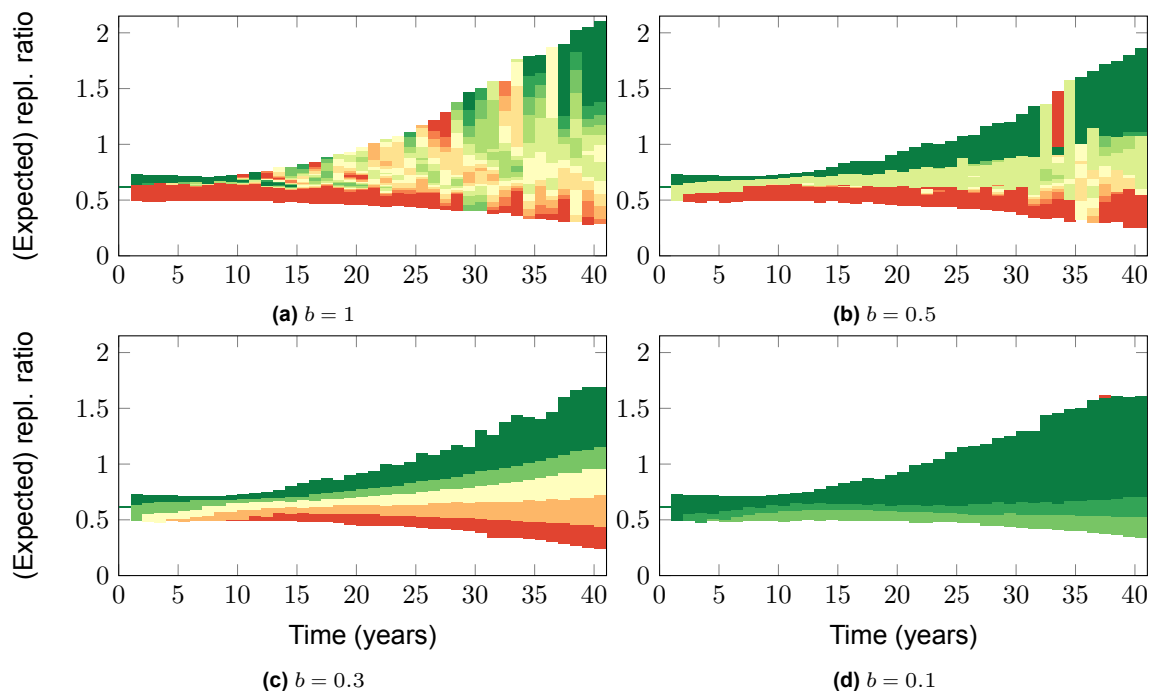


Figure 5.4: Asset allocation over time for different bandwidth limitations. Colours indicate the value of the return portfolio weight between 0 (green) and 1 (red), identical as in Figure 4.5. Figure 5.4a is the same plot as Figure 4.5.

5.3.1. Results

Figure 5.5 shows some results of dynamic programming strategy instances for different minimal improvement requirements. Table 5.3 shows that the stability of the allocation decisions improves when the minimal improvement requirement increases. This is also reflected in Figure 5.6.

Note that the average asset allocation jump increases for higher minimal improvement requirements. For a minimal improvement demand of 1%, the average asset allocation jump is more than 0.8. The investor is stuck with their initial allocation at first, as can be seen in Figure 5.6d. If the required improvement can be found, it often requires a large change in asset allocation. Smaller improvement requirements allow the investor to switch more gradually between asset allocations.

Minimal improvement requirements limit investors in their decisions. Results in terms of shortage measure and various risk measures are worse compared to the results from Chapter 4. Figure 5.7 compares the results of the various methods of control smoothing to the results of rule-based strategies and the unrestricted dynamic programming strategy from Chapter 4.

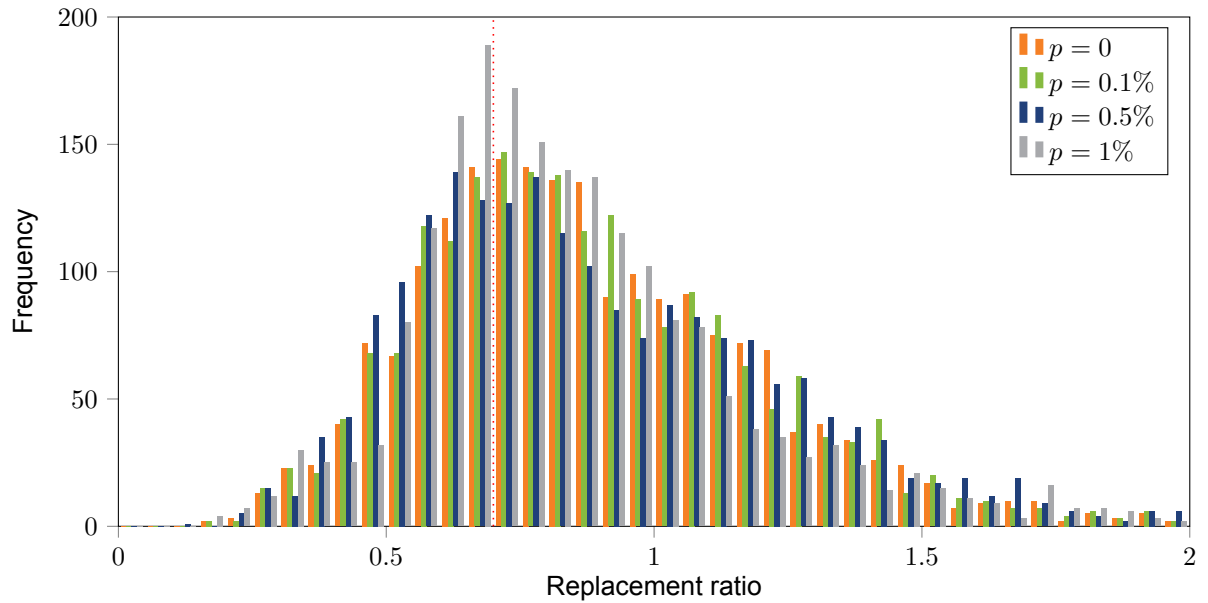


Figure 5.5: Replacement ratio distribution of dynamic programming strategy instances with different minimal improvement requirements applied. The orange histogram for $p = 0$ is the same as the dark-blue dynamic programming histogram seen in figure 4.4.

