

Goal priorities in multi-goal based planning

by

Femke J. Hermse

to obtain the degree of Master of Science
at the Delft University of Technology,
to be defended publicly on Wednesday September 16, 2020 at 15:00.

Student number: 4928210
Project duration: September 2, 2019 – September 16, 2020
Thesis committee: Prof. dr. ir. C.W. Oosterlee, TU Delft, supervisor
Dr. ir. J.T. van Essen, TU Delft
Dr. M. van der Schans, Ortec Finance, daily supervisor

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



Abstract

In this thesis we build a goal based planning framework that takes into account goal priorities. Goal based planning is a type of personal wealth planning, with a focus on the feasibility of an investor's goals. Currently, to take goal priorities into account a dynamic asset allocation is created. We focus on incorporating goal priorities by ensuring goal cash flows reflect investors' preferences, no matter the asset allocation. No money should be spent on less important goals if it jeopardises achieving more important goals. With such a goal cash flow strategy, an investor gains valuable insight into the feasibility of her goals. We build such a framework in two ways. First, we use utility theory. With utility theory, we can quantify an investor's satisfaction with the goal cash flows. For that purpose, we create a total utility function that takes into account goal priorities. Then, we optimise over the goal cash flows such that it results in the highest expected total utility. Second, we use conditional expectations. With conditional expectations, we determine the relationship between earlier wealth and future goal achievability. Therefore, we create a framework with a spending rule based on the conditional expectation. We implement conditional goal cash flows. With conditional cash flows, we create a goal based planning framework that takes into account goal priorities for multi-goal investors.

Preface

This thesis has been submitted for the degree of Master of Science in Applied Mathematics from the Delft University of Technology. During the last twelve months, I worked at Ortec Finance. For three days a week, I worked on my thesis within the Goal Based Planning solution. The other two days a week, I first joined the Climate and ESG solution and then the Scrum team of Pension Strategy. With the end of my thesis and my time at Ortec Finance, I would like to say some thanks.

First of all, I would like to thank my daily supervisor Martin van der Schans from Ortec Finance. Without your help, this thesis would not have been the same. You have the incredible capability of explaining everything in the easiest of ways. After our meetings, I always felt assured and completely capable of tackling the next chapter of my thesis. Additionally, I would like to thank professor Kees Oosterlee. From the start, you were part of this project and with your help I got the opportunity to work and complete my thesis at Ortec Finance. During our meetings, you provided valuable feedback and interesting new avenues to pursue. Lastly, I would like to thank Theresia van Essen for being part of my thesis committee.

Second of all, I would like to thank everyone at Ortec Finance. In particular, I would like to thank three teams. Firstly, a thank you to the GBP team. Even with all the Belgian jokes (looking at you Arnoud), you made this year and this thesis more fun. Secondly, a thank you to the Scrum team. The never-ending bugs in the code would have been more annoying if you were not part of the process. Lastly, a thank you to the CES team. On my first day at Ortec, I was immediately taken in by the chaos and energy that is this team. From day one, you gave me the opportunity and responsibility to create something of value.

Lastly, I would like to thank my friends and family. All of you have been part of and responsible for the highlights of this year, either online or in person. I especially want to thank my parents and my sisters. You always had time for me, whether to listen to my complaints, cheer me on when things were going well, or distract me when I needed a break.

*Femke J. Hermse
Rotterdam, September 2020*

Contents

1	Introduction	1
1.1	Goal based planning framework	2
1.2	Outline	3
2	Total utility function	5
2.1	Expected utility theory	5
2.2	Cumulative prospect theory	7
2.3	Utility functions in goal based planning	8
2.3.1	Step utility	8
2.3.2	CRRA utility	9
2.3.3	Prospect utility	10
2.4	Expected total utility	11
2.4.1	Marginal utility function	11
2.4.2	Total utility function	12
2.5	Conclusion	12
3	Utility theory framework	13
3.1	Upper bound on the first goal's cash flows	15
3.2	Maximum utility barrier	16
3.3	Conclusion	18
4	Conditional expectation	21
4.1	Regressed conditional expectation	22
4.1.1	Creation of datasets	22
4.1.2	Fit of regression algorithms	23
4.1.3	Robustness of lasso spline regression	27
4.2	Algorithm for lasso spline regression	28
4.3	Conclusion	30
5	Conditional expectation framework	31
5.1	Conditional expectation barrier	32
5.2	Barrier comparison	34
5.3	Early-exercise options	37
5.3.1	Bermudan options	37
5.3.2	Swing options	39
5.3.3	Relation early-exercise options and goal based planning	41
5.4	Multi-goal investors	42
5.5	Conclusion	45
6	Conclusion and recommendations	47
6.1	Conclusions	47
6.2	Recommendations for future research	48
6.3	Recommendations for Ortec Finance	48
A	KNW Scenarios	51
B	Regression Algorithms	53
B.1	Ordinary Least Squares Regression	53
B.2	Penalized Ordinary Least Squares Regression	54
B.2.1	Ridge Regression	54
B.2.2	Lasso Regression	55
B.2.3	Regularization Parameter	55

B.3 Support Vector Regression	55
B.4 Spline Regression	58
Bibliography	61

Introduction

In this thesis, we develop a goal based planning framework that takes into account goal priorities. Goal based planning is a type of personal wealth planning focused on an investor's goals. An investor determines, potentially with a financial advisor, her goals, the timings, and the amounts. In some goal based planning frameworks the investor also prioritises the individual goals. For instance, an investor might want to save up 10K for a car in two years and 50K for a more important retirement goal in twenty years. In goal based planning there are two approaches: holistic and bucketing. In the holistic approach, an investor's goals are considered as one entity. A single portfolio is created to determine the feasibility of an investor's goals. In the bucketing approach, also referred to as mental accounting, an investor's goals are considered separately. Per goal, a specific portfolio is created to reach that goal. An investor then divides her initial wealth and future deposits over the portfolios and determines the goals' feasibility separately. Both approaches have their advantages [10], but we take a holistic approach to goal planning. Based on the investor's risk profile, a model portfolio is made up of different financial products. The model portfolio follows one of three investment strategies: buy-and-hold, rebalancing, or life-cycle/dynamic. Uncertainty is incorporated into the modelling of the portfolio to reflect the missing information on the future, often in scenarios or scenario trees to emulate different futures. These scenarios contain diverse return values, modelling different future economic behaviours of the portfolio. In each scenario, the deposits and goal cash flows are simulated. Then, over all scenarios, the feasibility of an investor's goals are determined. In this thesis, we require three investor determined parameters: initial wealth, future deposits, and prioritised goals. We require an investor to prioritise her goals as the aim of this thesis is to create a goal based planning framework that focuses on goal priorities.

In literature, there are examples of holistic goal based planning frameworks that incorporate goal priorities. Dempster and Medova [6] create per goal a utility function that translates a goal's cash flows to utility. The goal priority is incorporated in the slope of the function. Then they optimise over the expected utility of lifetime spending, which is a function of the individual goal functions. Fowler and de Vassal [10] forgo utility theory. They use a sequence of optimisations to maximise the probability of goal success, starting with the highest priority goal. The second optimisation maximises the probability of success of the second-highest priority goal. To take into account the most important goal, they place a minimum constraint on its probability of goal success. This process is repeated for all goals, adding an extra constraint for each iteration. Kim et al. [19] use multi-stage stochastic programming in combination with goal programming. They use stochastic programming to incorporate uncertainty in the framework. Multiple futures are presented in the form of a scenario tree. Goal programming falls under multi-objective optimisation. It is an extension on the linear programming framework and is capable of solving multiple objectives with conflicting constraints. They sequentially maximise the expected total spending over all scenarios, defining in each iteration the optimal cash flows for a particular goal. They start with the most important goal. In the next iteration, they again maximise total spending, while placing strict cash flow constraints on higher priority goals.

All of the previous frameworks incorporate goal priorities in function of portfolio selection. The result of the optimisation problems, no matter the objective function, is the most optimal (dynamic) asset allocation. In this thesis, we take a different approach to incorporate goal priorities. We do not optimise in function of asset allocation or any investor parameter. The goal is to create a framework in which all goal cash flows align with the investor's goal priorities, no matter the investment strategy. We want to prevent that current cash flows put future goal achievability in jeopardy, especially when more important goals occur later in time. We return to the investor described in the first paragraph, who has a car goal and a retirement goal. The retirement goal occurs later in time, but is more important than the car goal. Following the priorities, no money should be spent on a car if that prevents the investor from obtaining her retirement goal. We want to exclude the focus on short-term goals from the framework.

To place more emphasis on short-term objectives at the expense of long-term objectives is a common trait in people [21]. People often forgo long-term plans if that means they can obtain something right now. This occurs in investing as well, where investors focus on short-term goals instead of long-term goals. Research has shown that on a corporate level, this behaviour is not beneficial to the health of the company [2]. Following that logic, such behaviour is also not beneficial to individual investors. Therefore, we build a framework in which goal cash flows align with the investor's priorities. If the later goals are more important, the framework and goal cash flows should reflect that. In return, the goal based planning framework presents the investor with the feasibility of her goals, while incorporating a healthy investment and cash flow strategy.

The aim of this thesis is twofold:

- Determine the best way to build a framework that takes into account goal priorities in the context of goal cash flows.
- Verify whether a framework that takes into account goal priorities is an improvement compared to a framework which does not.

We explore several directions to develop a suitable framework. We cannot use the same techniques as Fowler and de Vassal [10] or Kim et al. [19]. The objective functions all incorporate the present value of the goals. With the present value, we cannot gain information on sequential goal cash flows, which we are interested in. Dempster and Medova use multi-period cash flows to calculate the expected utility. Following that logic, we create a utility function over time and evaluate a framework that optimises over the expected total utility. Additionally, we develop a framework that uses conditional expectations. With conditional expectations, we determine the relation between present and future. Initially, we focus on a two-goal portfolio, of which the second goal is more important. Later, we expand the framework to multiple goals.

1.1. Goal based planning framework

For the purpose of this thesis, we create a simple goal based planning framework in Python. The primary inputs for the framework are an investor's initial wealth, future deposits, and prioritised goals. Priorities are denoted with a number of stars out of five. As stated above, we take a holistic approach to goal based planning. We consider all goals as one and create a single portfolio to achieve them. To determine the feasibility of the goals, we use scenario analysis. These scenarios contain returns and portray different future behaviours of the portfolio and underlying asset classes. We do not optimise on the optimal allocation strategy. Instead, we work with portfolios consisting of a ratio of stocks to 10-year bonds. This allocation does not change over the investment horizon. We create these scenarios with the data provided by "de Nederlandsche Bank" (DNB). For more information on the scenario set, we refer to Appendix A. Per investor, we simulate 2000 scenarios using the client's initial wealth, deposits, and prioritised goals. Per goal, we determine the goal completion in each scenario and from those values, the goal achievability.

Definition 1.1 (Goal Completion). *Suppose we have a goal at time t of value v and a goal cash flow of c . Then the goal completion for goal t , denoted γ_t , is equal to*

$$\gamma_t = \frac{c}{v}.$$

Definition 1.2 (Goal Achievability). *Suppose we have a goal at time t with a value of G . For a number of scenarios ω , we determine the goal achievability as*

$$I(\omega, t) = \begin{cases} 1 & \text{if } W(\omega, t) \geq G \\ 0 & \text{otherwise} \end{cases}$$

for $\omega \in \{\omega_1, \dots, \omega_n\}$ with the indicator function and $I(\omega, t)$, the goal completion as defined in Definition 1.1 at time t in scenario ω .

The difference between a framework that does not take into account goal priorities and one that does is how goal cash flows are determined. For a framework that does not take into account goal priorities, all the money available is spent on each goal, limited to the specified goal amount. Clearly, the future is not taken into consideration; each goal is evaluated separately. This means that the framework potentially disregards the investor's priorities, especially in cases where more important goals occur later in time. In a framework that does take into account goal priorities, goal cash flows are determined by considering current and future goals.

1.2. Outline

In this thesis, we attempt to build a goal based planning framework that aligns goal cash flows with goal priorities. We explore two distinct ways to create such a framework: with utility theory and with conditional expectation. In Chapters 2 and 3 we use expected utility theory and cumulative prospect theory to create a priority-based framework. In Chapter 2, we develop a measure that quantifies an investor's satisfaction with the goal completions: the total utility function. We create the function using expected utility and prospect theory. The theories describe how an investor assesses a risky decision. In Chapter 3, we use several optimisation techniques with the expected total utility as objective function. In Chapters 4 and 5 we use conditional expectations to create a priority-based framework. In particular, we use the conditional expectation

$$E[I(\omega, t) | \mathcal{F}_t],$$

where $I(\omega, t)$ is the indicator function, $W(\omega, t)$, the wealth at time t and T respectively, and G the goal amount at time t for $t < T$. In Chapter 4, we create an efficient and robust heuristic to calculate the conditional expectation for varying investor parameters. In Chapter 5, we use the heuristic to create a priority-based framework. We start with two-goal portfolio and expand the framework to multi-goal portfolios. In Chapter 6, we conclude the thesis and give recommendations for future research.

2

Total utility function

The main purpose of this thesis is to create a goal based planning framework that takes into account goal priorities as described in Chapter 1. We modify the framework of Section 1.1 such that goal cash flows adhere to the goal priorities. Suppose an investor has two goals, of which the most important goal occurs later in time. In the framework, no money should be spent on the first goal if that puts achieving the second goal at risk. As a first attempt to build such a framework, we use optimisation to determine the optimal cash flows on the first goal. As an objective function, we need a measure which quantifies an investor's satisfaction with the goals' feasibility. This leads us to expected utility theory and cumulative prospect theory. Both theories describe how an investor assesses outcomes of a risky decision. Within the context of this thesis, we use them to create a total utility function. In goal based planning, the most important factor that defines success is goal completion as defined in Definition 1.1. As such, we use goal completion as the main input for the utility function. Additionally, we incorporate the relative preferences of goals to account for goal priorities in the framework, which are denoted with a number of stars out of five. Firstly, we introduce expected utility and cumulative prospect theory. Secondly, we discuss these theories within the context of goal based planning. Lastly, we introduce and elaborate on the total utility function specific to this thesis.

2.1. Expected utility theory

In this section, we introduce the components of expected utility theory (EUT) that are necessary for this thesis. Suppose an investor has two options to invest €1. One option is to invest the €1 and get back €2 guaranteed. The second option is to invest the €1 and then have a 50/50 percent chance of receiving €5 or losing €1. The classic theory for evaluating such investing decisions, where there is uncertainty concerning the return, is EUT. Simply put, the choice is between a guaranteed €1 or an expected €1 increase in wealth. This evaluation, however, does not take into account the true value of the potential €5. Depending on the current wealth of the investor, it might be worth to take the risk and have this €5 as an option. This characteristic and more is incorporated in EUT. We now describe the components of the theory, as stated in [3].

Definition 2.1. Let C be a set of choices $\{c_1, \dots, c_n; c_{n+1}, \dots, c_m\}$ with $c_i \succ c_j$, $c_k \succ c_l$ for $\{1, \dots, n\}$ and $\{n+1, \dots, m\}$ with $\sum_{i=1}^n p_i = 1$. For a choice c_i , the variables o_j represent the outcomes of the choice and the probabilities of the outcomes.

Definition 2.2. A preference relation, denoted \succsim , determines an ordering over a set of choices C . Given $c_x, c_y \in C$, if $c_x \succsim c_y$, we say that choice x is at least as preferred as choice y .

Assumption 2.1. Every investor can fully define her preferences so that she has a preference relation \succsim .

For clarity sake, we present two notations:

- $x \succ y$: the agent prefers x over y
- $x \sim y$: the agent is indifferent to the choice of x over y

Definition 2.3 (Rationality). *A preference relation is rational if and only if it satisfies the following properties:*

- Reflexivity: for every x , it holds that $x \succsim x$
- Completeness: for every x, y , it holds that $x \succsim y$ or $y \succsim x$
- Transitivity: for every x, y, z such that $x \succsim y$ and $y \succsim z$, it holds that $x \succsim z$

Definition 2.4 (Continuity). *A preference relation is continuous if for every x, y, z , such that $x \succ y$ and $y \succ z$, there exists a scalar $\alpha \in [0, 1]$ such that*

$$\alpha x + (1 - \alpha)z \succ y.$$

Definition 2.5 (Independence). *A preference relation satisfies independence if for all x, y, z , and $\alpha \in [0, 1]$, we have that $x \succ y$ if and only if*

$$\alpha x + (1 - \alpha)z \succ \alpha y + (1 - \alpha)z.$$

If the preference relation of an investor exhibits rationality, continuity, and independence, then EUT can be applied to the investor.

Assumption 2.2. *The preference relation of an investor satisfies rationality, continuity, and independence.*

Definition 2.6 (Expected Utility Function). *Given a preference relation and Assumption 2.2, there exists a function that assigns to a choice the value*

$$u(x) = u(x),$$

where u is the marginal utility function of an investor under certainty.

Theorem 2.1 (Expected Utility Theory). *Given a set of choices as defined in Definition 2.1, a preference relation, and Assumption 2.2, there exist S scalars α_i , $i = 1 \dots S$ such that for every*

$$x \succ y, \quad \alpha_i x + (1 - \alpha_i)z \succ \alpha_i y + (1 - \alpha_i)z,$$

where α_i is defined as in definition 2.6. See [3, p. 22] for the proof.

The marginal utility function of an investor evaluates how satisfied an investor is with that outcome. The function incorporates an investor's risk profile, which can be either risk-averse, risk-neutral, or risk-seeking. An investor's risk profile determines the curvature of the marginal utility function, as can be seen in Figure 2.1. Not only does the curvature of the marginal utility reflect an investor's risk profile, it also incorporates an investor's propensity for risk relative to her current wealth. Since most investors are risk averse, the degree of relative risk is focused on risk aversion. We define two quantities, the coefficient of absolute and relative risk aversion, to clarify the degree of risk-averseness.

Definition 2.7 (Absolute Risk Aversion (ARA)). *For an investor with utility function u , the coefficient of ARA is defined as*

$$A(x) = -\frac{u''(x)}{u'(x)}.$$




Figure 2.1: Risk Profiles

Definition 2.8 (Relative Risk Aversion (RRA)). For an investor with utility function u , the coefficient of RRA is defined as

$$RRA(x) = - \frac{u'(x)}{u(x)} \cdot x.$$

Conclusively, EUT as a way of choosing between choices $\{x_1, \dots, x_n; y_1, \dots, y_n\}$ is based on three principles [17]:

1. Expected utility of a choice: $u(x_1, \dots, x_n; y_1, \dots, y_n) = p_1 u(x_1) + \dots + p_n u(x_n)$.
2. A choice is acceptable if the utility resulting from it increases the current utility: at wealth x , we only choose y if $u(x_1, \dots, x_n; y_1, \dots, y_n) > u(x)$.
3. An investor is risk averse: u is concave ($u'' < 0$)

2.2. Cumulative prospect theory

Another way to evaluate risky investment decisions is with cumulative prospect theory (CPT) [17][18]. Classic EUT considers final states: the increase or decrease in absolute wealth resulting from a decision. CPT considers reference points and classifies outcomes as gains or losses. This new theory was created to get rid of four flaws in EUT. For concrete examples, see [17]. For this thesis, the most important of these flaws are the following two. Firstly, investors often overweight the utilities of certain outcomes. Secondly, when it comes to losses, investors with a propensity for risk aversion actually become risk loving. They often prefer a loss that is merely probable over a smaller loss that is certain. This goes against the principle that the utility function of a risk averse investor is concave.

CPT gets rid of these four flaws of EUT by introducing four elements, as summarized by [1]:

1. reference dependence,
2. loss aversion,
3. diminishing sensitivity,
4. probability weighting.

The first three elements specific to CPT are all incorporated in the value function, which evaluates the outcomes of a choice. The most fitting value function, deduced experimentally, is defined as

$$v(x) = \begin{cases} x^\alpha & \text{for } x \geq 0 \\ -\lambda |x|^\beta & \text{for } x < 0 \end{cases}, \tag{2.1}$$

where α is the coefficient of loss aversion, β the one of gain satiation, and λ that of loss satiation. It is generally assumed that $0 < \alpha < 1$, $\beta = 1$, and $2 < \lambda < 4$. In Figure 2.2, we clearly see the distinction between gains and losses with respect to a reference point. The function exhibits different characteristics for negative and positive outcomes. The observation that investors become risk loving for losses underlies the convexity of the value function in the loss region. The addition of λ incorporates people's tendency to be more sensitive to losses than gains of equal size. The aspect of diminishing sensitivity is reflected by the decreasing steepness of the function in both regions. The last element, probability weighting, deals with the observation that investors overweight low probabilities and underweight high probabilities.




Figure 2.2: The value function (2.1), with $\alpha = 0.88$ and $\beta = 2.25$

Conclusively, combining these four elements, we get that for a choice $(x_1, w_1; \dots; x_n, w_n)$, where $x_i < 0$ for $i < n$ is evaluated as

$$v(x),$$

where $v(x)$, representing the value function, is increasing with $v(0) = 0$ and the w_i are decision weights.

2.3. Utility functions in goal based planning

In Sections 2.1 and 2.2, we explain how an investor evaluates an investment choice and subsequently makes a decision. In theory, the utility function is specific to an investor. An investor clearly lays out her preferences and evaluates choices, from which a utility function can be inferred. Then, an investor chooses the investment decision that results in the highest utility. In turn, this results in the highest satisfaction for an investor. In practice, the utility function is difficult to define. Instead, a general utility function is preferred, based on common tendencies in investors. In the context of goal based investing, the utility is derived from goal completion (see Definition 1.1) instead of accumulated wealth.

In this section, we discuss three utility functions appropriate to the context of goal based planning. These utility functions are derived from EUT and CPT. The utility function has to behave such that utility increases as goal completion increases. To achieve this characteristic, all functions have two elements. First, the input parameter is the goal completion. Within goal based planning, the ultimate goal of an investor is to acquire enough money to complete the goal. Therefore, goal completion is a crucial element to evaluate an investor's satisfaction with an investment decision. Second, if the goal completion is equal to zero, the utility is zero. Additionally, the utility value increases as it reaches a goal completion equal to one. The functions differ in how they increase, i.e. the marginal benefit of increasing goal completion. The functions, however, are only valid as long as the utilities between no satisfaction and complete satisfaction are within the values of these regions [12]. The following three utility functions all portray the two necessary characteristics.

2.3.1. Step utility

A straightforward example of a utility function whose value increases as goal completion increases is a step utility function [12]. It is defined as,

$$v(x) = \begin{cases} 0 & \text{if } x < 1 \\ 100 & \text{if } x = 1, \end{cases} \quad (2.2)$$

where x is goal completion (see Definition 1.1). As seen in Figure 2.3, it is at a constant low when the goal is not satisfied and at a constant high when it is. Only when enough money is saved at

the appropriate time do we attain maximum utility. The advantage of the step utility (2.2) is that no assumptions are made about an investor’s personal preferences. The only input used is an investor’s goal amount. It is thus universal and easily explainable. The disadvantage of the step utility is that it attributes no utility to almost reaching a goal. It interprets especially well those goals which require an exact amount of money, e.g. an investor’s mortgage payments. Anything less than the value required truly means the goal was not achieved. Although when it comes to other goals, this abrupt cut-off is extreme. Suppose an investor wants to remodel her house for €20K. Would she truly be completely unsatisfied should she only have €19,999 available? Or would she be partially satisfied with €15K? We define these goals as bounded goals. For these bounded goals, a discontinuous utility function is not the best option.

Definition 2.9 (Bounded Goals). *Suppose an investor has a goal at time t . The investor determines a range and a goal amount . The investor prefers a goal amount of , but is satisfied with a goal cash flow $[(1 -) ,]$. We call the goal at time t a bounded goal, defined by a lower and upper bound.*




Figure 2.3: Step utility function (2.2).

2.3.2. CRRA utility

For bounded goals as defined in Definition 2.9, we prefer to evaluate the goal completion with a continuous utility function. In expected utility theory, one of the most common families of utility functions is the power family or the family of constant relative risk aversion (CRRA) functions. It is defined as

$$U(c) = \frac{c^{1-c}}{1-c}, \tag{2.3}$$

where c represents goal completion as defined in Definition 1.1 and $(1 - c) > 0$ the investor’s risk aversion as defined in Definition 2.8. The CRRA function is commonly used since the concept of constant RRA accurately reflects investors’ propensity for risk. It implies that no matter investors’ absolute level of wealth, they will keep the same investment profile. However, as noted in [4], the CRRA function does not work in the context of goal completion. When keeping with the traditional concept of diminishing marginal utility (i.e. $c > 0$), most of the utility is already achieved when completing only part of the goal. Suppose we set risk aversion $c = 0.5$. We see in Figure 2.4 that $U(0.5) \approx 0.7$. Almost 70% of the possible total utility is gained when only half of the goal is completed.

The CRRA function has good and bad aspects. On the one hand, the CRRA function (2.3) portrays what the step utility (2.2) lacks: $U'(c) > 0$ for $0 < c < 1$. On the other hand, achieving the utility occurs too quick. This leads to poor spending decisions. If an investor is faced with multiple goals that she cannot achieve simultaneously, the CRRA utility function results in more summed utility for all partially funded goals [4]. Instead, an investor should pursue only some goals which she can fully achieve. Next, we consider another continuous utility function that gains utility slower.

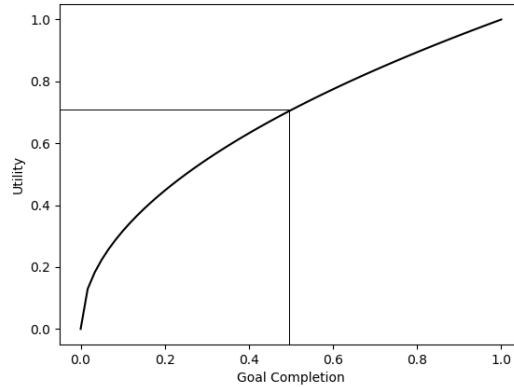


Figure 2.4: The CRRA function (2.3) with $\alpha = 0.5$

2.3.3. Prospect utility

Another continuous utility function to evaluate bounded goals is the value function (2.1) of CPT. As described in Section 2.2, it eliminates certain critiques on EUT. Furthermore, it is regarded as the best possible description of how investors make decisions under risk [1]. The most difficult aspect of CPT is defining the reference point and the regions for gains and losses. In the context of goal based planning, however, the reference point is clear: $\alpha = 1$, with goal completion of Definition 1.1. As supported by research, goals truly serve as reference points and goals inherit the properties of CPT's value function [14]. Taking goal completion α as an input parameter, as defined in [4], the value function becomes

$$v(\alpha) = \begin{cases} (-\alpha^{0.88} + 1) + 2.25 & \text{for } \alpha < 1, \\ (\alpha - 1) + 2.25 & \text{for } \alpha > 1. \end{cases} \quad (2.4)$$

Following the original study on CPT, we choose the parameter values as $\alpha = 0.88$ and $\beta = 2.25$. This function is an improvement on the previous CRRA function, especially when it comes to the utility for $0 < \alpha < 1$. Additionally, it aligns with original CPT, which describes how investors evaluate choices and make decisions in practice.




Figure 2.5: Utility function (2.4) with $\alpha = 0.88$ and $\beta = 2.25$

When it comes to utility functions in goal based planning, we consider three options: the step function (2.2), the CRRA function (2.3), and the value function of CPT (2.4). The step function has the advantage that it is universal and solely dependent on the goal amount of an investor. It only considers success, however, when the goal is completely achieved. This works well for exact goals, but fails for bounded goals. For bounded goals, we consider continuous utility functions. Although the CRRA function can be

used, it has the negative characteristic that utility is gained quickly for low goal completions. The value function of CPT on the other hand is similar to the CRRA function, in that it attributes utility to partly completed goals. The value function, however, assigns utility to partial goal completion slower.

2.4. Expected total utility

As stated in the beginning of this chapter, we use utility theory to create a measure to evaluate an investor’s satisfaction with the feasibility of her goals. Since we operate within the context of goal based planning, the most important factors in the evaluation are the goal completions (see Definition 1.1) and the priority of these goals. To create a total utility function that measures an investor’s satisfaction with her goals’ feasibility, we define two essential components: the marginal utility function and the incorporation of goal priorities. We choose these components such that a higher goal completion results in a higher utility and higher prioritized goals carry more weight. First, based on the functions described in Section 2.3, we define a marginal utility function. This function evaluates a single goal. Second, we define the total utility function, which determines the utility derived from multiple goals. Here we incorporate the goal priorities. Ultimately we use the total utility function to evaluate the expected total utility.

2.4.1. Marginal utility function

For this thesis, we define a marginal utility function based on the utility functions in goal based planning discussed in Section 2.3. For exact goals, we choose the step utility function (2.2). For bounded goals, as defined in Definition 2.9, we want a function that gains utility for partial goal completion. From the discussion above, the prospect utility function (2.4) serves that purpose. This function, however, limits an investor’s input. For some goals an investor is satisfied with a goal completion of 90%, yet for others that percentage could be as low as 50%. Therefore, we define a utility function that on top of goal completion, takes in an investor-specified range as well. The function is

$$\begin{aligned}
 & 0 \text{ for } x < (1 - r) \\
 (x, r) = & 100 \left[\frac{x - (1 - r)}{r} \right]^r \text{ for } (1 - r) < x < 1 \\
 & 100 \text{ for } x \geq 1,
 \end{aligned} \tag{2.5}$$

where x is the goal completion of Definition 1.1 and r the investor-defined range. We use a power of 1.2 to preserve the near-linearity of the original value function (2.1) in the losses region. In Figure 2.6, we see an example of the function for a range of 50%. In practice, we let an investor determine r . We revert to a simple step function if $r = 0$ and an investor wants the exact goal amount. If $r > 0$, we use CPT’s value function to determine the utility gained for $(1 - r) < x < 1$. Suppose $r = 0.3$ and $x = 0.9$. Then $(0.9, 0.3) = 100 \cdot 0.67^r$, since the goal completion falls two-thirds of the way between 70% and 100%, which is the investor-specified success range.




Figure 2.6: Utility function (2.5) with $r = 0.5$

The utility function (2.5) is not the same as the original value function of CPT (2.4), yet it incorporates

the most important factors:

- gains and losses reflect a reference point,
- convex with CRRA in losses and concave in gains,
- it is steeper for losses than for gains.

When it comes to the gains region, we cap the maximum utility at goal completion equal to one. We do this for two reasons. One, an investor will not spend more money than the original goal amount. Thus, she should not receive utility for saving more than she will spend. Second, any excess money is again invested over the period of time until the next goal. Thus, any excess money saved is already incorporated in the utility of the subsequent goal. Gaining utility for > 1 would be counting it twice. Next we incorporate the marginal utility function in the total utility function.

2.4.2. Total utility function

In this section, we define the total utility function which quantifies an investor's satisfaction with the goal completions in a single scenario. In Section 2.4.1, we define the marginal utility function (2.5) for individual goals. To evaluate a scenario, we sum up the marginal utilities, while keeping in mind the goals' priorities. Since an investor wants to achieve her goals simultaneously, we assume linear additivity for the individual goals. As described in Section 1.1, investors denote priority with stars out of five. Then, to incorporate the priority per goal, we multiply the marginal utility of each goal with the number of stars. This way, high-priority goals carry more weight in the total utility function, according to the explicit preferences of an investor. Thus, we define the total utility function

$$(x_1, \dots, x_G) = \sum_{i=1}^G (s_i \cdot u_i(x_i)), \quad (2.6)$$

where G represents the number of goals, s_i the number of stars and x_i the goal completion of goal i , and u_i the utility function (2.5). We do not know the probability distribution of the specific outcomes, thus we run the simulation over multiple scenarios as explained in Section 1.1. The expected total utility of the framework is then the expectation over these scenarios:

$$E[U] = \sum_{j=1}^J (p_j \cdot U(x_1^j, \dots, x_G^j)), \quad (2.7)$$

with U the total utility as defined in equation (2.6) for scenario j . With the expected total utility, we can quantify an investor's satisfaction with the goals' feasibility over all scenarios.

2.5. Conclusion

In this chapter, we created a total utility function (2.6). With the total utility function, we determined the expected total utility (2.7), which quantifies an investor's satisfaction with the goals' feasibility. The total utility evaluates a single scenario and is a priority weighted summation of the marginal utility functions (2.5). The priority weights are the number of stars out of five, since in the framework of Section 1.1, priority is denoted as such. The marginal utility function evaluates an investor's satisfaction with a single goal. We created a marginal utility function based on common utility functions used in the context of goal based planning. It accepts as a main input the goal completion as defined in Definition 1.1, which is essential within goal based planning. Additionally, the marginal function can handle both exact and bounded goals (see Definition 2.9). For that purpose, an investor provides a range of satisfaction, $[a, b]$. If $a = 0$, the investor prefers the exact amount. If $a > 0$, the investor accepts any goal amount in the range $[(1 - a) \cdot b, b]$. For instance, suppose an investor's preferred goal amount is €20K. The investor, however, is also satisfied with anything as of €15K. Then the user specified range is $a = .25$. The marginal utility function does not attribute full utility for goal completion $0.75 < 1$. It increases between that range, with maximum utility at €20K. In the next chapter, we build a goal based planning framework that takes into account goal priorities using the expected total utility.

3

Utility theory framework

The aim of this thesis is to create a goal based planning framework that takes into account goal priorities as stated in Chapter 1. In this chapter, we incorporate the expected total utility (2.7) in the framework of Section 1.1. Using the expected total utility, we can maximise utility by altering the portfolio mix, the periodic deposit amounts and times, the initial wealth, the goal amounts, etc. to ensure optimised goal based plans. Optimising on these components falls short when the priority lies on the second goal. In the original framework of Section 1.1, at each goal time, the money available is spent. Depending on an investor's preferences, however, it might not be the most optimal strategy. Cash flows in the near future influence wealth accumulation for the far future. Keeping in mind total utility and its incorporation of goal priorities, a higher utility could be obtained if we adjust the first goal's cash flows. In this chapter, we use optimisation in combination with the expected total utility function to build such a framework. We consider investors with two goals, of which the second is more important. We ensure that the methods for optimisation are feasible, accurate, and computationally efficient before adding more complexity. First we evaluate the characteristics of the expected total utility in the setting of the original framework. Next, we build a priority-based framework by optimising over the cash flows on the first goal. Lastly, we build a framework where we prevent some cash flows on the first goal in favour of the second goal.