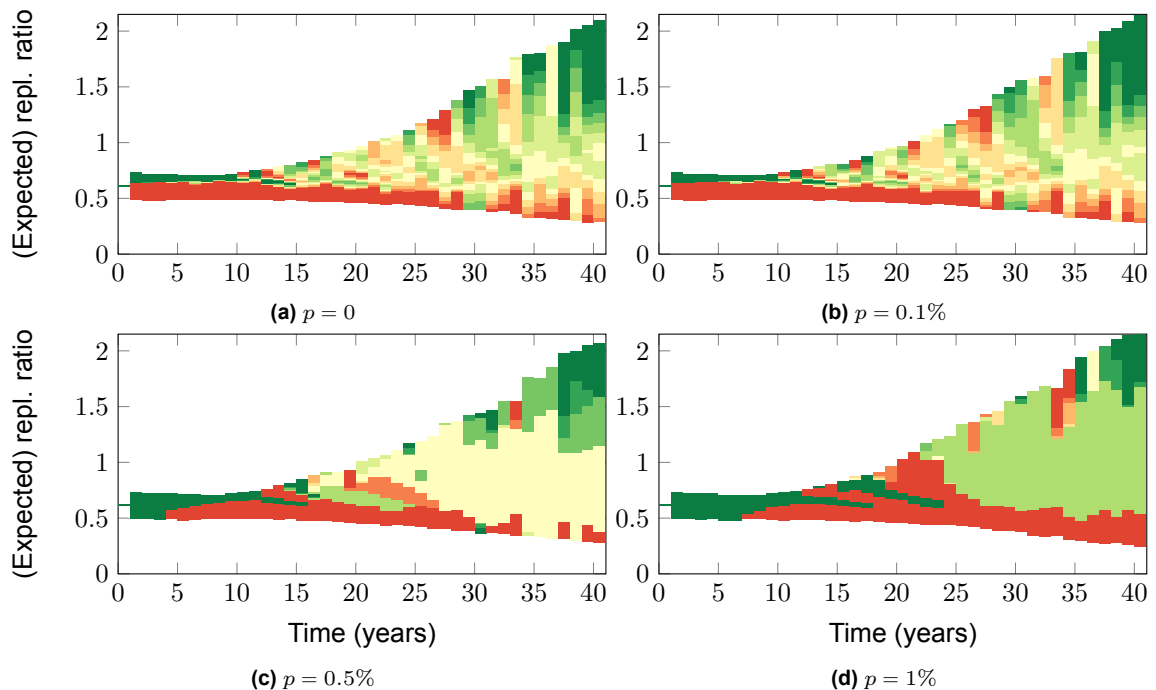


Figure 5.6: Asset allocation over time for different minimal improvement requirements. Colours indicate the value of the return portfolio weight between 0 (green) and 1 (red), identical as in Figure 4.5. Figure 5.6a is the same plot as Figure 4.5.