Before optimising on the first goal's cash flows, we consider the expected total utility for two investor's cases in the original framework of Section 1.1. In both cases, we have an investor with two goals, of which the second goal is more important. In the original framework, removing the first goal is the only way to prevent non-optimal cash flows on the first goal. With the expected total utility, we can determine whether an investor prefers to drop her first goal. As per the make-up of the function, it should take into account goal priorities, but not subsequently disregard goal completion defined in Definition 1.1. In both example investor cases, we calculate the expected total utility for an investor who compares the outcome of the framework with two goals and with only the long-term goal. Additionally, we calculate the utility achievability of each goal. With these parameters, we evaluate the expected total utility.

Definition 3.1 (Utility Achievability). *Suppose we have a goal at time t and a utility function, denoted U , which takes as input goal completion as defined in Definition 1.1. For a number of scenarios Ω , we determine the utility achievability for a goal at time t as*

$$[((\cdot) > 0)],$$

where I is the indicator function and C is the goal completion of the goal at time t in scenario $\omega \in \Omega$.

For the first investor's case, we have an investor with goals as in Figure 3.1. The investor invests 30% in stocks and 70% in bonds. The investor saves for the exact goal amount, not satisfied with anything less. In Table 3.1, we see that the investor does not have enough money to achieve both goals with satisfying numbers. As reflected by the expected total utility, the investor is better off not spending money on the first goal. When the investor spends no money on the first goal, she gains no utility from it. The lack of marginal utility of the first goal, however, is compensated by the higher marginal utility of

the second, more important goal. In the next case, we analyse the expected total utilities when more money is available.



Figure 3.1: Investor's initial wealth and goals on a timeline, depicting amount and priority in stars.

	Money spent on goal 1	No money spent on goal 1
Expected total utility	279	402
Achievability goal 1	87%	0%
Achievability goal 2	4%	80%

Table 3.1: Expected total utility calculated with equation (2.7) and goal achievabilities as defined in Definition 3.1 for the investor as in Figure 3.1.

For the second investor's case, the investor's goals are exactly the same as in Figure 3.1. Only now, the investor deposits €300 a year for thirty years as seen in Figure 3.2. The investor again invests 30% in stocks and 70% in bonds. As we can see in Table 3.2, the expected total utility is higher when money is spent on both goals. In the previous example, the higher marginal utility of the second goal compensated for the lack of utility of the first goal. In this case, if money is only spent on the second goal, the marginal utility does not compensate for not spending on the first goal.

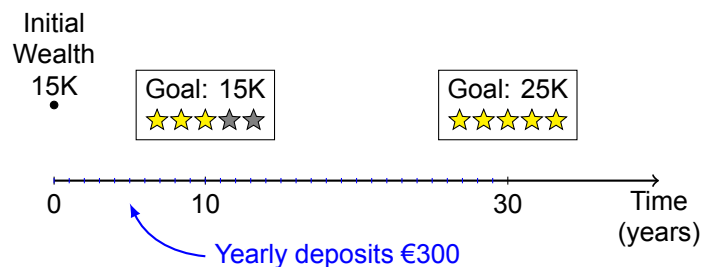


Figure 3.2: Investor's initial wealth, deposits, and goals on a timeline, depicting amount and priority in stars.

	Money spent on goal 1	No money spent on goal 1
Expected total utility	493	482
Achievability goal 1	99%	0%
Achievability goal 2	39%	96%

Table 3.2: Expected total utility calculated with equation (2.7) and goal achievabilities as defined in Definition 3.1 for the investor with profile as in Figure 3.2.

With the measure of expected total utility, we can quantify when spending on the first goal makes the more important, second goal sufficiently infeasible. In that case, the investor should not spend money on the first goal. As seen in the previous examples, the expected total utility balances achievability and priority. When there is enough wealth, either initially or through periodic deposits, the utility is maximal.

when spent on both goal. If there is not enough wealth, utility correctly interprets a higher utility for just the long-term goal. Now we verify if we can achieve a higher utility in a framework where cash flows on the less important goal are guided by the more important goal. Using such a framework, we can compare different goal cash flow patterns, instead of only comparing utilities and achievabilities of an investor with two goals versus just one.

3.1. Upper bound on the first goal's cash flows

In this section, we consider investors with two goals, of which the first goal is bounded (see Definition 2.9) and the second goal is more important. For this type of investor, we can optimise the upper boundary of the first goal's cash flows. With the optimisation, we determine if an investor is satisfied with spending less than preferred on the first goal. Since the first goal is bounded, spending less on the first goal can result in a higher achievability for the more important, second goal, while still obtaining marginal utility for the first goal. We remain within the framework described in Section 1.1. The optimisation works by placing an upper bound on the first goal's cash flows. We optimise for the value $[(1 - \alpha) \cdot C_1, C_1]$, with C_1 the investor specified range of the first goal with value C_1 . As per the marginal utility function (2.5), an investor only gains utility for the first goal in this range. We get the minimisation problem:

$$\begin{aligned} \min & \quad - [U(C_1, C_2)] \\ \text{s.t.} & \quad (1 - \alpha) \cdot C_1 \leq C_1 \\ & \quad 0 \leq C_1 \leq C_1 \quad \text{for } \{1, 2000\}, \end{aligned} \tag{3.1}$$

with C_1 the cash flow on the first goal in scenario $\{1, 2000\}$ and $[U(C_1, C_2)]$ the expected total utility (2.7). If the optimal value is $C_1 = (1 - \alpha) \cdot C_1$, an investor should only spend money on the second goal. In each iteration of the framework, we calculate the goal completion of the first goal as $\frac{C_1}{C_1}$. This way we take note of an investor's original preferences.

For the optimisation, we use Brent's method for minimisation, which combines the golden-section search method and inverse parabolic interpolation [5]. Within the interval $[(1 - \alpha) \cdot C_1, C_1]$, the negative of the expected total utility behaves differently depending on an investor's goals and portfolio. Most often, it behaves in one of three ways. Either the function is strictly decreasing ($\frac{dU}{dC_1} < 0$) or strictly increasing ($\frac{dU}{dC_1} > 0$). If $\frac{dU}{dC_1} = 0$, the function has one extremum, a minimum. Brent's method is capable of handling all these cases.

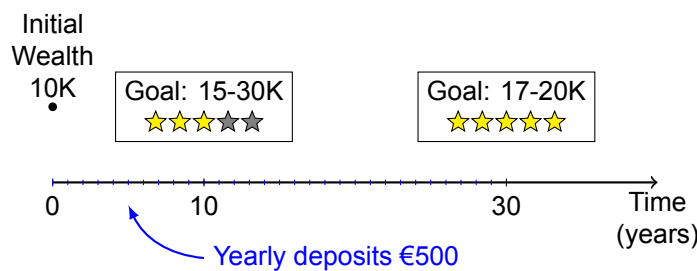


Figure 3.3: Investor's initial wealth, deposits and goals on a timeline, depicting amount and priority in stars.

We consider three similar investor's cases to evaluate the optimisation on the first goal amount. We compare the goals' achievabilities and expected total utilities for the original framework of Section 1.1 with and without optimisation. In each case, the investor invests 30% of her wealth in stocks and the remaining 70% in bonds. For the first investor's case, an investor has two goals, as specified in Figure 3.3. In Table 3.3, we see that it is optimal for the investor to only spend money on the second goal. It results in a higher expected total utility. In the second investor's case, we have a similar investor as in Figure 3.3, only now $C_1 = 20K$. We start with more wealth. With expected total utility, we determine if $\frac{dU}{dC_1} > 0$. From the results in Table 3.4, we see that it is indeed optimal to spend money on the first goal, but not the preferred amount as $C_1 = 22K < 30K$. Since the investor's first goal is bounded,

she gains marginal utility for goal cash flows lower than the preferred amount. In the third investor's case, an investor has goals similar as in Figure 3.3, only $\alpha = 20$ and $\beta = [14, 20]$. The investor's case is similar to the second, but the preferred amount for the first goal is smaller. As expected, the optimal value $\alpha = 20K$. In the second investor's case, $\alpha = 22K$. Now that α is smaller than in the previous example, expected utility is highest when the first goals' cash flows are unbounded.

	Preferred amount	Optimal amount
Upper bound cash flows goal 1	30K	0K
Expected total utility	144	362
Achievability goal 1	91%	0%
Average spent on goal 1	19K	0
Achievability goal 2	25%	78%
Average spent on goal 2	19K	19.8K

Table 3.3: Results of optimising on the upper bound of the first goal's cash flows (see minimisation problem (3.1)) for an investor as in Figure 3.3. We compare it to the framework where we do not place an upper bound on the cash flows of the first goal. Goal achievabilities are defined in Definition 3.1. Averages are calculated over the scenarios in which the marginal utility (2.5) is greater than zero.

	Preferred amount	Optimal amount
Upper bound cash flows goal 1	30K	22K
Expected total utility	477	552
Achievability goal 1	100%	100%
Average spent on goal 1	29K	22K
Achievability goal 2	52%	91%
Average spent on goal 2	19K	20K

Table 3.4: Results of optimising on the upper bound of the first goal's cash flows (see minimisation problem (3.1)) for an investor as in Figure 3.3 with $\alpha = 20K$. We compare it to the framework where we do not place an upper bound on the cash flows of the first goal. Achievabilities are defined in Definition 3.1. Averages are calculated over the scenarios in which the marginal utility (2.5) is greater than zero.

From the investor's cases above, we gather that the optimal upper boundary for the first goal's cash flows depends on the feasibility, range, and relative importance of the goals. In the first investor's case, the optimisation results in $\alpha = (1 - \alpha)$. In the second, the optimisation results in $\alpha = ((1 - \alpha), \alpha)$. In the third, the optimisation results in $\alpha = \alpha$. With this optimisation, we stay in the original framework of Section 1.1. This means that all available money is spent. We only curb the amount that can be spent. In the next section, we create a spending rule for the first goal. With the spending rule, we restrict spending on the first goal in some scenarios.

3.2. Maximum utility barrier

In this chapter, we create a goal based planning framework that takes into account goal priorities using the expected total utility function (2.7). In Section 3.1, we place an upper boundary on the cash flows on the first goal to create such a framework. In this section, we incorporate a spending rule into the framework of Section 1.1. We implement conditional cash flows, such that money is spent if the wealth fulfills a certain condition. Depending on the type of an investor's first goal, we implement the exact or bounded spending rule.

Definition 3.2 (The Exact Spending Rule). *Suppose an investor has a goal at time t with value α . We have N scenarios, with α_i . We determine a value Y . Then the exact spending rule states that in each scenario $\{1, \dots, N\}$, the amount α_i is spent on the first goal if $\alpha_i \geq Y$, with α_i the wealth at time t in scenario i .*

Definition 3.3 (The Bounded Spending Rule). *Suppose an investor has a bounded goal at time t with bounded value $[(1 - \alpha)W_t, W_t]$, with α a range determined by the investor. We have N scenarios, with $\{1, \dots, N\}$. We determine a value Y . Then the bounded spending rule states that in each scenario i , the amount $[(1 - \alpha)W_t, W_t]$ is spent on the first goal if $W_t \geq Y$, with W_t the wealth at time t in scenario i .*

With a correct value Y , we prevent cash flows on the first goal in scenarios with too little money for both goals. In scenarios where there is enough, money is spent on both goals. In this sense, the expected total utility of the investor is increased. We optimise on the expected total utility to calculate the optimal value Y . We define this value as the maximum utility barrier (MU barrier).

Definition 3.4 (The Maximum Utility Barrier). *Suppose an investor has a goal at time t . The goal based planning framework implements the spending rule defined in Definition 3.2 or Definition 3.3, depending on the type of the first goal. The maximum utility barrier for goal t is the value Y in the spending rule such that the expected total utility of equation (2.7) is maximized using equation (3.2).*

The MU barrier is determined by analysing the behaviour of all scenarios. It is applied in every scenario, regardless of their individual performance. The behaviour of the expected total utility in this optimisation is similar to its behaviour in the optimisation of Section 3.1. Therefore, we use the same optimisation technique as above: bounded Brent’s method. Suppose we have two goals, at time t and T , of which the first goal is less important. We optimise expected total utility in the following minimisation problem:

$$\begin{aligned} \min & \quad - \\ \text{s.t.} & \quad 0 \leq Y \leq 2W_t, \end{aligned} \tag{3.2}$$

with U_t the expected total utility function (2.7) and W_t the second goal amount. The framework used in this optimisation implements the exact or bounded spending rule, dependent on the type of the first goal. The optimal value Y is the MU barrier. We optimise for $Y \in [0, 2W_t]$. The upper bound is determined under the assumption that for $W_t \in (0, 2W_t]$, we expect invested wealth to at least retains half its value. If $W_t = 10$, then $W_t \geq 5$ for $W_t > 0$.

We consider an investor with goals and financial parameters as seen in Figure 3.2. The investor invests 30% in stocks and 70% in bonds. In Table 3.2, we see that spending on both goals results in more expected total utility than only spending on the most important goal. Now we run the framework with spending rule as defined in Definition 3.2 and value Y equal to the MU barrier. We see the results of introducing the spending rule in Table 3.5. The investor gains more utility if she spends money on both goals, but spends according to her preferences. When she focuses on the second goal and only spends money on the first goal if $W_t \geq 15 + 6.33$, she attains more satisfying results. Unfortunately, there is a slight loss in the achievabilities of both goals. This loss, however, is negligible when considering the overall improvement.

	No barrier	MU barrier
Barrier value	0	6.33 K
Expected total utility	493	535
Achievability goal 1	94%	58%
Achievability goal 2	39%	72%
Achievability both goals	39%	32%

Table 3.5: Running the framework for an investor as in Figure 3.2 with spending rule of Definition 3.2 with Y equal to the MU barrier defined in Definition 3.4. Expected total utility is calculated using equation (2.7), with achievability as defined in Definition 3.1.

Now we consider the investor’s case for an investor as in Figure 3.3, but with $W_t = 20K$. The investor invests 30% in stocks and 70% in bonds. We considered this investor’s case in Section 3.1. We compared the results of the original framework with and without placing an upper boundary on the first

goal's cash flows in Table 3.4. We now compare those results with the framework that implements the bounded spending rule defined in Definition 3.3 with the MU barrier. In Table 3.6, we see that the framework with spending rule outperforms the two previous frameworks. With the spending rule, the expected total utility is higher and the utility achievabilities equal or higher.

	No barrier	MU barrier
Barrier value	0	6.73 K
Expected total utility	477	608
Achievability goal 1	100%	98%
Average spent on goal 1	29K	24K
Achievability goal 2	52%	93%
Average spent on goal 2	19K	20K
Achievability both goals	52%	92%

Table 3.6: Running the framework for an investor as in Figure 3.3, but $Y = 20K$, with spending rule of Definition 3.3 with Y equal to the MU barrier defined in Definition 3.4. Expected total utility is calculated using equation (2.7), with goal achievability as defined in Definition 3.1.

From the previous investor's cases, we see that implementing a spending rule is beneficial. The framework with a spending rule outperforms the original framework, where money is spent without any regard for future effects on achievability. Additionally, it outperforms the framework with an upper boundary on the first goal's cash flows. It results in higher utilities and higher utility achievabilities. Overall, we see that the total and utility achievabilities improve when an investor considers how current goal cash flows affect future achievabilities.

3.3. Conclusion


In this chapter, we built a goal based planning framework that takes into account goal priorities with the expected total utility function (2.7). We focused on investors with two goals, of which the second is more important than the first. First, we evaluated the behaviour of the function for different investor cases with the framework of Section 1.1. With the expected total utility, we determined that in some cases, the investor should only spend money on the second goal. This occurred when the investor had too little wealth and her investment profile was too risk-averse. Although with more money, the utility was higher when the investor spent money on both goals. Therefore, we attempted to build a framework in which money is spent on both goals, while cash flows are in line with the investor's priorities. In this chapter, we experimented with two different optimisations with the expected total utility as objective function. Firstly, we maximised total utility with an upper bound on the first goal's cash flows. With this optimisation, we determined if the investor was satisfied with spending less on the first goal. Secondly, we optimised over the optimal value Y in the spending rule as defined in Definition 3.2 or 3.3. This resulted in the MU barrier as defined in Definition 3.4. The goal cash flows were allocated such that the investor was most satisfied. The framework with a spending rule and the MU barrier worked better than placing an upper bound on the first goal's cash flows. It resulted in higher expected total utilities and utility achievabilities (see Definition 3.1). With the expected total utility, we have shown that priority-based frameworks perform better than one in which priority is not considered. Additionally, with utility we could allocate goal cash flows according to the investor's priorities, but also determine when it was optimal to do so and when not.

In the next chapters, we develop a different approach to the spending rule. Instead of quantifying people's preference into stars and incorporating those into a utility function, we ask for a preferred achievability for their most important goal. We step away from utility theory for two reasons:

1. Communication: While utility functions are common in financial decision making, they are inherently abstract and individual. Utility functions have obscure parameters which makes communicating the meaning of these functions and their role within the framework difficult towards investors [9]. Additionally, we assume a universal utility function with limited individual input, which goes directly against the nature of the function.

2. Optimisation: The behaviour of the utility function is irregular. In most cases the function only has one minimum, resulting in easy optimisation. In some cases, however, it is not smooth and portrays local minima, which results in slow and sometimes non-optimal optimisation results.

In combination with focusing on achievability, we introduce conditional expectations to the framework. With conditional expectations, we consider the relation between current states and future behaviour of scenarios. Thus we can make decisions now, keeping in mind the knowledge of the future. In the next chapter, we focus on calculating the conditional expectation efficiently and ensuring it is robust. In the chapter thereafter, we create a spending rule similar to those of Definition 3.2 and 3.3 using the information of the conditional expectation. Additionally, we verify whether it leads to better results than the previous optimisation techniques.



Conditional expectation

For this thesis, we build a framework in which we take goal priorities into account as described in Chapter 1. This is especially important for investors with more important goals that occur later in time. Money should not be spent on less important goals if that is detrimental to the achievability of later goals. We achieve this by creating a spending rule per goal, similar to those defined in Definition 3.2 and 3.3. With these spending rules, we implement conditional cash flows, such that money is spent if the wealth fulfills a particular condition. We want to use conditional expectation to create such a spending rule. In this thesis, we use the conditional expectation

$$E[W_T - G_T | W_t], \quad (4.1)$$

where W_t , W_T are the wealth at time t and T , G_T the goal amount at time T , and $W_t < G_T$. This determines the relationship between current wealth and future achievability. Using the expectation, an investor only spends money on goal X if she has enough money left such that she maintains the achievability of the next goal. This extra money, or the barrier, required per goal is obtained from the conditional expectation. The spending rule prevents spending on earlier goals if it puts the achievabilities of later goals in jeopardy; creating a framework taking into account goal priorities.

First, we need to accurately determine the conditional expectation numerically. Suppose an investor has two goals, one at time t_1 with goal amount G_1 and one at time t_2 with goal amount G_2 . We could do the calculation with nested simulation: for each wealth at t_1 in path $\{1, 2000\}$, we calculate

$$E[W_{t_2} - G_2 | W_{t_1}],$$

with W_{t_1} the wealth at t_1 and r_{t_1} the return at time t_1 in path j for $\{1, 2000\}$. Then, we calculate $E[W_{t_2} - G_2 | W_{t_1}]$ for $\{1, 2000\}$. Practically, this is infeasible as the cost of nested simulation grows exponentially. Instead, we estimate the conditional expectation function (CEF) with conditional sample paths and a regression algorithm. For instance, for the same investor as above, we calculate

$$W_{t_2} - G_2 = W_{t_1} (1 + r_{t_1}),$$

for $\{1, 2000\}$. Then we calculate $E[W_{t_2} - G_2 | W_{t_1}]$ and regress on the results of the indicator function; in effect creating the CEF of equation (4.1). This is a well-known technique, first applied to value American options [20]. The relationship between conditional expectation and regression is defined in the following theorem.

Theorem 4.1 (The CEF Decomposition Property). *Given two random variables Y and X , where Y is dependent and X is a vector of covariates, the CEF decomposition property states that*

$$E[Y|X] = E[Y] + \epsilon,$$

where (a) ϵ is mean-independent of X and (b) ϵ is uncorrelated with any function of X .

Proof. () Using properties of conditional expectation, we show that

$$E[\epsilon|X] = E[Y - E[Y|X]|X] = E[Y|X] - E[Y|X] = 0.$$

Additionally,

$$E[\epsilon] = E[E[\epsilon|X]] = E[0] = 0.$$

Thus $E[\epsilon|X] = 0$, which proves mean-indifference.

() Using the law of iterated expectations and the results from (), for any function $f(\cdot)$

$$E[f(\epsilon)] = E[f(\epsilon) E[\epsilon|X]] = 0.$$

Using the information from (), $E[f(\epsilon)] = 0$, which proves uncorrelation. \square

Using Theorem 4.1, we define the CEF as $E[Y|X] = \alpha + \beta X$, where $\alpha = E[Y - \beta X]$ and $\beta = \frac{Cov(Y, X)}{Var(X)}$ as in (4.1). In theory, regressing over the results of the indicator function can approximate the conditional expectation. In practice, we first determine the correct type of regression to approximate (). In this chapter, we start with a simple investor's portfolio to determine the most accurate regression. We then apply it on multiple types of returns, ensuring that it is robust. Next, we define a heuristic algorithm with regression and make it more efficient.

4.1. Regressed conditional expectation

In this section, we discuss several regression techniques for the CEF. We determine which of the techniques results in the closest approximation of the CEF in equation (4.1). Initially, we create three datasets containing different types of returns: normal, autoregressive, and those based on the original DNB data (see Appendix A). We start with the normal returns, as we can calculate the analytical conditional expectation; the best base of comparison for our simulation. Afterwards, we check whether the regression results in an accurate conditional expectation for the other underlying returns.

4.1.1. Creation of datasets

Before we determine which regression technique is most accurate, we create three different sets of portfolio returns: normally distributed, autoregressive, and ones based on the DNB data. We use the normal returns to establish which regression technique is most accurate in Section 4.1.2. From the DNB data, we create returns for an investor's portfolio comprised of 70% stocks and 30% bonds. Thus, we combine 70% of the DNB's stock returns and 30% of the 10-year bond returns. On this set we base the normal and autoregressive returns.

For the autoregressive returns, we follow the model with one lag:

$$r_t = \alpha + \beta r_{t-1} + \epsilon_t, \quad (4.2)$$

with α , β constants and ϵ_t (0, σ^2) a random variable representing noise. We estimate the parameters α , β and σ^2 using the following formulas:

$$\begin{aligned} \hat{\alpha} &= \bar{r} - \hat{\beta} \bar{r} \\ \hat{\beta} &= \frac{Cov(r_t, r_{t-1})}{Var(r_{t-1})} \end{aligned}$$

and

$$\begin{aligned}
 () &= (+ +) \\
 () &= \frac{1}{1 - }
 \end{aligned}$$

We need the expected stationary mean, $[]$, and variance, $()$, of the DNB data. In Figure 4.1, we see that the DNB returns become stationary as of year forty. To calculate $[]$ and $()$, we take the mean and standard deviation over all scenarios between years forty and sixty. This results in $[] = 0.047$ and $() = 0.0071$. Then we choose the parameter $= 0.04$, such that $= 0.14$ and $= 0.084$. We choose the parameter c such that $| | < 1$ to preserve stationarity. Thus we get that the autoregressive returns follow the model:

$$= 0.14 + 0.04+ ,$$

where $(0, 0.084)$.




Figure 4.1: DNB mean returns for a portfolio 70% invested in stocks and 30% in bonds over a period of 60 years.

For the normal returns, we calculate $[]$ and $()$, with $$ the returns in year forty of the DNB returns for $\{1, 2000\}$; the normal returns follow a distribution of $(0.044, 0.079)$. For the normal and autoregressive returns, we create 10,000 scenarios with returns over sixty years. For the DNB data, we have 2000 scenarios.

4.1.2. Fit of regression algorithms

In this section, we evaluate several regression techniques described in Appendix B. To determine the regression that results in the best fit for the CEF, we start with a base case using the set of normally distributed returns of Section 4.1.1. There are only three time steps in our model setting:

- is the starting point,
- is the timing of the first goal,
- is the timing of the second goal.

Suppose the investor starts with $= 1$. The value of the second goal $= 1.1$. Following this base case using the normal returns, we calculate the CEF analytically. We need to calculate the minimum necessary return, $$, such that $=$, i.e.

$$\begin{aligned}
 [(1.1) | =] &= (1.1 | =) \\
 &= (\frac{1.1}{- 1}) \\
 &= 1 - (\frac{1.1}{- 1}) ,
 \end{aligned} \tag{4.3}$$

where F is the cumulative distribution function of $(0.044, 0.079)$. Using the analytical CEF, we determine the accuracy of the regressed CEF. We create a fine wealth grid of 10000 points at $t=1$ over interval $[0.8, 1.4]$. With knowledge of the standard deviation of $(0.044, 0.079)$, we know this interval covers all wealth that either will statistically never, sometimes, and always achieve W . Thus, it fully captures the CEF. We determine wealth at $t=1$ such that $W = W_0(1 + r)$, with r the returns of the first year of the scenario set for $\{1, 10000\}$. We then calculate $(W, F(W))$. We call the results of the indicator function the regression dataset. Next, we run several regression techniques on this dataset.

Ordinary least squares regression

The first type of regression we evaluate is non-linear ordinary least squares (OLS) regression. For details on OLS, see Section B.1. With normal underlying returns, the best functional fit is $1 - e^{-\beta W}$ as per equation (4.3). This does not apply for the other return datasets created in Section 4.1.1, making it not robust. Instead, since the sample points are the result of an indicator function, we use a logistic function with two parameters

$$F(W) = \frac{1}{1 + e^{-\beta(W - \mu)}} \tag{4.4}$$

The two additional parameters μ and β are essential. Parameter β adjusts the steepness of the logistic function, while parameter μ ensures horizontal translation. Since we know that the function's range is equal to $[0, 1]$, we do not add a parameter for vertical stretching or translation. In Figure 4.2, we see the fit of the OLS using a logistic function on the regression dataset. The fit is accurate near the inflection point and the concave section of the function. It is unable, however, to correctly capture the convex section of the function. The inability is not caused by the dataset, since for other data points this issue persists. In effect, it is caused by the functions' different rates to zero.

- $F(W) \rightarrow 0$ at the rate of $e^{-\beta W}$, for $W \rightarrow -\infty$
- $F(W) \rightarrow 1$ at the rate of $e^{-\beta(W - \mu)}$

The logistic function goes to zero slower, hence OLS regression with a logistic function form does not fit correctly.

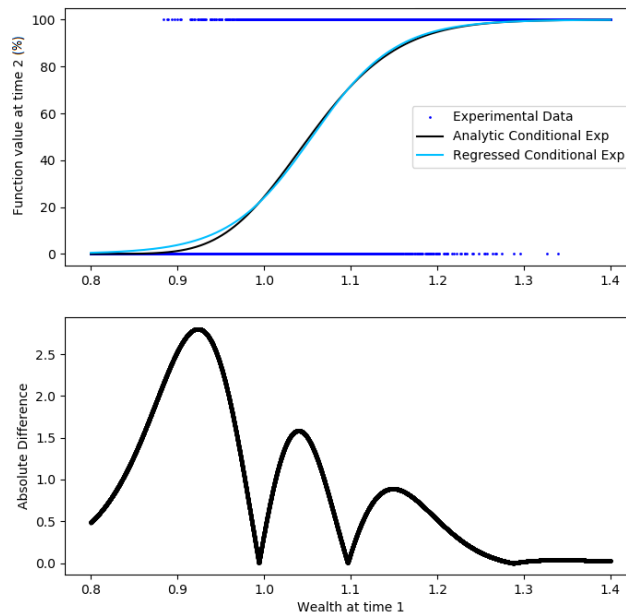


Figure 4.2: Non-linear OLS regression using the logistic function (4.4). The top graph shows the fit of the regression (light blue) compared to the analytic conditional expectation (black). The bottom graph shows the absolute difference between the regression and the analytic function.

Instead of regression with a specified functional form that fulfils the specifications of the data, we evaluate two regression techniques that do not require a predetermined functional form: support vector regression using radial basis functions and spline regression.

Support vector regression

Since regular non-linear OLS regression does not fit correctly, we assess regression techniques that do not require a predefined function. One technique is support vector regression (SVR). For the non-linear data, we use SVR in combination with radial basis functions (RBF) (see Section B.3). The values of γ and C are integral to the workings of this algorithm. The parameter γ determines the width of the RBF, subsequently deciding which support vectors are considered for each x . The parameter C controls the trade-off between flatness and the tolerance for deviations. The SVR fits a sum of peaks, the RBFs, which results in a smooth function. To prevent the function from rising above one or declining below zero, we fix two aspects:

- if $f(x) < 0$ then we fix $f(x) = 0$,
- if $f(x) > 1$ then we fix $f(x) = 1$.

In Figure 4.3, we see the best SVR fit on the regression dataset. The effect of fixing the minimum and maximum value results in two distinct flat sections of the regressed conditional expectation. Although it aligns near the inflection point of the curve, it does not produce the right curvatures nearing zero or hundred. The binary data points complicate the fit. If the RBF is too narrow, the regression only takes into consideration the closest points, which are most likely of the same value. This results in a too steep conditional expectation. If the RBF is too wide, the regressed conditional expectation becomes more averaged and underfits. Adjusting the value of γ cannot overcome these tendencies. In effect, it suffers from the same characteristic. It either pulls the regression too steep or causes underfitting.




Figure 4.3: Support vector regression using RBF for $C = 50$, $\gamma = 7$, $\epsilon = 0.1$. The top graph shows the fit of the regression (light blue) compared to the analytic conditional expectation (black). The bottom graph shows the absolute difference between the regression and the analytic function.

SVR in combination with RBFs does not work for the binary dataset. Now, we evaluate a second type of non-functional regression: spline regression. It builds a piecewise linear function. The idea is to focus on different intervals of the dataset at a time, correctly interpreting each piece.

Spline Regression

Spline regression is a non-functional type which results in a piecewise linear function. The result is piecewise linear, since the spline regression algorithm separates the data and fits each interval separately using breakpoints. The accuracy of the regression is highly dependent on the placement of the breakpoints. Spline regression can be used in combination with different types of basis functions. For more information on spline regression and the basis functions we use in this thesis, we refer to Section B.4. As we state in the beginning of this section, we use 10,000 wealth points over $[0.8, 1.4]$ at $t = 1$. Initially, we determine ten equidistant breakpoints over the interval $[0.8, 1.4]$. We know, however, that due to the basis function, the regression is equal to zero at the left-most breakpoint. Therefore, we determine eleven breakpoints over $\{0\} \cup [0.8, 1.4]$, since $(0) = 0$ and $(0) = 0$.

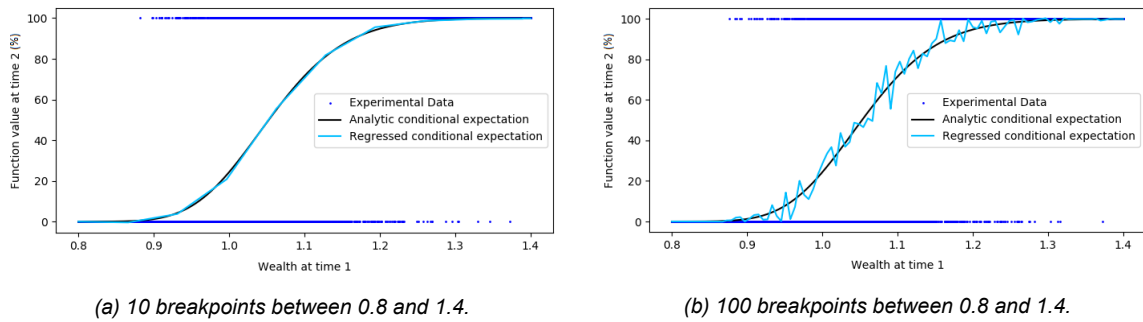


Figure 4.4: Spline regression using OLS coefficients. The graph shows the fit of the regression (light blue) compared to the analytic conditional expectation (black).

Using regular OLS coefficients, we get the result in Figure 4.4a. As expected, we get a piecewise linear function with a discontinuity at the rightmost breakpoint. The absolute differences are small, thus the breakpoints were chosen correctly. Still, using only ten breakpoints is not robust, as the accuracy of the regression is highly dependent on the placement. Instead, we prefer more breakpoints to reduce the importance of the placement. As we see in Figure 4.4b, however, the spline regression using OLS coefficients overfits for a hundred breakpoints. This is a common problem when using OLS. Therefore, we test two types of penalized OLS regressions that overcome this problem: ridge and lasso regression. Both algorithms incorporate penalties on the size of the coefficients. The purpose of using penalized OLS coefficients is to flatten out the peaks we see in Figure 4.4b. For the theory on ridge and lasso regression, we refer to Sections B.2.1 and B.2.2 respectively.

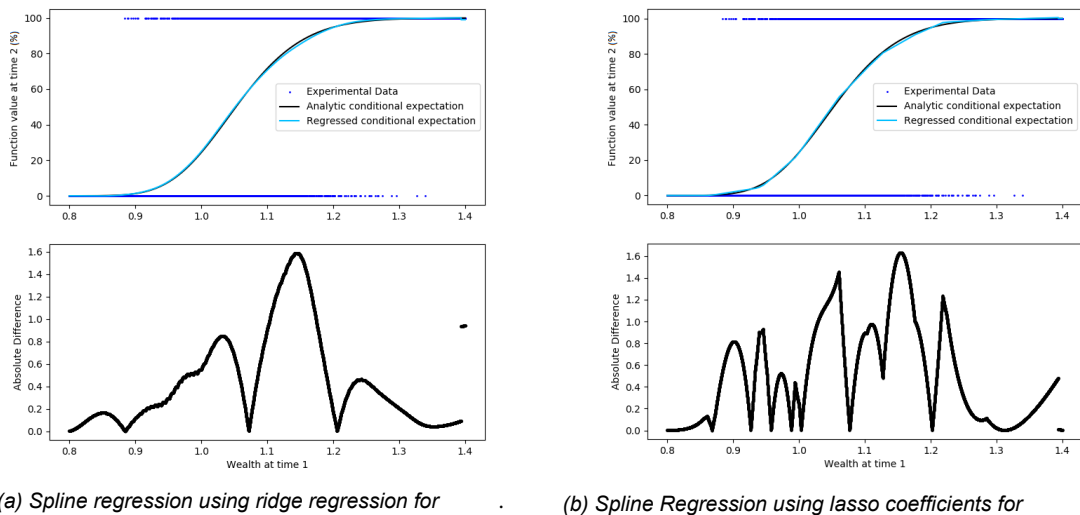


Figure 4.5: Spline regression using two types of penalized ordinary least squares regression. The top graph shows the fit of the regression (light blue) compared to the analytic conditional expectation (black). The bottom graph shows the absolute difference between the regression and the analytic function.