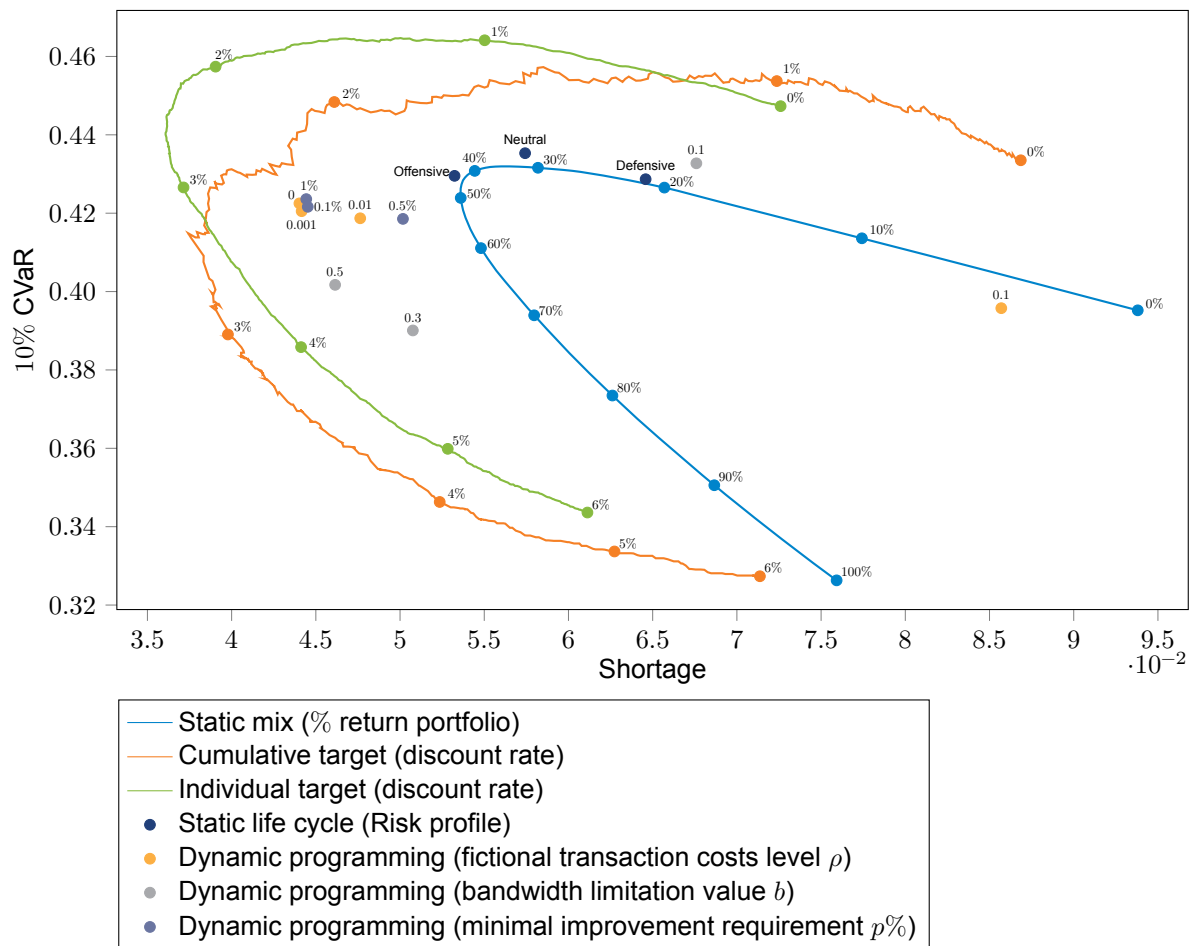


Figure 5.7: Risk measure comparison between static mixes, static life cycles, rule-based strategies, the dynamic programming strategy and various methods of control smoothing applied to the dynamic programming strategy.

5.4. Conclusion

This chapter showed various methods to stabilize the allocation decision over time. Three methods have been considered: The introduction of fictive transaction costs in Section 5.1, a bandwidth limitation in Section 5.2 and a minimal improvement requirement in Section 5.3. The stability of the asset allocation over time improves when one of these methods is applied, at the cost of optimality. Each method modifies the decision space in a different way.

Fictive transaction costs decrease annual turnover by introducing a penalty that increases linearly with the change in asset allocation. Results show that the stability of the allocation decision over time improves when fictive transaction costs are applied. Fictive transaction costs can be too high, causing investors to be stuck with their initial asset allocation.

A bandwidth limitation allows the investor to change freely between asset allocations in a limited interval. A downside of this approach is that big jumps are not allowed in this setting, even if they offer big improvements. Results show that a bandwidth limitation decreases the annual turnover and improves the stability of the allocation decision over time. Investors get stuck around the initial asset allocation when the bandwidth is too small.

Last, a minimal improvement requirement increases the expected utility of the current allocation, such that it becomes more attractive for investors to maintain the current asset allocation. This approach also decreases annual turnover and improves the asset allocation over time. A downside of this method is that when the required improvement has been found, it often comes with a large change in asset allocation. It is less likely that investors switch gradually between asset allocations.

Results of fictive transaction costs or bandwidth limitation mechanisms are more practical. Investors can still move gradually between different asset allocations in these settings, while big jumps in asset allocation are often found for the minimal improvement method.

All three methods restrict the solution space compared to the unrestricted problem solved in Chapter 4. Results found in this chapter are therefore less optimal compared to the results of Chapter 4.

To summarize,

- All three methods can be used to decrease annual turnover and improve the stability of the asset allocation over time.
- Methods that use fictive transaction costs or apply bandwidth limitations are more practical compared to methods that use minimal improvement requirements.
- Smoothing methods restrict the solution space: less optimal solutions are found compared to the unrestricted case shown in Chapter 4

5.A. Strategy results

	Static mix	Dynamic Programming			
	49%	$\alpha_t \in \{0, 0.1, \dots, 1\}$			
		$\rho = 0$	$\rho = 0.001$	$\rho = 0.01$	$\rho = 0.1$
Averages (RR)					
Mean	0.838	0.894	0.895	0.870	0.728
10% CVaR	0.425	0.423	0.421	0.419	0.396
5% CVaR	0.381	0.363	0.359	0.358	0.364
Percentiles (RR)					
Median	0.778	0.836	0.840	0.824	0.668
10% VaR	0.498	0.518	0.519	0.514	0.448
5% VaR	0.438	0.447	0.446	0.433	0.404
Goal (70% RR)					
Shortage	0.054	0.044	0.044	0.048	0.086
Goal reached (%)	0.626	0.696	0.699	0.678	0.437
Estim. error (RR)					
Mean	0.063	0.073	0.072	0.069	0.026
Std dev.	0.182	0.189	0.189	0.184	0.118
Asset allocation ($\Delta\alpha_t$)					
Average weight change	0	0.175	0.164	0.065	0.009
Average jump	0	0.283	0.272	0.196	0.124
Asset allocation(% of adjusting times)					
Asset weight changes	0	0.618	0.604	0.333	0.076
Maximum asset weight changes	0	0.029	0.026	0.003	0.000

Table 5.1: Simulation statistics for dynamic programming strategies with a replacement ratio target and different levels of fictive transaction costs.