In Figure 4.5, we see the results of spline regression with 100 breakpoints using penalized OLS coefficients. Looking at the absolute differences between the analytic and regressed conditional expectation, both of the regression techniques produce an accurate fit. For the ridge regression, we get a smoother conditional expectation, while the lasso is piecewise linear. Still, we prefer the lasso regression, since it has the ability to set coefficients to zero. Therefore, random and poor data samples affect it less. To verify the robustness of the spline regression using lasso coefficients, we check the technique for different underlying returns and longer time horizons. For convenience sake, we refer to this technique as the lasso spline regression or LSR for short.

4.1.3. Robustness of lasso spline regression

As indicated in Section 4.1.2, the LSR produces an accurate fit for the conditional expectation (4.1) for a simple case of three time steps and normal returns. We want a regression technique that performs well for different returns, time horizons, goal amounts, etc. The regression technique needs to be robust. To check whether the LSR performs well for different parameters, we test it on a more complex case as seen in Figure 4.6. We test this case using the three scenario sets with different returns created in Section 4.1.1.




Figure 4.6: Investor’s initial wealth and goals on a timeline, depicting amount and priority in stars.

Due to the longer time horizon, we can no longer calculate the conditional expectation function $() = [(20) | =]$ analytically. Instead, we simulate an approximation of the conditional expectation using nested simulation with 100 equidistant wealth points at time ten over the interval $[0, 40]$. We use 2000 scenarios. Thus we create the CEF

$$() = [(, 20) | =], \tag{4.5}$$

for $\{1, 2000\}$. We compare the simulated and regressed conditional expectations for the more involved case on the three different returns.

In Figure 4.7, we see the results of the LSR for different underlying returns. We use determine 10,000 wealth points in $[0, 30]$. From the simulated CEF (4.5), we know this is where the gradient varies most. The normal and autoregressive returns have 10,000 scenarios. In the DNB returns, we only have 2000. Therefore, we randomly sample from the existing scenarios (allowing repeats) such that we create 10,000 scenarios. For all returns we use 150 breakpoints and an alpha value of 0.01. Overall, the regression is less accurate than for the simpler case. This occurs since we use the same amount of data points, but over a larger interval. Additionally, the LSR suffers from convergence issues due to the larger domain. Still, it is sufficiently accurate for the priority-based framework as the absolute difference stays below 5% in most cases. By running fifty LSR regressions for an investor as in Figure 4.6 with different sets of normal returns, we determine that the absolute difference stays below 5% in 85% of the regressions. For an investor, the difference between 80% and 83% is unimportant. Before we implement a spending rule based on the conditional expectation, we create a robust algorithm using the LSR. In the previous cases, we specified the wealth interval for the regressions based on the knowledge of the simulated CEF. The algorithm should determine the domain. Additionally, the algorithm should be efficient, without sacrificing too much precision.




Figure 4.7: Spline regression with lasso coefficients for different underlying returns: Normal, AR, and DNB returns. The top graph shows the fit of the regression (light blue) compared to the simulated conditional expectation (black). The bottom graph shows the absolute difference between the regression and the simulated function.

4.2. Algorithm for lasso spline regression

We want to use the conditional expectation to create a spending rule for a priority-based framework. First, we build an algorithm around the LSR. The algorithm should result in an accurate conditional expectation, no matter the input parameters. Additionally, it should be efficient. Ideally, we determine the wealth over the time horizon $0 = < < < =$ such that

$$, = , (1 + ,),$$

with $=$ the initial wealth, $=$ the return at time t , and $\{1, 10000\}$. Then we determine any conditional expectation (4.1) with these wealth points. This is the most consistent and efficient. The accuracy of the LSR, however, is highly dependent on its domain of wealth. It results in more accurate conditional expectations when the data points are where the gradient varies most. Contrarily, if most of the data points are where the function has gradient zero, the regression will be less accurate. We thus cannot rely solely on wealth over the time horizon. Instead, we use the same scenario set as the simulation, but specify the wealth intervals per regression. We determine a heuristic to choose these points. The domain needs to be sufficiently large to capture the complete CEF, but not too large to be detrimental to the accuracy. Additionally, a smaller domain results in a more efficient regression. To build a framework that implements a spending rule based on the conditional expectation, we use the normally distributed returns. We thus create an algorithm for the LSR with those returns in mind.

First, we calculate the leftmost breakpoint, $=$, such that $[()] = 0$. Since the underlying returns are distributed as (4.4%, 7.9%), we use a best case return of $= 1 + (.044 + 3 \cdot 0.079) = 1.281$. Then, taking into account the goal amount, the investment time, and the deposits, the present value equals

$$\begin{aligned}
 &= (1.281) + \left(\frac{1}{1.281}\right) (1.281) \\
 &= (1.281) + \left(\frac{(1.281) - 1}{.281}\right) (1.281) \\
 &= \frac{ - \left(\frac{.}{.}\right) (1.281)}{(1.281)},
 \end{aligned} \tag{4.6}$$

where $=$ is the goal amount, $=$ the present wealth, $=$ the time invested, and $=$ the time deposited. The factor (1.281) is necessary for deposits that end before goal time. We set $= (, 0)$ since we do not include debt in this framework. For the rightmost breakpoint, we use a similar approach.

To calculate the rightmost breakpoint, w_{right} , we need a wealth such that $|(f(w_{right}) - 1)| = 1$. Similarly to the leftmost breakpoint, we calculate the present value for the worst case scenario that still achieves the goal. The return value is $r = 1 + (0.044 - 3 \cdot 0.079) = .807$. The behaviour for increasing t is different for the best and worst case scenarios:

- w_{right} at an exponential rate as $t \rightarrow \infty$,
- w_{left} at an exponential rate as $t \rightarrow \infty$.

For the worst case return, this leads to large wealth intervals, convergence issues, and inefficient LSR. To counteract this, we choose p based on the investment time. The longer that time, the closer the return value is to one. Thus, based on similar logic as in equation (4.6), the present value of the wealth at the rightmost boundary point is given by

$$w_{right} = \frac{1 - (r^{-t})^p}{1 - r^{-t}}, \tag{4.7}$$

where p is determined by the length of the time invested. Again, we set $w_{left} = (w_{right}, 0)$. The number of breakpoints is $2 \cdot p - 1$, such that a larger domain has more breakpoints.

Now that we have a heuristic for the LSR's domain and the breakpoints, we speed up the regression using two techniques: scaling and double regression. Scaling the data has two benefits. It speeds up the process and it eliminates convergence issues. For the scaling, we determine a new goal amount,

$$g' = \frac{g}{10^{\log(w_{right})}} \in [0, 1].$$

We then divide all other parameters by that same factor. The scaling has no effect on the end result, as long as we remain consistent with the factor.

The second technique to speed up the regression is double regression. First, we calculate

$$f(w) = \sum_{i=1}^p \beta_i B_i(w),$$

with B_i and β_i the basis functions as described in Section B.4. The coefficients β_i are determined by running OLS regression on the wealth w , and indicator function $(1, \dots, 1)$ for $w \in \{1, 10000\}$. Then we calculate

$$f'(w) = \sum_{i=1}^p \beta_i' B_i(w),$$

with B_i and β_i' the same basis functions. The coefficients β_i' are determined by running lasso least squares regression on the wealth w and indicator function $(1, \dots, 1)$ for $w \in \{1, 10000\}$. We show the advantage of the double regression with results from concrete cases with varying goal amounts and investment times for LSR with and without double regression. No case involves the use of deposits. We calculate the difference in mean-squared errors and calculation times. For the mean-squared error (MSE), we use

$$MSE = \frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}_i)^2, \tag{4.8}$$

where n is the number of wealth points, y_i is the simulated conditional expectation, and \hat{y}_i is the regressed conditional expectation. The calculation time is measured in seconds, from the beginning of the regression. We see the results of these cases in Table 4.1. The differences are $\Delta = t_{double} - t_{LSR}$. A negative value is thus an improvement of the double regression on the LSR. From Table 4.1a, we gather that the LSR is slightly more accurate. Although in combination with the results in Table 4.1b, we prefer the double regression. In some cases, the LSR is more accurate and quicker. This difference, however, is slight for both deltas. Contrarily, although the double regression is slightly less accurate,

		Invested time		
		[5y,10y]	[10y,20y]	[10y,30y]
Goal amount	1	0.017	-0.009	-0.115
	3	0.009	-0.029	0.017
	5	0.014	0.022	0.018
	7	0.018	0.028	0.007
	9	0.007	0.012	0.041

(a) Mean-squared error

		Invested time		
		[5y,10y]	[10y,20y]	[10y,30y]
Goal amount	1	0.57	0.65	0.78
	3	0.49	0.28	-3.35
	5	-0.83	-1.39	-6.3
	7	-3.92	-1.86	-15.16
	9	-6.75	-11.84	-16.68

(b) Time (s)

Table 4.1: This table shows the absolute difference in MSE and time between lasso spline regression and lasso spline regression after OLS spline regression (double regression). The differences are = - . Note, both methods use a scaling factor.

it can drastically improve calculation time. Especially for conditional expectations with large wealth domains, i.e. for large goal amounts or long investment times, the double regression performs well.

We have created an algorithm for the LSR that can handle different goal values, investment times, deposits, etc. and does so efficiently without the accuracy suffering. The earlier convergence issues have also been resolved by the scaling factor. In effect, we should call the actual regression method that we use the double spline regression with lasso and OLS coefficients. Still, the most important aspect of this regression is the use of the lasso coefficients in combination with the spline regression. The use of an additional spline regression with OLS coefficients is merely for efficiency's sake.

4.3. Conclusion

In this chapter, we showed that spline regression using lasso coefficients accurately and efficiently calculates the conditional expectation (4.1). For the regression technique, we refer to Section 4.1.2. We used regression to avoid nested simulation, which is practically infeasible. Instead, we used conditional sample paths from goal to goal to determine the conditional expectation. For a simple case using normal returns, we compared the results of the regression with the analytical conditional expectation. For 10,000 scenarios, the regression produced an accurate fit. To show that it was independent of the underlying returns and capable of dealing with complex cases, we ran the regression technique for normal, autoregressive, and DNB returns for the investor's case shown in Figure 4.6. We could no longer determine the accuracy by comparing to the analytical conditional expectation. Instead, we compared it to a simulated conditional expectation. For practical usage, we did not want the absolute difference to be larger than 5%. Although the regression was less accurate than for the simple case, the regression was still sufficiently correct. With the LSR, we created a heuristic algorithm. Not only can the algorithm deal with varying parameters such as goal amount and investment times, it does so quickly. We sped up the heuristic with the help of a scaling factor and an additional run of spline regression using OLS coefficients. We managed to speed up the calculation without sacrificing too much accuracy.

In the next chapter, we use the conditional expectation to create a framework that takes into account goal priorities. We create a spending rule such that we do not spend money on earlier goals if it negatively affects the achievability of later, more important goals. As before, we first focus on two-goal investors, of whom the second goal is more important than the first. We then compare the results of a framework with the spending rule based on the conditional expectation to with the spending rule with the MU barrier defined in Definition 3.4. Later, we consider multi-goal investors.

5

Conditional expectation framework

In this chapter, we build a goal based planning framework that takes into account goal priorities as described in Chapter 1 using conditional expectations. The goal is to present an investor the feasibility of her goals, while the goal cash flows are in line with the investor's preferences. An investor will not spend money on a less important goal if it affects her later, more important goals. We build the framework by incorporating a spending rule for goals into the original framework of Section 1.1. The spending rule is similar to those of Definition 3.2 and 3.3. We base the spending rule on the conditional expectation

$$I(\omega) = \begin{cases} 1 & \text{if } W_t \geq W_T + G_T \\ 0 & \text{otherwise} \end{cases}, \quad (5.1)$$

where $I(\omega)$ is the indicator function, W_t , W_T the wealth at time t and T respectively, and G_T the goal amount at time T for $t < T$. We call a spending rule that uses information from the conditional expectation (5.1) a conditional expectation spending rule (CE spending rule). With the conditional expectation, we determine the relationship between earlier wealth (and cash flows) and future goal achievability as defined in Definition 1.2. We require an investor to provide preferred achievabilities for goals. Then, in a framework with a CE spending rule, money is spent on goal X if there is enough money left to obtain the preferred achievabilities for upcoming goals. We use the conditional expectation to determine the extra money needed extra at time t before money is spent on the goal. We call this value the conditional expectation barrier (CE barrier) for goal t .

Definition 5.1 (Conditional Expectation Barrier). *Suppose an investor has two goals: one at time t with amount G_t and one at time T with amount G_T . Then the conditional expectation barrier (CE barrier), denoted B_t , satisfies*

$$I(\omega) = \begin{cases} 1 & \text{if } W_t \geq W_T + G_T - B_t \\ 0 & \text{otherwise} \end{cases}, \quad (5.2)$$

with the indicator function, W_t , W_T the wealth at time t and T respectively, and B_t the preferred achievability of goal T .

In this chapter, we first evaluate the framework with the spending rule for different investor cases. We want the framework to exhibit two characteristics: prioritise obtaining the preferred achievability for the second goal and only spend excess money on the first goal. We limit our focus to two-goal investors, to whom the second goal is more important. For this framework, we use the set of 10,000 scenarios with normal returns that we created in Section 4.1.1. We compare the results of the framework with spending rule to the original framework of Section 1.1. We check whether a spending rule provides a better insight into the feasibility of the investor's goals. Secondly, we determine whether the framework with CE spending rule outperforms the framework with the spending rule using the maximum utility barrier. Thirdly, we introduce early-exercise options, which share some characteristics with an investor's goals. We investigate whether the optimal exercise strategy for early-exercise options can be used within a goal based planning framework. Lastly, we expand the framework and accompanying spending rule to multi-goal investors.

5.1. Conditional expectation barrier

In this section, we evaluate the framework that implements a CE spending rule. In each investor's case, an investor has two goals of which the second goal is more important. Therefore, we prioritise obtaining the preferred achievability of the second goal. For the simulation of the portfolio, we use the data set of 10,000 scenarios of normal returns as specified in Section 4.1.1. With these returns, we efficiently calculate the conditional expectation using the heuristic created in Section 4.2. Firstly, we implement the first CE spending rule.

Definition 5.2 (The First CE Spending Rule). *Suppose we have two goals: one at time t with value W_t and one at time T with value W_T . The goal at time T is more important and has a preferred goal achievability of α (see Definition 1.2). The framework implements the spending rule defined in Definition 3.2. We set the value Y in the spending rule equal to the CE barrier W_t^* as defined in Definition 5.1. Then the first CE spending rule states that in each scenario $\omega \in \Omega$, the amount Y is spent on the first goal if*

$$W_t + [(W_T - W_t) | \mathcal{F}_t] = Y - [(W_T - W_t) | \mathcal{F}_t] .$$

The first investor's case we evaluate is for an investor with goals as in Figure 5.1. The stars merely denote which goal is more important. We focus on the preferred achievabilities. Following the spending rule defined in Definition 5.2, for each scenario $\omega \in \Omega$, money is spent on the first goal if

$$[(W_T - 20) | \mathcal{F}_t] = W_t - 20 \geq 70\% .$$

For this specific case, we calculate a CE barrier $W_t^* = 15.12K$ for a preferred achievability of 70%. Thus the first goal amount of 20K is spent on the first goal in a scenario if $W_t \geq 35.12K$. For a framework implementing the first CE spending rule, we get the goal achievabilities as presented in Table 5.1. The spending rule works, as the framework prevents cash flows on the first goal to achieve a higher achievability on the more important goal. We want, however, the second goal to have an achievability of 70%, not of 94%. This occurs due to our specific spending rule and how we incorporate the information from the conditional expectation.




Figure 5.1: Investor's initial wealth and goals on a timeline, depicting amount and preferred achievability.

	No spending rule	The first spending rule
Achievability goal 1	65%	3%
Achievability goal 2	0%	94%

Table 5.1: Achievabilities for the investor with profile as in Figure 5.1. Achievability is calculated as in Definition 1.2. No spending rule refers to the original framework of Section 1.1. The spending rule refers to the first CE spending rule as defined in Definition 5.2.

To understand the results in Table 5.1, we first explain how to interpret the conditional expectation (5.1). As we state in the beginning of this chapter, the conditional expectation determines the relation between current wealth and future achievability. It is calculated over 10,000 scenarios. A single scenario has no probability of achieving a goal. It either does or does not. Only multiple scenarios have a probability of achieving a goal. As an example, we take the investor and goals of Figure 5.1. We calculate the CE barrier at time ten satisfying

$$W_t^* = 15.12 \text{ such that } [(W_T - 20) | \mathcal{F}_t] = 15.12 = 70\% .$$

Thus we know that for 70% of scenarios with $\text{CE} = 15.12\text{K}$ applies $\text{CE} = 20\text{K}$. Now we come back to the results in Table 5.1. As per the first CE spending rule defined in Definition 5.2, in scenario the first goal amount 20K is only spent on the first goal if $\text{CE} = 35.12\text{K}$. This means that for the 3% of scenarios in which money is spent, at least 70% of those will achieve the second goal. For the other scenarios, we cannot know for sure. In those scenarios, however, no money was spent on the first goal; all money is retained for the second goal. This causes the high achievability for the second goal.

The goal based planning framework with the first CE spending rule does not provide better insight in the investor's goals. The investor sees a higher achievability than preferred for the second goal in combination with an extremely low achievability for the first goal. From the framework, however, we gather that for 95% of the scenarios applies $\text{CE} = 15.12\text{K}$. This leads us to the second CE spending rule, which uses flexible goals.

Definition 5.3 (Flexible Goal). *Suppose an investor has a goal at time t . The goal is flexible if the investor prefers a goal amount of G but allows a goal cash flow of $C < G$ or $C > G$.*

Definition 5.4 (The Second CE Spending Rule). *Suppose we have two goals: one flexible goal (see Definition 5.3) at time t with preferred value G_t and one a time T with value G_T . The goal at time T is more important and has a preferred goal achievability of A_T (see Definition 1.2). We calculate the CE barrier B as defined in Definition 5.1. Then the second CE spending rule states that in each scenario $\{1, \dots, N\}$, the amount $|C_t - G_t|$ is spent on the goal at time t only if $C_t > B$.*

Now we build a framework that implements the second CE spending rule as defined in Definition 5.4. With flexible goals, only providing the goals' achievabilities is not enough. It does not wholly reflect what occurs at the first goal. In addition to achievabilities, we denote the worst, average, and best cash flows. We calculate the values by taking the means of the worst 25%, mid 50%, and top 25% of cash flows respectively.

We again consider an investor with goals as in Figure 5.1. Now we run a framework that implements the second CE spending rule. In Table 5.2, we see the achievabilities and cash flows. As we expected, the achievability of the first goal is the same as for the first CE spending rule (see Table 5.1). The achievability for the second goal is now 70% instead of 94%. This is due to two conditions: a flexible first goal and in 95% of the scenarios $\text{CE} = 15.12\text{K}$. Without these conditions, we cannot guarantee an achievability of 70% of the second goal. Contrary to the framework with the first CE spending rule, this framework does provide a clear overview of the feasibility of the investor's goals. Although the achievability of the first goal is still small, the investor is informed on what "not achieving" means. With this information, the investor can adjust her portfolio and goals accordingly. Additionally, we obtain the preferred percentage for the second goal. On top of that, the framework spend the extra money on the first goal if more wealth comes available.

	The second spending rule
Achievability goal 1	3%
Worst cash flow goal 1	1.9K
Average cash flow goal 1	7.7K
Best cash flow goal 1	15.7K
Achievability goal 2	70%

Table 5.2: Achievabilities for the investor with profile as in Figure 5.1. Achievability is calculated as in Definition 1.2. The spending rule is the second CE spending rule as defined in Definition 5.4. The worst cash flow is the average of the bottom 25% of cash flows, the average that of the mid 50%, and the best that of the top 25%.

With a framework that incorporates goal priorities, we want to prioritise the investor's preferred achievability for the second goal and only spend excess money on the first goal. The framework that implements the second CE spending rule has both characteristics. The first we showed in the previous investor's case. The second, we show for an investor with the same goals as in Figure 5.1. Additionally, the investor deposits 200 each year over fifteen years. From the results in Table 5.2, we know that we

can achieve 70% achievability for the second goal. The investor in this case has more wealth. As we see in Table 5.3, the framework indeed reaches the achievability of 70% for the second goal and only excess money is spent on the first goal. The achievability of the first goal is higher than in Table 5.2. The framework that implements the second CE spending rule thus has both characteristics.

	Flexible spending rule
Achievability goal 1	9%
Worst cash flow goal 1	4.7K
Average cash flow goal 1	11.2K
Best cash flow goal 1	19.7K
Achievability goal 2	71%

Table 5.3: Achievabilities for the investor with profile as in Figure 5.1, with yearly deposits of 200 over the first fifteen years. Achievability is calculated as in Definition 1.2. The spending rule is the second CE spending rule as defined in Definition 5.4. The worst cash flow is the mean of the bottom 25% of cash flows, the average that of the mid 50%, and the best that of the top 25%.

The framework incorporating the second CE spending rule works as we want it to. If possible, the investor's preferred achievability for the most important goal is obtained. Additionally, money is spent on the less important goal only after the achievability of the more important goal is reached. With this framework, the investor gains a clear overview of the feasibility of her goals, taking into account goal priority. On top of achievabilities, we present a more complete cash flow profile for the flexible goal. On the basis of this profile, the investor can adjust her goals and portfolio. The downside of the framework with this spending rule is that the accuracy depends on the wealth at t . The preferred achievability for a goal at time T is reached if in all scenarios $\{1, 10000\}$ the wealth at time t , $W_t = W_0 e^{rt}$, where c is the CE barrier (see Definition 5.1). We have shown that the framework that incorporates the second CE spending rule exhibits the two characteristics. Now we verify whether this framework outperforms the framework with spending rule and MU barrier as defined in Definition 3.4.

5.2. Barrier comparison

In this section, we compare frameworks with spending rules and barriers: the CE barrier defined in Definition 5.1 and the MU barrier defined in Definition 3.4. The latter barrier is calculated by optimising over the expected total utility (2.7). In this chapter, we take a step back from utility theory and focus on preferred achievabilities. Therefore, we define a new barrier called the root barrier, calculated similarly to the MU barrier.

Definition 5.5 (The Root Barrier). *Suppose an investor has two goals: one at time t with a value of c_t and one at time T with a value of c_T . The goal at time T is more important and has a preferred goal achievability of α (see Definition 1.2). The goal based planning framework implements the spending rule defined in Definition 3.2. The root barrier for goal t , denoted Y , is the value Y in the spending rule such that the achievability of goal T , if possible, is equal to α .*

The root barrier is similar to the MU barrier in that we also calculate it by simulating the framework multiple times. We calculate the root barrier by implementing Brent's root finding algorithm [5, Chapter 3-4]. For each iteration of the algorithm, we simulate the framework with spending rule for a different value Y . We continue until $|\text{ach}(c_t, Y) - \alpha| < \epsilon$, with α the preferred achievability of the goal at time T . We implement the bounds $[0, 2c_t]$ under the assumption that invested wealth retains at least half its value over time. With the root barrier, we define the root spending rule. The root spending rule is similar to the first CE spending rule defined in Definition 5.2. Therefore, we first compare the results of the frameworks implementing those spending rules.

Definition 5.6 (The Root Spending Rule). *Suppose an investor has two goals: one at time t with a value of c_t and one at time T with a value of c_T . The goal at time T is more important and has a preferred goal achievability of α (see Definition 1.2). The goal based planning framework implements the spending rule defined in Definition 3.2. We set the value Y in the spending rule equal to the root barrier defined in Definition 5.5.*

We compare the frameworks with spending rules and their barriers for the same investor as in Figure 5.1. In Table 5.4, we combine the results for a framework with the root spending rule and the first CE spending rule. The preferred achievability is reached with the root spending rule, while a decent achievability for the first goal is obtained. That is because β is a third of α . In the former framework, far less scenarios are restricted from spending money on the first goal. This results in a higher achievability. Even when we implement the second CE spending rule, the achievability of the first goal does not increase (see Table 5.2). It seems that the framework with the root spending rule outperforms the framework with the first CE spending rule. However, before validating that statement, we check the conditional achievabilities.

	Root spending rule	CE spending rule
Barrier value	5.7K	15.1 K
Achievability goal 1	30%	3%
Achievability goal 2	70%	94%

Table 5.4: Achievabilities and barrier values for the investor with profile as in Figure 5.1. Achievability is calculated as in Definition 1.2. We run the original framework of Section 1.1 with the spending rules. The root spending rule is as defined in Definition 5.6. The CE spending rule is as defined in 5.2.

At first glance the framework with the root spending rule outperforms the frameworks with either CE spending rules, since the preferred achievability of the second goal is reached while a high achievability for the first is obtained. We want, however, a framework that exhibits two characteristics: prioritise obtaining the preferred achievability of the second goal and only spend excess money on the first goal. The framework with root spending rule exhibits the first characteristic. To verify the second, we look at the conditional achievabilities. We remain in the investor's case where the investor has goals as in Figure 5.1. In Table 5.5, we take a closer look at these values for three spending rules: the root spending rule, the first CE spending rule, and the second CE spending rule. The most relevant values are in the top row. This row denotes the percentage of scenarios in which the second goal was not achieved, given that the first goal was. For the framework with the second CE spending rule, we only include those scenarios in which more than 20K is spent on the first goal. In this aspect, the frameworks with CE spending rules outperform the framework with root spending rule. In the framework with the latter spending rule, 81% of the scenarios in which the first goal is achieved, the second goal is not. This is exactly what we aim to prevent with a framework that takes into account goal priorities. This brings us back to our favoured framework with the second CE spending rule. It performs better on this aspect. The percentage of scenarios is higher than for the first CE spending rule, but remains below 30%. Considering we set the barrier such that $\lfloor \frac{\alpha}{\beta} \rfloor = 15.12 = 70\%$, the maximum percentage is $\lfloor \frac{\alpha}{\beta} \rfloor = 30\%$. We consider a second example to verify this behaviour.

	Root rule	First CE rule	Second CE rule
$\lfloor \frac{\alpha}{\beta} \rfloor$	80%	13%	27%
$\lfloor \frac{\alpha}{\beta} \rfloor$	20%	87%	73%
$\lfloor \frac{\alpha}{\beta} \rfloor$	8%	6%	30%
$\lfloor \frac{\alpha}{\beta} \rfloor$	92%	94%	70%

Table 5.5: Conditional achievabilities for the investor with profile as in Figure 5.1. Achievability is calculated as in Definition 1.2. The root rule is the spending rule as defined in Definition 5.6, the first CE rule as defined in Definition 5.2, and the second CE rule as defined in 5.4. We look at the effect of cash flows on the first goal (G1) on the achievability of the second goal (G2). The symbol '!' in front of G means a goal was not achieved.

For the second example, we evaluate an investor with goals as in Figure 5.2. We compare the framework with the root spending rule and with the second CE spending rule. We run the two frameworks and calculate all necessary values to reach an informed decision on the preferred spending rule. In Table 5.6, we see that a framework with the root spending rule performs better than in the previous example. An achievability of 70% is reached for the second goal, which does not occur with the second

CE spending rule. Additionally, in far less scenarios the second goal is not obtained given that the first goal was obtained. The conditional achievability, $[! | 1]$, confirms this observation. It decreased by 40% compared to the last example. Still, it is higher than preferred and still higher than for a framework with the second CE spending rule.




Figure 5.2: Investor's initial wealth and goals on a timeline, depicting amount and preferred achievability.

	Root rule	CE rule
Barrier value	13.8K	17.6 K
Achievability goal 1	17%	5%
Achievability goal 2	70%	64%
$[! 2 1]$	39%	30%
$[2 1]$	61%	70%
$[! 2! 1]$	28%	36%
$[2! 1]$	72%	64%

Table 5.6: Conditional achievabilities for the investor and portfolio as in Figure 5.2. Achievability is calculated as in Definition 1.2. The root rule is the spending rule as defined in Definition 5.6. The CE rule is the spending rule as defined in Definition 5.4. We look at the effect of cash flows on the first goal (G1) on the achievability of the second goal (G2). The symbol '!' in front of G means a goal was not achieved.

The goal of this section is to determine with which spending rule and barrier the framework performs best. Not only do we look at goal achievabilities, we also consider conditional achievabilities. With these values we can determine whether the framework performs well in the eyes of the investor, while adhering to goal priorities. The advantage of the framework with the root spending rule is its consistency in achieving exactly the preferred achievability, if possible for the portfolio. This is what an investor prefers to see. The disadvantage is that it permits cash flows on the less important goal that are detrimental to the achievability of the second goal. The framework with the second CE spending rule performs exactly opposite. It does not always reach the preferred percentage, but does limit jeopardising cash flows on the first goal. In the end, we go for the second CE spending rule. We accept that the framework sometimes fails to reach the preferred achievability in light of its consistent low conditional achievability. After all, that is the main objective in building a framework that takes into account goal priorities. On top of that, the calculation of the root barrier is computationally intensive, since it requires multiple iterations of the framework. Using the heuristic of Section 4.2, we calculate the CE barrier far more efficiently.

In conclusion, we prefer a framework that takes into account goal priorities by implementing the second CE spending rule as defined in Definition 5.4. The framework exhibits both characteristics: prioritise obtaining the preferred achievability and only spend excess money on the first goal. The most important change is that this framework regards goals as flexible (see Definition 5.3). This is not a disadvantage. We reason that with the spending rule and the cash flow profile per goal, an investor has a clear overview of the feasibility of her goals. With the profile of the cash flows, she can adjust the goal values or portfolio parameters accordingly. The framework that implements the spending rule runs forward in time. Goal cash flows are determined according to the goals' timing. The aim of the next section is to analyse if a framework with a goal cash flows strategy that runs backwards in time results in

better goal cash flows. For that purpose, we introduce early-exercise options and the optimal exercise strategies.

5.3. Early-exercise options

In this section, we discuss the strategy used in the method of Longstaff and Schwartz [20] to price early-exercise options and consider the transferability of the strategy to the goal based planning framework. For early-exercise options, such as Bermudan and swing options, the holder may exercise the option earlier than the agreed maturity time. In the method, the decision to exercise early is dependent on the value of a conditional expectation. This is similar to the framework that takes into account goal priorities by implementing the second CE spending rule (see Definition 5.4). In Section 5.2, we determined that the framework works well if the goal cash flows are determined by conditional expectation (5.1). Given this key commonality, of basing the optimal decision on the value of a conditional expectation, we introduce the topics of Bermudan and swing options. For each type of option, we discuss the use of the Longstaff-Schwartz method in the valuation of the options. In particular, we focus on the optimal strategy implemented for the valuation. We then explore the similarities between swing options and the simulation of an investor’s goals within the framework. We determine whether we can use a strategy, similar to the optimal exercise of swing options, for optimal goal cash flows.

5.3.1. Bermudan options

Before we consider swing options, we introduce Bermudan options and their valuation. Bermudan options are similar to their European counterparts, only the holder can exercise the option any time before the predetermined exercise date. The value of European options can be calculated analytically by means of the Black Scholes model [22, p. 62]. The value of Bermudan options cannot be calculated analytically, due to the addition of early-exercise rights. There are many numerical approaches, however, to price Bermudan options. We focus on approaches using Monte Carlo simulations and regression techniques, in particular the Longstaff and Schwartz’s Least Squares Monte Carlo (LSM) approach [20]. There are similar approaches, but the LSM is well known within Bermudan option pricing and interesting for this thesis.

Problem formulation

To value a Bermudan option, we need to:

1. find the optimal exercise rule,
2. compute the expected discounted payoff under this rule.

Before we explain the LSM approach, we present the problem formulation. Let the risk-free asset, denoted $B(t)$, be modelled by

$$dB(t) = rB(t)dt,$$

with r the risk-free interest rate. Let $S(t)$ be a sufficiently differentiable function of time t and stock price $S = S(t)$. Under the risk neutral measure \mathbb{Q} , the stock price is modelled by

$$dS(t) = \mu(S(t), t)dt + \sigma(S(t), t)dW(t),$$

with $\mu(S(t), t)$ the general drift term and $\sigma(S(t), t)$ the volatility term. Suppose that $V(S(t), t)$ satisfies the partial differential equation

$$-\frac{\partial V}{\partial t} + rV - \frac{1}{2}\sigma^2(S(t), t)\frac{\partial^2 V}{\partial S^2} - \mu(S(t), t)\frac{\partial V}{\partial S} = 0,$$

with final condition given by $V(S(T), T) = \max\{0, S(T) - K\}$. Then the solution $V(S(t), t)$ at any $t < T$ for the filtration \mathcal{F}_t is given by

$$\begin{aligned} V(S(t), t) &= E^{\mathbb{Q}} \left[\frac{1}{B(t)} \max\{0, S(T) - K\} \mid \mathcal{F}_t \right] \\ &= E^{\mathbb{Q}} \left[\exp(-\int_t^T r(u)du) \max\{0, S(T) - K\} \mid \mathcal{F}_t \right]. \end{aligned} \tag{5.3}$$

For options, the value (S, t) is the payoff at maturity time T , denoted

$$(S, t) = \begin{cases} \max(S - K, 0) & \text{for a call option,} \\ \max(K - S, 0) & \text{for a put option,} \end{cases}$$

with K the strike price of the option.

For Bermudan options, the price (S, t) is given by

$$(S, t) = \sup_{\tau \in [t, T]} [\exp(-r(\tau - t)) (S, \tau)], \quad (5.4)$$

where the supremum is obtained for the optimal stopping time

$$\tau = \inf\{t \geq 0 : (S, t) \leq (S, t)\}. \quad (5.5)$$

The optimal stopping time τ is the first time at which the price of the option is smaller than the payoff at this time [11].

Least squares Monte Carlo approach

Following Longstaff and Schwartz [20], we present the LSM approach. Bermudan options can be exercised at a restricted amount of times over the time horizon $[0, T]$. We define M times at which the holder has the right to exercise, given by $0 < t_1 < \dots < t_M = T$. We then create N Monte Carlo paths of stock prices over discrete time steps with $\Delta t = T/M$. Path $\{1, \dots, N\}$ is created by

$$S(t) = S(0) + \sigma W(t), \quad (5.6)$$

with $S(0) = S_0$ and with $W(t)$ dependent on the dynamics of $S(t)$. For the LSM approach, we define the following notations:

- $S_j(t)$: the value of the stock at time t in path $\{1, \dots, N\}$,
- $P_j(t, T) = (S_j(t), T)$ the value of the payoff of asset S ,
- $D_j(t, T) = \exp(-r(T - t))$ the discount value between time t and T ,
- $C_j(t, T) = (S_j(t), T)$ the cash flow generated by the option in path j at time t , given that the option is not exercised at or prior to time t and the holder follows the optimal stopping strategy for all exercise dates between $\{t_1, \dots, t_M\}$,
- $V_j(t, T) = (S_j(t), T)$ the value of the cash flows in path j at time t .