	Static mix	Dynamic Programming			
	49%	$\alpha_t \in \{0, 0.1, \dots, 1\}$			
		$b = 1$	$b = 0.5$	$b = 0.3$	$b = 0.1$
Averages (RR)					
Mean	0.838	0.894	0.841	0.833	0.718
10% CVaR	0.425	0.423	0.402	0.390	0.433
5% CVaR	0.381	0.363	0.333	0.329	0.394
Percentiles (RR)					
Median	0.778	0.836	0.818	0.833	0.687
10% VaR	0.498	0.518	0.518	0.492	0.494
5% VaR	0.438	0.447	0.421	0.410	0.444
Goal (70% RR)					
Shortage	0.054	0.044	0.046	0.051	0.068
Goal reached (%)	0.626	0.696	0.682	0.695	0.464
Estim. error (RR)					
Mean	0.063	0.073	0.063	0.064	0.016
Std dev.	0.182	0.189	0.172	0.175	0.108
Asset allocation ($\Delta\alpha_t$)					
Average weight change	0	0.175	0.096	0.066	0.019
Average jump	0	0.283	0.314	0.264	0.101
Asset allocation(% of adjusting times)					
Asset weight changes	0	0.618	0.306	0.249	0.185
Maximum asset weight changes	0	0.029	0.000	0.000	0.000

Table 5.2: Simulation statistics for dynamic programming strategies with a replacement ratio target and different bandwidth limitations applied.

	Static mix 49%	Dynamic Programming $\alpha_t \in \{0, 0.1, \dots, 1\}$			
		$p = 0$	$p = 0.1\%$	$p = 0.5\%$	$p = 1\%$
Averages (RR)					
Mean	0.838	0.894	0.895	0.894	0.862
10% CVaR	0.425	0.423	0.422	0.419	0.424
5% CVaR	0.381	0.363	0.360	0.362	0.342
Percentiles (RR)					
Median	0.778	0.836	0.838	0.823	0.799
10% VaR	0.498	0.518	0.522	0.504	0.542
5% VaR	0.438	0.447	0.446	0.441	0.448
Goal (70%RR)					
Shortage	0.054	0.044	0.045	0.050	0.044
Goal reached (%)	0.626	0.696	0.696	0.661	0.659
Estim. error (RR)					
Mean	0.063	0.073	0.074	0.072	0.056
Std dev.	0.182	0.189	0.191	0.195	0.167
Asset allocation ($\Delta\alpha_t$)					
Average weight change	0	0.175	0.176	0.111	0.115
Average jump	0	0.283	0.294	0.587	0.814
Asset allocation(% of adjusting times)					
Asset weight changes	0	0.618	0.598	0.189	0.142
Maximum asset weight changes	0	0.029	0.035	0.055	0.047

Table 5.3: Simulation statistics for dynamic programming strategies with a replacement ratio target and different minimal improvement requirements.

6

Combination strategy

Chapter 4 shows that dynamic programming can be used in a simple setting. The pension investor setting which has been used in Chapter 3 has too much complexities to find an acceptable dynamic programming solution. This chapter combines the results from Chapters 3 and 4 into a strategy that uses dynamic programming to make investment decisions for each individual contribution the investor has made to the pension savings. A similar investment strategy has not been found in existing literature.

6.1. Combination strategy

The combination strategy is evaluated in the same setting as is used in Chapter 3. The individual wealth targets as defined in Equation (3.5) are also used for the combination strategy. The investment decision of the individual target strategy for a contribution c_k at time t can be summarized as follows:

$$\alpha_t = \begin{cases} 0, & \text{if } \exists \tau \leq t \text{ s.t. } W_\tau^{(k)} \geq W_k^*(\tau), \\ 1, & \text{otherwise} \end{cases},$$

where $W_\tau^{(k)}$ is the portfolio wealth at time τ of sub-problem k as stated in Section 3.2.

The 0 – 1 decision is one of the main disadvantages of the previously shown rule-based strategies. The combination strategy improves the allocation decision by using dynamic programming results from chapter 4.

Chapter 4 applies dynamic programming to an investor that has an initial wealth and a fixed wealth target over the time horizon. No contributions are made to the investment during the investment period. The individual target strategy from Chapter 3 models the pension investor as n separate investment problems that each have an initial wealth (the contribution made at time t) and a wealth target (the individual wealth targets).

Individual wealth targets defined in Chapter 3 are not constant over the time horizon. They are defined using expected inflation and use the market value factor at time t . The individual wealth targets are also not constant between scenarios. They are dependent on the height of the contribution. The dynamic programming approach requires a fixed wealth target in order to evaluate the expected utility. This is solved by changing the state variable of the dynamic programming algorithm. Instead of choosing $Z_t = W_t$ as in Chapter 4, Z_t is the ratio between the portfolio wealth and the wealth target,

$$Z_t = \frac{W_t}{W^*(t)}.$$

The target state is constant in time: $Z_t^* = 1$.

An advantage of choosing this state variable is that, theoretically, the dynamic programming solution only has to be computed once. The investment decisions for the first contribution of the investor can be used for all other contributions. Dynamic programming results for the first contribution span the time horizon $[t_0, T]$. At each rebalancing time, the optimal investment decision depending on state Z_t has

already been computed. Because a ratio is used, these investment decisions can also be used for the contributions in later years.

This is not the case in practice. Algorithm 4.1 does not fully convergence due to the limited sample size and due to sets of observations changing multiple times per iteration. Running algorithm 4.1 separately for each premium gives better decision rules. Separate runs of algorithm provide a better approximation of the optimal solution on average.

The asset allocation for the combination strategy has more degrees of freedom compared to the dynamic programming strategy of Chapter 4. The asset allocation of the overall portfolio is no longer bound to the discretization. Instead, the asset allocation of the overall portfolio is the sum of the asset allocations for each individual contribution.