The value (S, t) will be non-zero for only one $\{1, \dots, N\}$ in path $\{1, \dots, N\}$, since the Bermudan option can be exercised once. With the notation clear, we explain the LSM approach.

The LSM approach goes backwards in time. At each point in time, the same exercise rule is followed - only exercise early if the payoff from immediate exercise is larger than the continuation value.

Definition 5.7 (Continuation Value). For a Bermudan option exercisable over time $0 < t_1 < \dots < t_M = T$

and N Monte Carlo simulated paths over the discrete time steps, the continuation value at time t in path $\{1, \dots, N\}$ is given by

$$V_j(t, T) = [C_j(t, T) | \mathcal{F}_t],$$

with \mathbb{Q} the risk neutral measure, $C_j(t, T)$ the value of the cash flows in path j at time t , and \mathcal{F}_t the filtration at time t .

At time $t_M = T$, the value of the option is simply equal to the payoff $P_j(t_M, T)$. At time t_{M-1} , the decision is made by comparing the continuation value $V_j(t_{M-1}, T)$ and the immediate payoff $P_j(t_{M-1}, T)$. The value of the immediate payoff is known, but the continuation value is not. At this stage, regression is utilised.

As the name suggests, least squares regression is an integral part to the LSM approach. Regression is used to approximate the continuation value $V_j(t, T)$ at each time step $t < T$. For the in-the-money paths

$\{1, \dots\}$ at t , the approximation of the continuation value, denoted \hat{V}_t , is calculated using a linear combination of functions over the asset price S_t given by

$$\hat{V}_t = \sum_{i=1}^M \beta_i \phi_i(S_t), \tag{5.7}$$

with ϕ_i a set of basis functions and β_i the regression coefficients at time t . The approximation would be equal to V_t if an infinite amount of basis function were used; for practical purposes, $M < \infty$. Only in-the-money paths are used since these paths are relevant to the early-exercise decision. To estimate the coefficients β_i , least squares regression is run on the in-the-money stock prices at t and the discounted cash flows, i.e. (S_t, C_t) with $C_t = \sum_{j=1}^N \frac{C_{t,j}}{(1+r)^{t-j}}$.

Now, at time t , the early-exercise decision is made by comparing the approximated continuation value \hat{V}_t and the immediate payoff V_t . If $\hat{V}_t > V_t$, the option is exercised early in path j and the cash flow decision at t is set to zero, since the Bermudan option can be exercised once. This exercise strategy is applied for all time steps $t > 0$. After completing the exercise strategy at $t = 0$, the value of the Bermudan option at $t = 0$ is evaluated as the average over the discounted cash flows from each path, given by

$$V_0 = \frac{1}{N} \sum_{j=1}^N \sum_{t=0}^{T-1} \frac{C_{t,j}}{(1+r)^t}, \tag{5.8}$$

with β_t the discount factor between t and $t+1$ and (S_t, C_t) the cash flow generated in path n for $n \in \{1, \dots, N\}$.

5.3.2. Swing options

With knowledge of Bermudan options, we introduce swing options and their valuation using an approach similar to the LSM approach in Section 5.3.1. Swing options are multiple early-exercise contracts that are often used in the energy market [16]. Instead of only having one exercise right over a time horizon $[0, T]$, like with a Bermudan option, the holder of a swing option has multiple exercise rights over $[0, T]$. At each exercise date, the holder has the right to purchase a different amount of energy. This provides the holder flexibility. First, we discuss the setting and parameters of the swing options. Second, we discuss the expansion of the LSM approach to swing options and multiple early-exercise options.

Problem formulation

Before we discuss the valuation of swing options, we define the components. According to Jaillet, Ronn, and Tompaidis [16], there are many different types of swing options, but each type has the following characteristics. If the contract is written at $t = 0$, then the option comes into effect during a period $[t_1, t_2]$ such that $0 < t_1 < t_2 < T$. During this period, the holder can exercise a right up to N times, at $t_1, t_1 + \Delta t, \dots, t_2$. The holder can exercise only once at each $t_i \in \{t_1, t_1 + \Delta t, \dots, t_2\}$. Determined with the swing option is the amount of energy, Q_i , that can be purchased at each date t_i . On each exercise date t_i , the holder can purchase a different amount with increment $\Delta Q_i \in \{1, \dots, Q_i\}$. The holder can purchase more, $Q_i + \Delta Q_i$, or less, $Q_i - \Delta Q_i$, than the predefined amount. This is the “swing” aspect of the option. The increment at each date t_i is restricted, such that the value Q_i lies between the following intervals:

$$[Q_i - \Delta Q_i, Q_i + \Delta Q_i],$$

with $Q_i > 0$.

Additionally, in a typical swing option, there is a restriction on the total volume delivered over the period $[t_1, t_2]$. Any deviation outside of this restriction is penalized at t_2 . The penalty can be determined at t_2 or can be dependent on the state at t_2 . In this thesis, we discuss only swing contracts that have a local effect. This implies that if a right is exercised of an amount $Q_i + \Delta Q_i$ on date t_i , it has no impact on the original energy amount Q_i that can be purchased on a later date. The deviation $\Delta Q_i \in \{+1, -1\}$ is still determined with respect to the value Q_i .

There are two difficulties with the valuation of swing options:

- modelling the underlying price process,
- determining the optimal multiple early-exercise strategy.

In this thesis, we do not consider the modelling of the underlying price process. We focus on the optimal strategy in combination with Monte Carlo simulation paths. We consider in specific the extension of the LSM approach of Section 5.3.1 to multiple early-exercise options in [7].

Expansion of the least squares Monte Carlo approach

We can use an approach similar to the LSM approach of Section 5.3.1 to calculate the value of a swing option. Dörr [7] demonstrates that the LSM approach can handle the multiple early-exercise characteristics of swing options, if we incorporate an extra dimension to the problem: the number of exercise rights left. We call this approach the extended LSM (ELSM) approach. To illustrate the basic structure of the ELSM approach, we consider a swing contract over a period $[0, T]$, with L exercise dates $t_1 < t_2 < \dots < t_L$. The holder of the option has M exercise rights (with $M \geq L$). We only allow positive increments $x_i \in \{0, 1, \dots, M\}$ and we do not enforce a penalty. On an exercise date t_i , $x_i \in \{0, 1, \dots, M\}$, the holder can purchase a certain amount of the underlying for a specified strike price K_i . Similar to the LSM approach, we create N number of Monte Carlo paths defined by equation (5.6). We use similar notation as for the LSM approach.

For the ELSM, we have three dimensions we need to consider:

- the number of Monte Carlo paths, N ,
- the number of time steps, L ,
- the number of exercise rights left, which is at most M .

The stock value S and payoff value P at time t_i in path $\{1, \dots, N\}$ do not depend on the amount of exercise rights left. Neither does the discount factor β between time t_i and t_{i+1} . We do need to adjust the notation of the cash flows, the value of the discounted cashflows, and the continuation value of Definition 5.7 to take into the third dimension, i.e. the exercise rights left:

- $C_j(t_i, S, M)$ the cash flow generated by the option in path j at time t_i , conditional on the option being exercised x_i times at or prior to t_i and the holder following the optimal exercise strategy for all opportunities between t_i and t_L ,
- $C_j^x(t_i, S, M)$ the value of the cash flows in path j at time t_i given that there are x_i exercise rights left.
- $V_j^x(t_i, S, M)$ the continuation value at time t_i with x_i exercise rights left.

The value $V_j^x(t_i, S, M)$ will be non-zero for at most x_i dates $\{t_1, \dots, t_{x_i}\}$ in path $\{1, \dots, N\}$. With the notation clear, we explain the ELSM approach.

Just as the LSM approach, the ELSM approach runs backwards in time. Initially, the cash flows $C_j^x(t_i, S, M)$ are determined for $\{1, \dots, N\}$, $\{1, \dots, N\}$, and $\{1, \dots, N\}$. These cash flows are considered simultaneously, since if there are x_i exercise rights left and x_i exercise dates between t_i and t_L , then the option is exercised early if the payoff $C_j^x(t_i, S, M) > 0$. Now the cash flows for t_{i+1} are determined. First, the continuation value $V_j^{x_i+1}(t_{i+1}, S, M)$ is calculated for one exercise right left. This is similar to a Bermudan option, thus the LSM approach is applied. The cash flows for $\{2, \dots, N\}$ are already known, so no early-exercise decisions need to be made.

At t_{i+1} , the cash flows for $x_{i+1} = 2$ and $x_{i+1} = 1$ are calculated. First, the cash flows are considered for the largest number of exercise rights left, i.e. $x_{i+1} = 2$. The early-exercise condition reads

$$C_j^{x_{i+1}}(t_{i+1}, S, M) + \beta V_j^{x_{i+1}-1}(t_{i+1}, S, M) > V_j^{x_{i+1}}(t_{i+1}, S, M), \quad (5.9)$$

with $C_j^{x_{i+1}}(t_{i+1}, S, M)$ the payoff and $V_j^{x_{i+1}-1}(t_{i+1}, S, M)$, $V_j^{x_{i+1}}(t_{i+1}, S, M)$ the continuation values for one and two exercise rights left respectively, at time t_{i+1} in path $\{1, \dots, N\}$. Before the condition can be evaluated, the continuation values

V_1 and V_2 are calculated. They are calculated similarly as in the LSM approach, with linear regression of the value of the cash flows at t_1 on the basis functions. If condition (5.9) holds for paths $\{1, 2\}$, then $(V_1, V_2, V_3) = (V_1, V_2, V_3)$. In the LSM approach, the cash flows for $t_1 > t_2$ are set to zero in the paths where the early-exercise right is exercised at t_1 . Due to the extra dimension, however, the ELSM approach deviates from the LSM approach in the adjustment of previously determined cash flows. In the ELSM approach, the cash flows are set equal to the cash flows of the previous iteration for one exercise right less, i.e.

$$(V_1, V_2, V_3) = (V_1, V_2, V_3) \text{ and } (V_1, V_2, V_3) = (V_1, V_2, V_3)$$

for $t_1 > t_2$. For $t_1 = t_2$, the ELSM follows the LSM approach. The cash flows for $\{3, 4\}$ remain unchanged.

The algorithm continues for $t_1 > t_2 > t_3$. When $t_1 > t_2$, no cash flows remain unchanged and we apply the following steps for all $\{2, 3\}$. At t_1 , the early-exercise right is exercised if the condition

$$V_1 + V_2 > V_3$$

holds. This is the general version of equation (5.9). We set cash flows $(V_1, V_2, V_3) = (V_1, V_2, V_3)$ for $\{1, 2, 3\}$. Since we set cash flows equal to the value of the cash flow of the previous iteration with one less exercise right, it is important to start each step of the ELSM with the largest number of t_1 . For $t_1 = t_2$, the ELSM follows the steps of the LSM.

The value of the swing option is equal to the average of the discounted cash flows for M exercise rights left over all paths given by

$$V_0 = \frac{1}{M} \sum_{i=1}^M V_i$$

For the ELSM approach for general swing options, with positive and negative volume increments and a penalty function, we refer to [7].

5.3.3. Relation early-exercise options and goal based planning

We explore the transferability of the optimal strategy implemented to value swing options in the ELSM approach onto a framework that takes into account goal achievabilities. When regarding swing calls and an investor's goals, there are several similarities. There are set dates at which a cash flow decision needs to be made and those decisions affect future decisions. For swing calls, the decision to exercise early need to be made, knowing that there is a limited amount of exercise dates. For an investor's goals, the goal cash flows need to be determined, knowing that it effects the wealth for future goals. Additionally, due to the nature of call options and the use of increments, Δ 's, the value purchased differs. Again, this is similar to an investor's goals with different goal amounts. With these similarities in mind, we analyse whether we can modify the strategy as in the ELSM approach such that it results in the optimal goal cash flows.

In both cases, a strategy is implemented to ensure a satisfying result under constraints. The exercise strategy for swing options is constrained by the number of exercise rights and the amounts purchased. The cash flow strategy for an investor's goals is constrained by a limited amount of wealth available over the investment horizon. In Section 5.2, we determined that the goal cash flows align well with an investor's preferences if the decision to spend is based on the conditional expectation (5.1). In the optimal exercise strategy of the ELSM approach, the early-exercise decision is based on the continuation value (see Definition 5.7); also a conditional expectation. There is, however, a clear difference in the behaviour of the underlying over the time horizon for swing options, the stock price $S(t)$, and for goal cash flows, the wealth $W(t)$. This is why, despite the similarities between swing options and an investor's goals, the optimal strategy of swing options does not work for an investor's goals in a goal based planning framework. We elaborate on this statement below.

Even though the strategies for call options and investors' goals determine the best cash flows (whether on energy or on goals), the implementation of the exercise strategy for swing options cannot be modified to a goal based planning framework. A backwards approach is not possible, since the decision for goal

cash flows affects the underlying (S_t) . Therefore, Monte Carlo paths of the wealth cannot be simulated over the time horizon following equation (5.6). Suppose we do follow a backwards approach. An investor has two goals, one at t_1 with amount G_1 and one at t_2 with value G_2 . First, wealth is simulated over the time horizon $0 < t_1 < t_2$ for paths $\{1, \dots, N\}$, such that

$$W(t) = W(0) (1 + R(t)), \quad (5.10)$$

with $W(0) = W_0$ the initial wealth and $R(t)$ the return of the portfolio at time t . At t_1 , money is spent if wealth $W(t_1) < G_1$. No money is spent until time t_2 . Following the ELSM approach, the goal cash flows are determined based on the conditional expectation (5.1). If we can still satisfyingly achieve the last goal, money is spent. The action, however, affects $R(t)$ for $t < t_2$. The cash flow $C(t)$ is no longer relevant, since the decision was based on parameters that are no longer valid. This issue is not present in swing options, since the underlying, the stock price (S_t) , is not affected by a decision to exercise or not. The stock price can be simulated over time, without knowledge of the exercise dates.

We cannot use a similar approach to the optimal exercise strategy of swing options on an investor's goal. A backwards approach fails, since wealth cannot be simulated over time. Instead, we focus on a forward approach. Suppose an investor has M goals, at times $t_1 < t_2 < \dots < t_M$ and with goal amounts $\{G_1, \dots, G_M\}$. She invests in a model portfolio. Then wealth is simulated in N paths such as in equation (5.10) until t_M . At t_M , money is spent if $W(t_M) < G_M$. We follow this pattern until t_1 , at which time money is spent if $W(t_1) < G_1$. In Section 5.2, we showed that the framework works well if we base value V on conditional expectation (5.1) and regard goals as flexible (see Definition 5.3). The framework implements the second CE spending rule defined in Definition 5.4. Therefore, we determine a goal cash flows strategy for multi-goal investors based on this framework in the next section.

5.4. Multi-goal investors

In this section, we create a goal based planning framework that takes into account the priorities of multi-goal investors. We determined in Section 5.3 that we cannot determine the goal cash flows by going backwards in time. Therefore, we expand the forward facing framework that implements the second CE spending rule, as defined in Definition 5.4, to multi-goal investors. Suppose an investor has M goals, ordered by occurrence in time. Each goal $\{1, \dots, M\}$ has a goal amount of G_m and preferred achievability of A_m as defined in Definition 1.2. For multi-goal investors, we create a spending rule such that for each goal m we obtain the preferred achievability A_m . Similar to before, the framework should exhibit two characteristics: prioritise obtaining the preferred achievability in order of importance and only spend excess money on less important goals. In Section 5.2, we determined that the framework implementing the second CE spending rule exhibits the characteristics for two-goal investors. Based on that rule, we define the multi-goal spending rule for multi-goal investors.

Definition 5.8 (The Multi-Goal Spending Rule). *Suppose an investor has M goals, ordered by occurrence in time. Each goal $\{1, \dots, M\}$ has a goal amount of G_m . We have N scenarios, with $N > 0$. We determine values $\{A_1, \dots, A_M\}$. Then the multi-goal spending rule states that for each goal $\{1, \dots, M\}$, in each scenario $\{1, \dots, N\}$, the amount $|W(t_m) - G_m|$ is spent on goal m only if $W(t_m) < G_m$, with $W(t_m)$ the wealth at the time of goal m in scenario i . The goals $\{1, \dots, M\}$ are flexible, as defined in Definition 5.3.*

The goal of this section is to determine the correct values A_m , such that the framework exhibits the two required characteristics. We start with a framework for three-goal investors. This framework, however, is easily expanded to investors with more goals. In previous chapters we limit our focus on two-goal investors, of whom the most important goal occurs last. In this section, we maintain that aspect. We also keep in mind efficiency. We build the framework while considering the computational costs. In the first investor's case, we have an investor as seen in Figure 5.3. We thus need to determine values A_m , for the multi-goal spending rule.

The first values of A_m are based on the concept that an investor stays on track. Suppose an investor has one goal in twenty years. He has a preferred achievability of 75% for that goal. Then for $t \in [1, 19]$ and $\{1, 10000\}$, $[(W(t) - G_1) | W(t) < G_1] = 75\%$. For the investor and portfolio as in Figure 5.3, we calculate three values:




Figure 5.3: Investor's initial wealth and goals on a timeline, depicting amount and preferred achievabilities.