Figure 6.1: Risk measure comparison between static mixes, static life cycles, rule-based strategies and combination strategies.

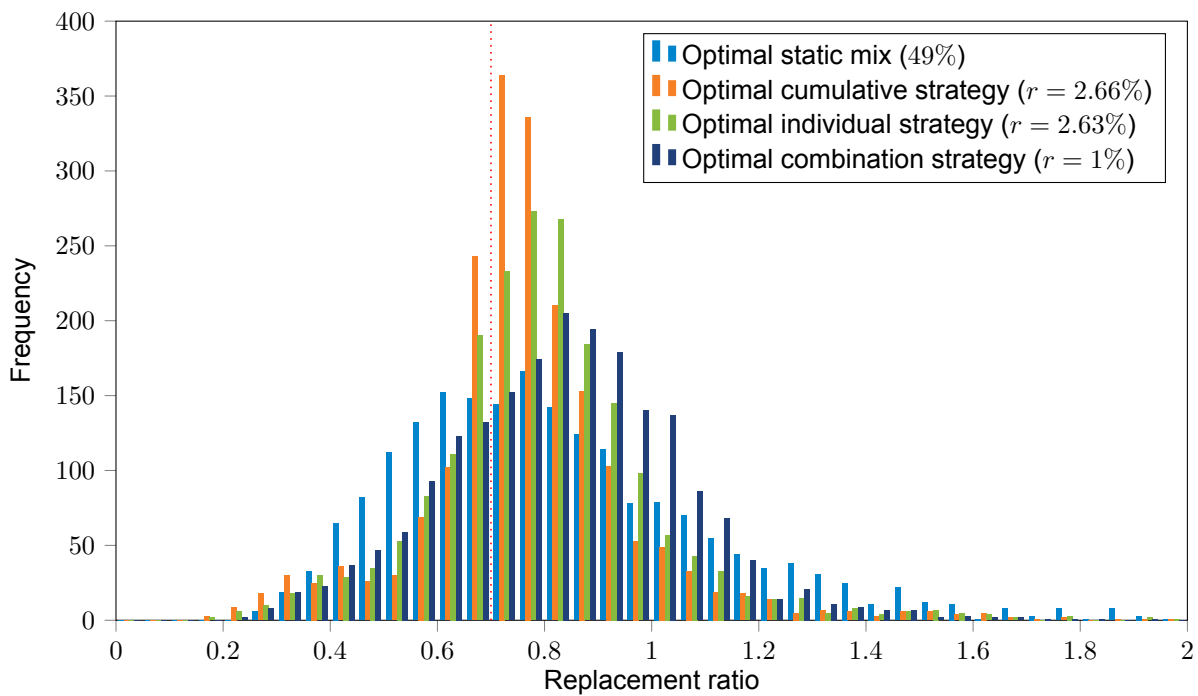


Figure 6.2: Replacement ratio distribution for the optimal static mix, optimal results from Chapter 3 and the combination strategy. The optimal instance of a strategy is the instance with the lowest shortage measure.

6.2. Results

Figure 6.2 shows the distribution of the replacement ratio for the combination strategy. Results from Chapter 3 have been added for comparison. Table 6.1 also compares results of the combination strategy to results from previous chapters.

The asset allocation over time is more stable for the combination strategy compared to a dynamic programming strategy. A dampening effect can be observed when combining the asset allocations of dynamic programming strategies for different contributions.

Figure 6.1 compares the shortage measure and 10% CVaR of various instances of the combination strategy. Although the discount rate interval is the same as used for rule-based strategies, the results of the combination strategy are much more clustered together.

Compared to rule-based strategies, the combination strategy assigns more wealth to the matching portfolio for high discount rates.

Local regressions used in the dynamic programming algorithm are limited to the sample. For the combination strategy, the sample at time t is the set of ratios between portfolio wealth W_t and wealth target $W^*(t)$. Wealth target $W^*(t)$ increases for a higher discount rate, causing the ratio between portfolio wealth and wealth target to decrease. The sample range becomes smaller. Therefore, regressions for high discount rates are more limited compared to regressions for low discount rates.

Dynamic programming sub-problem results always have a large matching portfolio allocation for the scenarios with high state variables and a low matching portfolio allocation for the scenarios with a low state variable, regardless of the state variable range. Allocation decisions at time t for the same value of state variable Z_t may therefore differ for different discount rates. In general, the matching portfolio allocation increases for the same value of state variable Z_t increases for a higher discount rate.

6.3. Academic model

Figure 6.3 shows the results of the combination strategy for different discount rates. Figure 6.4 compares the distribution of outcomes of optimal rule-based strategies compared to the combination strategy. Results are comparable to the results of Section 6.2 and are therefore not discussed in detail.




Figure 6.3: Risk measure comparison between static mixes, static life cycles, rule-based strategies and combination strategies using scenarios generated by the academic model.

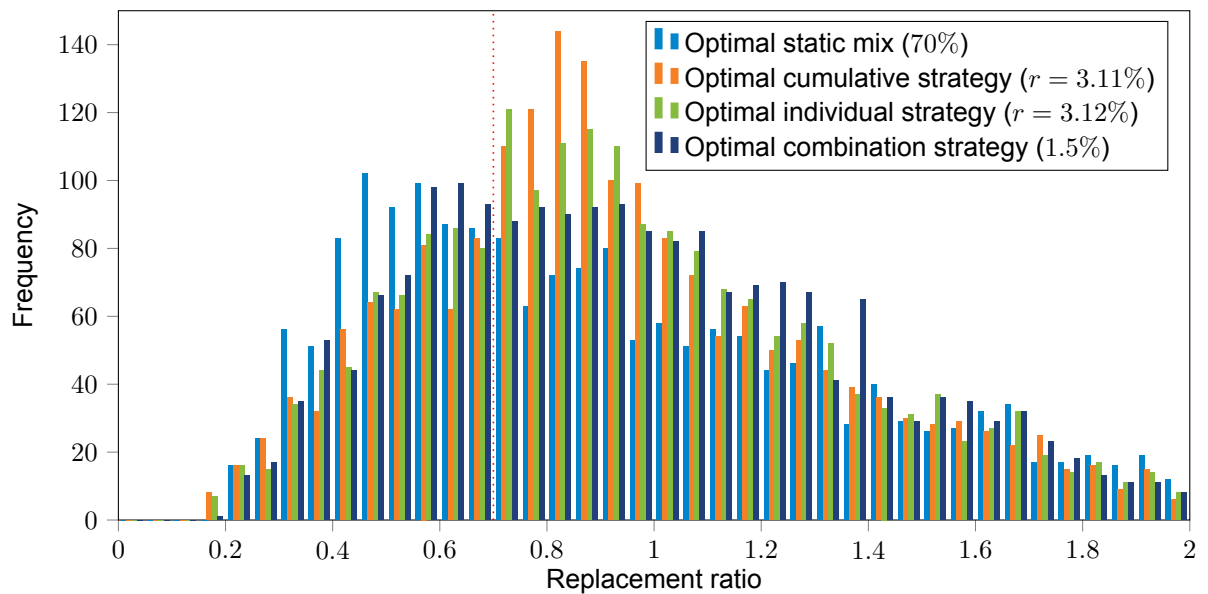


Figure 6.4: Replacement ratio distribution for the optimal static mix, optimal results from Chapter 3 and the optimal combination strategy using scenarios generated by the academic model. The optimal instance of a strategy is the instance with the lowest shortage measure.

6.4. Conclusion

This chapter combines conclusions from Chapter 3 and 4 into a combination strategy that solves some of the restrictions found in Chapters 3 and 4.

Chapter 3 concludes that investment results improve when contributions are invested individually. The combination strategy also looks at contributions individually.

Chapter 4 concludes that dynamic programming can be used for a setting with an initial wealth and a fixed wealth target. The combination strategy uses each contribution as initial wealth and uses the ratio between portfolio wealth and wealth target to create a fixed wealth target.

The combination strategy does not have the disadvantage of a 0-1 decision that the rule-based strategies from Chapter 3 have. A rule-based strategy only switches to the matching portfolio once the target value has been reached. If no targets are reached, the investor remains with a risky allocation of only the return portfolio. The combination strategy can switch to the matching portfolio before the target is reached.

Chapter 4 also applies dynamic programming to the pension case. A required change in state variable made it difficult to find the relationship between asset allocation and change in state variable. The combination strategy uses a more simple state variable. The relationship between the asset allocation and change in state variable is better observable for the dynamic programming algorithm when using a portfolio wealth based state variable.

Another restriction of the dynamic programming strategy of Chapter 4 is the unstable asset allocation over time. The asset allocation over time of the combination strategy is much more smooth. Wild asset allocations of each individual contribution cancel out once they are combined in the overall portfolio.

The combination strategy does not use the information that a target pension income can be secured until the end of the investment horizon by keeping a bond until final time T . The rule-based strategies from Chapter 3 do make use of this mechanism. The combination strategy could be further improved by making use of this mechanism.

To summarize,

- Compared to rule-based strategies, the combination strategy does not have the disadvantage of a 0 – 1 decision.
- Compared to the dynamic programming strategy, the combination strategy has improved results and the stability of the asset allocation over time has increased significantly.
- The combination strategy can be improved by securing pension income in the matching portfolio, just like rule-based strategies.

	Static mix	Rule Based		Dyn. Prog.	Combination
	49%	Cum. Target	Ind. Target	$\alpha_t \in [0, 0.1, \dots, 1]$	1%
Averages (<i>RR</i>)					
Mean	0.838	0.776	0.801	0.894	0.831
10% CVaR	0.425	0.415	0.440	0.423	0.448
5% CVaR	0.381	0.328	0.363	0.363	0.385
Percentiles (<i>RR</i>)					
Median	0.778	0.756	0.787	0.836	0.834
10% VaR	0.498	0.568	0.564	0.518	0.554
5% VaR	0.438	0.423	0.459	0.447	0.461
Goal (70% <i>RR</i>)					
Shortage	0.054	0.037	0.036	0.044	0.037
Goal reached (%)	0.626	0.705	0.717	0.696	0.729
Estim. error (<i>RR</i>)					
Mean	0.063	0.011	0.021	0.073	0.112
Std dev.	0.182	0.118	0.115	0.189	0.139
Asset allocation ($\Delta\alpha_t$)					
Average weight change	0	0.034	0.046	0.175	0.072
Average jump	0	0.181	0.071	0.283	0.077
Asset allocation(% of adjusting times)					
Asset weight changes	0	0.188	0.653	0.618	0.934
Maximum asset weight changes	0	0.022	0.004	0.029	0.000

Table 6.1: Simulation statistics for the optimal combination strategy with a replacement ratio target of 70%.

	Static mix	Rule Based		Combination
	70%	Cum. Target	Ind. Target	1.5%
Averages (<i>RR</i>)				
Mean	1.088	0.972	0.981	0.983
10% CVaR	0.350	0.362	0.370	0.378
5% CVaR	0.300	0.293	0.304	0.319
Percentiles (<i>RR</i>)				
Median	0.904	0.885	0.905	0.925
10% VaR	0.433	0.472	0.478	0.479
5% VaR	0.358	0.378	0.386	0.383
Goal (70% <i>RR</i>)				
Shortage	0.069	0.051	0.051	0.052
Goal reached	0.652	0.738	0.728	0.705
Estim. error (<i>RR</i>)				
Mean	0.141	0.061	0.063	0.150
Std dev.	0.349	0.197	0.194	0.212

Table 6.2: Simulation statistics for the optimal combination strategy with a replacement ratio target of 70% using scenarios generated by the academic model.

7

Conclusion

This chapter contains the conclusions of this thesis and offers recommendations for future research.

7.1. Conclusion

This thesis showed different dynamic investment strategies for an individual DC pension investor. The goal of each of these investment strategies is to improve the target-reachability of the investor, while at least maintaining the same level of downside risk. Investment decisions that follow from the dynamic investment strategies should be applicable in practice.

In Chapter 3, two rule-based strategies have been explained. Both strategies define target pension incomes for each year in the investment horizon. For each year, a bond price that secures the target pension income is computed and compared to the portfolio value. The investor buys this bond if they have sufficient wealth.

Investment decisions for the rule-based strategies are applicable in practice. The investor buys return portfolio shares, unless they can buy the bond that secures the target pension income. A downside of the rule-based strategies is that investors maintain a risky allocation of only return portfolio shares if they are not able to buy the bonds.

Rule-based strategies offer an improvement in reachability of the target and have less downside risk compared to static allocations or static life cycles. Rule-based strategies that investment annual contributions independent of each other have better results compared to rule-based strategies that make investment decisions based on the complete portfolio.

In Chapter 4 dynamic programming is used to determine the optimal investment decision for every year in the investment horizon. Dynamic programming showed potential in a simple setting with an initial wealth and fixed wealth target. It is difficult to use dynamic programming in the setting of the DC pension investor, because the target of the investor is no longer fixed. It is also dependent on salary and inflation, making it more difficult to observe the relationship between investment decisions and investment progress.

Investment decisions found by using dynamic programming are not always applicable in practice. The investment decision is unstable and annual turnover can be large. Methods to stabilize the investment decision and to decrease the annual turnover are shown in Chapter 5.

Dynamic programming strategies offer improvements in reachability of the target and have similar downside risk compared to fixed allocations or static life cycles. Improvements are less compared to the improvements seen by rule-based strategies. Rule-based strategies secure pension incomes by keeping bonds until the end of the investment horizon, while the dynamic programming strategy does not use this information.

Investment results of dynamic programming strategies become worse after methods are applied that attempt to stabilize the investment decision. These methods restrict the solution space. Investment results found are therefore less optimal compared to the unrestricted case of Chapter 4.

Conclusions from rule-based strategies and dynamic programming have been combined in the combination strategy in Chapter 6. Investment results improve when contributions are invested separately. Dynamic programming can be used in a simple setting with an initial wealth and a fixed wealth target.

The combination strategy splits the DC pension investor problem into sub-problems for each year in the investment horizon. Each sub-problem has an initial wealth (the annual contribution of that year) and a wealth target (the target pension income as defined for rule-based strategies). Dynamic programming is used to find the optimal investment decisions for each sub-problem.

The investment decisions of the combination strategy are applicable in practice. Investment decisions for separate contributions can be unstable and sometimes have a large annual turnover, but this is not the case for the overall portfolio. When the investment decisions for the separate contributions are combined, a dampening effect is observed. The stability of investment decisions of the combination strategy is similar to the stability of the investment decisions of rule-based strategies.

The combination strategy offers an improvement in reachability of the target and has less downside risk compared to static allocations or static life cycles. Improvements are comparable to the best rule-based strategies, although it does not use the information that pension income can be secured by keeping a bond until the end of the investment horizon.

7.2. Recommendations

Three different types of dynamic investment strategies have been studied in this thesis. This section offers some areas of improvement for each of them.

First, rule-based strategies in Chapter 3 showed promising investment results. However, a disadvantage of these strategies is that the investor maintains a risky allocation of only return portfolio shares when they are unable to reach the target values. In such cases, additional rules could be implemented that do assign some wealth to the less risky matching portfolio. For example, such rules could only activate in the last 10 years of the investment horizon if no wealth has been assigned to the matching portfolio before this time. This way, investors are able to add some security to their pension savings during the final years of the investment period.

Second, dynamic programming has been studied in Chapter 4. It is difficult to implement dynamic programming for the DC pension investor setting. Investment decisions can be unstable and annual turnover can be large. Chapter 5 showed methods that can be applied to stabilize investment decisions, but this comes at the cost of investment results.

The stability of the investment decision can also be improved by regressing on a fixed mesh of state variables, instead of using the state variables provided by each scenario. This approach is costly in terms of computational time. Each point in the mesh needs to be evaluated for every scenario, for every year and for every available investment decision. In each of these computations, the scenario then needs to be evaluated until final time in order to compute the expected utility. This evaluation is costly for the DC pension investor setting.

The advantage of a fixed mesh of state variables is stability in the local regressions, which determine the investment decision.

Another area of improvement for the dynamic programming strategy is adding the option to secure a pension income in the matching portfolio, just like rule-based strategies do.

Third, dynamic programming is applied to each contribution separately in the combination strategy. The combination strategy combines the some of the best properties of rule-based strategies and dynamic programming. Just like the dynamic programming strategy, it can be improved by regressing on a fixed mesh of state variables. The combination strategy can be further improved by also securing a pension income in the matching portfolio.

Bibliography

- [1] Achmea. Life cycle gemiddeld, 2019. <https://www.centraalbeheer.nl/inkomen-en-pensioen/voor-later/Paginas/life-cycle-gemiddeld.aspx>.
- [2] Richard Bellman. *Dynamic Programming*. Dover Publications, 1957. ISBN 9780486428093. <https://www.bibsonomy.org/bibtex/29cdd821222218ded252c8ba5cd712666/pcbouman>.
- [3] J. C. Bogle. *Bogle on Mutual Funds: New Perspectives for the Intelligent Investor*. McGraw-Hill, 1993.
- [4] John Y Campbell and Luis Viceira. The term structure of the risk-return tradeoff. Working Paper 11119, National Bureau of Economic Research, February 2005. <http://www.nber.org/papers/w11119>.
- [5] Aegon cappital. Wat betekent het beleggingsprofiel voor uw pensioen?, 2019. <https://www.aegoncappital.nl/werknemer/beleggen-voor-pensioen/beleggingsprofiel-pensioen>.
- [6] William S. Cleveland. Robust locally weighted regression and smoothing scatterplots. *Journal of the American Statistical Association*, 74(368):829–836, 1979. doi: 10.1080/01621459.1979.10481038. <https://www.tandfonline.com/doi/abs/10.1080/01621459.1979.10481038>.
- [7] William S. Cleveland and E. Grosse. Computational methods for local regression. *Statistics and Computing*, 1(1):47–62, Sep 1991. ISSN 1573-1375. doi: 10.1007/BF01890836. <https://doi.org/10.1007/BF01890836>.
- [8] F. Cong and C.W. Oosterlee. Multi-period mean–variance portfolio optimization based on monte-carlo simulation. *Journal of Economic Dynamics and Control*, 64:23 – 38, 2016. ISSN 0165-1889. doi: <https://doi.org/10.1016/j.jedc.2016.01.001>. <http://www.sciencedirect.com/science/article/pii/S0165188916000026>.
- [9] Francis X. Diebold and Glenn D. Rudebusch. *Yield Curve Modeling and Forecasting: The Dynamic Nelson-Siegel Approach*. Princeton University Press, 2013. ISBN 9780691146805. URL <http://www.jstor.org/stable/j.ctt1r2dc4>.
- [10] Ortec Finance. Ortec finance scenario approach, 2018. <https://www.ortecfinance.com/en/insights/research/ortec-finance-scenario-approach>.
- [11] Peter Forsyth and Kenneth R. Vetzal. Optimal asset allocation for retirement saving: Deterministic vs. time consistent adaptive strategies. October 2018. <https://ssrn.com/abstract=2991828>.
- [12] Francis A. Longstaff and Eduardo S. Schwartz. Valuing american options by simulation: A simple least-squares approach. *Review of Financial Studies*, pages 113–147, 2001.
- [13] Robert Merton. Lifetime portfolio selection under uncertainty: The continuous-time case. *The Review of Economics and Statistics*, 51(3):247–57, 1969.
- [14] Petra Molenaar-Cox and Dennis Woestenburg. Pensioenaansprakenstatistiek 2015. September 2018. <https://www.cbs.nl/nl-nl/achtergrond/2018/39/pensioenaansprakenstatistiek-2015>.
- [15] Nationale Nederlanden. Lifecycle-methode, 2019. <https://www.nn.nl/Particulier/Beleggen/Beheerd-Beleggen/Lifecyclemethode.htm>.

- [16] Svetlozar T Rachev, Christian Menn, and Frank J Fabozzi. *Fat-tailed and skewed asset return distributions: implications for risk management, portfolio selection, and option pricing*, volume 139. John Wiley & Sons, 2005.
- [17] Rijksoverheid. Vertrouwen in de toekomst, 2017. <https://www.rijksoverheid.nl/documenten/publicaties/2017/10/10/regeerakkoord-2017-vertrouwen-in-de-toekomst>.
- [18] H. Steehouwer. *Macroeconomic Scenarios and Reality: A Frequency Domain Approach for Analyzing Historical Time Series and Generating Scenarios for the Future*. PhD thesis, Vrije Universiteit Amsterdam, 2005. Naam instelling promotie: VU Vrije Universiteit Naam instelling onderzoek: VU Vrije Universiteit.
- [19] Jules H. van Binsbergen and Michael W. Brandt. Solving dynamic portfolio choice problems by recursing on optimized portfolio weights or on the value function? *Computational Economics*, 29(3):355–367, May 2007. ISSN 1572-9974. doi: 10.1007/s10614-006-9073-z. <https://doi.org/10.1007/s10614-006-9073-z>.
- [20] Martin van der Schans and Hens Steehouwer. A frequency domain impulse response approach. 2012.
- [21] Ministerie van Financiën. Loonheffingen, inkomstenbelasting. pensioenen; beschikbare-premieregelingen en premie- en kapitaalovereenkomsten en nettopensioenregelingen (staffelbesluit pensioenen. *Staatscourant van het Koninkrijk der Nederlanden*, 68652, 2018.
- [22] J.M. Peeters Weem. Mean-variance optimization for life-cycle pension portfolios, 2017. <http://resolver.tudelft.nl/uuid:29181e27-003c-4ffa-9c77-aeb53783ad91>.
- [23] Rongju Zhang, Nicolas Langren, Yu Tian, Zili Zhu, Fima Klebaner, and Kais Hamza. Sharp target range strategies for multiperiod portfolio choice by decensored least squares monte carlo. April 2017.