1. such that $[(15) | =] = 70\%$,
2. such that $[(20) | =] = 80\%$,
3. such that $[(20) | =] = 80\%$.

We calculate the values $= 13.3K$ and $= 16.3K$ following the same logic as the CE barrier of Definition 5.1. We calculate the value $= 13.9K$ as a CE barrier neglecting the second goal that ensures we stay on track for the third goal. We determine the values $= 27.2K$ and $= 16.3K$. Money is spent in scenarios $\{1, 10000\}$ at time 5 if $> 27.2K$ and at time 10 if $> 16.3K$. In Table 5.7, we see the results of running the framework that implements the multi-goal spending rule in combination with these values. From the achievabilities in Table 5.7, there is a potential problem with the values since $[(15)] < 70\%$. A possible contributing factor to the lower achievability of the second goal is that there are scenarios such that $< 27.2K$. From the framework, however, we gather that for 99% of scenarios $> 27.2K$. Therefore, we deduce that the value is too small.

	Goal 1	Goal 2	Goal 3
Achievability	81%	64%	80%
Worst cash flow	7.6K	10.6K	18.1K
Average cash flow	15.9K	17.2K	20K
Best cash flow	26.1K	25.2K	20K

Table 5.7: Achievabilities and cash flow profiles for an investor with portfolio as in Figure 5.3. Achievability is calculated as in Definition 1.2. We apply the spending rule defined in Definition 5.8, with values $= \{27.2, 16.3\}$. The worst cash flow is the mean of the bottom 25% of cash flows, the average that of the mid 50%, and the best that of the top 25%.

The value is not enough to ensure a 70% achievability for the second goal and an 80% achievability for the third goal. This is due to our interpretation of staying on track. Suppose we set $= 13.9K$ for all $\{1, 10000\}$ and calculate $[(20) | =]$. This value is 80%, denoting that we indeed stay on track for goal three. Still, that does not mean that for all $\{1, 10000\}$ with $= 13.9$ applies $16.3 =$. Thus to secure an 80% achievability for the more important third goal, money necessary to achieve the second goal is set aside for the third goal. This results in a lower achievability for the second goal. Thus we need to calculate for the multi-goal spending rule differently.

We need to calculate values , such that we stay on track for the next goal in addition to having enough money for the upcoming value. In the previous example, we concluded that the achievability of the second goal was lower than preferred since not 70% of scenarios had $> 31.3K$ before spending. A simple and efficient solution is to calculate for each goal m the value such that

$$[((+)) | =] = ,$$

based on the logic of the CE barrier of Definition 5.1. Then, going back in time, we calculate for each goal only one conditional expectation. With this value, we ensure that in at least the preferred percentage of scenarios there is enough money for the next goal amount and value. The problem

is that we do not know the behaviour of the other scenarios. Suppose an investor has a portfolio with three goals. The second goal is small with a preferred achievability of 40%. The last goal is a large with a high preferred achievability. Then at the time of the first goal, enough money is set aside for the small goal and x_1 . Although only with an expected achievability of 40%. Thus at the time of the second goal, we only know for certain that 40% of the scenarios have wealth x_2 . This jeopardises the achievability of the large, third goal. Depending on the portfolio, the x_2 value underestimates the necessary money for upcoming goals.

To ensure that we stay on track for upcoming goals and above the upcoming x_2 values, we calculate x_3 , similarly to the first attempt. For each goal, we calculate two CE barriers: the barrier needed to stay on track for the next goal and the barrier needed to stay on track for the goals thereafter. This time, instead of interpreting “on track” as an expected achievability over time of X , we interpret it as the wealth in all scenarios above the CE barrier needed to achieve X . We calculate these values by starting at the last goal. For the investor with a portfolio as in Figure 5.3, we calculate three values:

1. x_1 such that $P(W_1 \geq 20 | W_0 = x_1) = 80\%$,
2. x_2 such that $P(W_2 \geq x_2 | W_1 = x_1) = 90\%$,
3. x_3 such that $P(W_3 \geq 15 | W_2 = x_2) = 70\%$.

We calculate values $x_1 = 16.3K$ and $x_2 = 13.3K$ as CE barriers for sequential goals. We calculate value $x_3 = 16.8K$ such that, with relative certainty, in all scenarios $W_3 \geq 16.3K$ after spending on the second goal. We implement a framework with the multi-goal spending rule of Definition 5.8 with $\alpha = \{30.1, 16.3\}$. In Table 5.8, we see the results of running this framework. As we expected, the x_3 with these α values, what is necessary to achieve the preferred achievabilities is overestimated. We obtain a larger achievability for the second goal than preferred. This is due to the separate calculation of x_2 . In five years, 16.8K is set aside to ensure that in 90% of the scenarios $W_2 \geq 16.3K$. On top of that, 13.3K is set aside such that in 70% of scenarios $W_3 \geq 15K$. In combination with the money set aside for x_1 , however, in some of those other 30% of scenarios now $W_3 \geq 15K$. Still, we rather have a higher than preferred achievability than a lower one.

	Goal 1	Goal 2	Goal 3
Achievability	67%	83%	80%
Worst cash flow	4.9K	13.4K	18.1K
Average cash flow	13.0K	20.8K	20K
Best cash flow	23.2K	29.5K	20K

Table 5.8: Achievabilities and cash flow profiles for an investor with portfolio as in Figure 5.3. Achievability is calculated as in Definition 1.2. The framework implements the multi-spending rule of Definition 5.8 with $\alpha = \{30.1, 16.3\}$. The worst cash flow is the mean of the bottom 25% of cash flows, the average that of the mid 50%, and the best that of the top 25%.

From the results in the previous example, we deduce that with the x_3 values for the multi-spending rule what needs to be reserved for future goals is overestimated. For many portfolios, it results in higher than preferred achievabilities. For some portfolios, however, where the last goal is large compared to the first goals and money is insufficient, this value results in the necessary amount. As an example, we take an investor with a portfolio as in Figure 5.4. Testing the framework for the multi-spending rule and the first two methods to calculate x_3 values, the achievability for the second goal is lower than preferred. For both these different x_3 values, the necessary money for the future is underestimated. As we see in Table 5.9, however, it is possible to obtain the preferred achievabilities. Although calculating the values in this way overestimates what is necessary in some cases, they are correct for cases in which the preferred achievabilities can just be obtained.

In this section, we create a multi-goal spending rule for multi-goal investors. We base the multi-goal spending rule on the second CE spending rule (see Definition 5.4) for two-goal investors. We test the spending rule for three-goal investors, but it can easily be expanded to multi-goal investors. We initially calculate x_3 values for the multi-goal spending rule of Definition 5.8 to stay on track. Although

just because we stay on track for future goals, does not mean the wealth in the scenarios are above the value for the next goal. Thus we calculate values based on two requirements: stay on track for the upcoming goal and ensure all scenarios are above the value for goals thereafter. For each goal j , we calculate two values:

1. such that $[() | =] = \%$,
2. such that $[() | =] = 90\%$.

Incorporating the multi-goal spending rule with these values into the framework results in higher than preferred achievabilities, but prevents underachieving which is most important. It thus exhibits the required characteristics: prioritise obtaining the preferred achievabilities in order of importance and only spend excess money.




Figure 5.4: Investor’s initial wealth and goals on a timeline, depicting amount and preferred achievabilities.

	Goal 1	Goal 2	Goal 3
Achievability	0%	48%	79%
Worst cash flow	0K	0.7K	35.3K
Average cash flow	0K	9.7K	40K
Best cash flow	2.0K	23.1K	40K

Table 5.9: Achievabilities and cash flow profiles for an investor with portfolio as in Figure 5.4. The framework implements the multi-spending rule defined in Definition 5.8 with $\alpha = \{31.9, 28.2\}$. The worst cash flow is the mean of the bottom 25% of cash flows, the average that of the mid 50%, and the best that of the top 25%.

5.5. Conclusion

In this chapter, we created a goal based planning framework that uses a spending rule based on conditional expectation (5.1). We took a step back from utility theory and focused on preferred goal achievabilities instead. At first, we considered two-goal investors, of whom the second goal was more important than the first. We determined that a framework with a spending rule based on the CE barrier (see Definition 5.1) with flexible goals (see Definition 5.3) works best. It portrayed the two features we required from the framework. It prioritised the investor’s preferred achievability for the second goal and it spent excess money on the first goal. Additionally, it provided the investor with insight into the feasibility of her goals. The framework with this spending rule also outperformed the framework with the root spending rule defined in Definition 5.6. We based this distinction on the goal achievabilities, the conditional achievabilities, and the computation time. The conditional expectation barrier especially performed better in the last two categories.


We introduced the topic of early-exercise options, such as Bermudan and swing options. In particular, we analysed whether the optimal exercise strategy in the ELSM approach to value swing options could be applied to an investor’s goals in a goal based planning framework. Swing options and an investor’s goals share the characteristic that earlier decisions constrain future decisions. For both, the decision to act is based on a conditional expectation: the continuation value of Definition 5.7 and the conditional expectation (5.1). The backwards optimal exercise strategy for swing options, however, cannot be implemented to determine the optimal goal cash flows. In the ELSM approach, the value of the underlying, the stock price, is known over the entire time horizon. For goal based planning the value of

the underlying, the wealth, cannot be simulated over the entire time horizon. The wealth is dependent on the goal cash flow decisions taken in the future. In goal based planning, we cannot guarantee that decision taken in the past are still optimal in the future. Therefore, we returned to the forward framework that implemented the second CE spending rule.

Based on the performance of the spending rule using the CE barrier for two-goal investors, we expanded it to multi-goal investors. We evaluated the framework with the multi-goal spending rule of Definition 5.8 for different α values. We preferred the spending rule with $\alpha = \{ \alpha_1, \dots, \alpha_n \}$ based on two values:

1. such that $\mathbb{P} (W_{t_j} \geq W_{t_{j+1}}) = \alpha_j$ %,
2. such that $\mathbb{P} (W_{t_j} \geq W_{t_{j+1}}) = 90\%$,

with the indicator function, I_{j+1} , the wealth at the time of the j^{th} and $(j+1)^{\text{th}}$ goal, and G_{j+1} the goal amount of the $(j+1)^{\text{th}}$ goal. With $\alpha_j = \alpha_{j-1} + \alpha_j$, the framework stays on track for the next goal and has relative certainty to achieve the α_j values for the next goals. When using the multi-goal spending rule with these values, the framework often overestimates. For many portfolios the framework obtains higher than preferred achievabilities. Barring numerical inaccuracies, however, it never underestimates achievabilities if wealth in the portfolio allows it. On top of that, the framework upholds the requirement that earlier cash flows do not put later goal achievability in jeopardy.



Conclusion and recommendations

In this chapter, we conclude the thesis and provide recommendations for future research and for Ortec Finance specifically.

6.1. Conclusions

In this thesis, we created a goal based planning framework that takes into account goal priorities. Specifically, we focused on investors with more important goals occurring later in time. The goal cash flows should reflect that; no money should be spent on less important goals if it is detrimental to achieving the more important goals. We built such a framework in two ways. Initially, we focused on investors with two goals, of which the second is more important than the first.

First, we built such a framework using utility theory. With utility theory, an investor can quantify which risky decision will bring her the most satisfaction; the one resulting in the highest utility. In Chapter 2, we created a total utility function (2.6). With the total utility function, we can quantify an investor's satisfaction with the goal completions in a single scenario. The total utility function takes into account individual goal completion and goal priorities. With the total utility function, we can calculate the expected total utility over all scenarios. An investor can then choose the simulation and goal cash flows that result in the highest expected total utility.

In Chapter 3, we used the total utility function to optimise the goal based planning framework such that it takes into account goal priorities. We used the expected utility in two ways. First, we optimised on an upper boundary of the first goal's cash flows. In the simulation, we restricted cash flows on the first goal to determine whether an investor was satisfied with spending less money on the first goal. The expected utility was higher when spending less on the first goal, if that resulted in a higher goal achievability for the second goal. Secondly, we optimised over the expected total utility to determine the maximum utility (MU) barrier defined in Definition 3.4. We implemented a spending rule in the goal based planning framework such that at a goal at time t with amount G_t , money was spent if $\text{wealth} + G_t \geq \text{MU barrier}$. The MU barrier is the value such that the framework reaches maximum expected utility. With the spending rule and MU barrier, we achieved higher expected total utilities than with an upper bound on the first goal's cash flows. We then took a step back from utility theory, however, because of issues with communication and optimisation. We continued with spending rules in combination with conditional expectations.

Second, we built such a framework using conditional expectation (4.1). Instead of attributing goals with priorities, an investor attributed preferred achievabilities. In Chapter 4, we determined that with conditional sample paths and spline regression in combination with lasso coefficients, we can accurately calculate the conditional expectation. Then, we created an efficient heuristic to calculate the conditional expectation for different investor parameters. In Chapter 5, we built a framework that takes into account goal priorities using a spending rule based on the conditional expectation. We determined that the framework with flexible goals (see Definition 5.3) and a spending rule (see Definition 5.4) worked best. It achieved the investor's preferred achievability for the second, more important goal and spent

excess money on the first goal. Next, we expanded the framework to multi-goal investors.

To expand the framework to multi-goal investors, we first looked at early-exercise options. Options such as Bermudan and swing options share commonalities with an investor's goals in the goal based planning framework. We analysed whether the optimal exercise strategy used to value the early-exercise options in the least squares Monte Carlo approach of Longstaff and Schwartz [20] or Dörr [7] could be applied to an investor's goals. It cannot be applied for goals, since the optimal exercise approach is determined by going backwards in time. Instead, we created a forwards approach, based on the framework with the spending rule for two-goal investors. The multi-goal spending rule determines when money can be spent on each goal. The conditional cash flows are based on the conditional expectation (4.1). Money can be spent on each goal, only if an investor stays on track for the upcoming goal and sets aside enough money for the goals thereafter.

Ultimately, we created a goal based planning framework that takes into account goal priorities for multi-goal investors. We created a multi-goal spending rule such that no money is spent on less important goals if that jeopardises achieving more important goals. It ensures that investors stay on track for each goal and set aside enough money for upcoming goals. With the multi-goal spending rule and the values we determined, the amount necessary is overestimated. The framework with this spending rule often results in higher than preferred goal achievabilities. It does not result, however, in lower than preferred goal achievabilities if possible. The goal cash flows are therefore in line with the investor's priorities, which was the aim of this thesis. With a framework that implements the multi-goal spending rule, an investor gains insight into the feasibility of her goals, while adhering to a healthy investment strategy. For each goal, a goal cash flow profile is portrayed such that an investor can accurately ascertain what "not achieving" a goal means. It is therefore an improvement compared to a framework which does not take into account goal priorities. There are, however, some points for future research before the framework is ready for application.

6.2. Recommendations for future research

In this thesis, we determined that a good way to incorporate goal priorities into the goal based planning framework is through conditional expectations. Furthermore, we showed that in such a framework goal cash flows are actually aligned with investors' preferences. This thus provides investors with the feasibility of their goals. There are two clear points of future research.

First, in Chapter 4, we established that spline regression in combination with lasso coefficients produces an accurate fit for normally distributed, autoregressive, and the original DNB (see Appendix A) returns. The heuristic we created in Section 4.2, however, only works for normally distributed returns. A heuristic for the other returns requires a different choice of left and right breakpoints. Additionally, further research could determine whether it is feasible to calculate the conditional expectations required using a single simulation of wealth over the time horizon. It is more efficient to use the same paths to determine the conditional expectations per goal, instead of simulating specific paths and data points each time.

Second, in Chapter 5, we determined a spending rule and accompanying values for multi-goal investors. These values, however, often overestimate what is necessary for future goals. This results in higher than preferred achievabilities for goals. Future research could determine whether a more accurate multi-spending rule is feasible. One that results in the exact preferred goal achievabilities. One possible avenue is to use a different conditional expectation to calculate the values for the multi-goal spending rule.


6.3. Recommendations for Ortec Finance

In this thesis, we determined that it is valuable to incorporate goal priorities into the goal based planning framework. Then, an investor gains insight into the feasibility of her goals, while the goal cash flows are in line with the investor's preferences. The recommendations for Ortec Finance are given, with an eye on the practical side.

In this thesis, we mostly explore incorporating goal priorities into the goal based planning framework through spending rules and conditional cash flows. In Chapter 5, we concluded that the spending rule

works best if the goal cash flows are conditional on the value of the conditional expectation barrier defined in Definition 5.1. The conditional expectation barrier, however, was calculated using a heuristic specific to normally distributed scenarios. The Ortec Finance Scenario Set is far more complicated. Therefore, we propose that Ortec Finance continues with generalising the root barrier of Definition 5.5.

For two-goal investors, the root barrier can be determined with OPAL's optimiser. The additional research is required for multi-goal investors. The values for the spending rule need to be determined iteratively, starting at the second to last goal and going backwards in time. The value for the root barrier for the second to last goal can be determined as for two-goal investors. For the goals earlier in time, the barrier must be calculated with the additional constraint of future goal achievabilities and root barriers.



KNW Scenarios

In this thesis, we use scenario analysis to determine the feasibility of an investor's goals. The scenarios portray different future behaviours of the portfolio and underlying asset classes. We limit portfolios to ratios of stocks to bonds. We use data provided by "de Nederlandsche Bank" (DNB). In particular, we use the DNB's stock returns and the parameters to create interest rate term structures for 2019 Q1 (<https://www.toezicht.dnb.nl/2/50-233246.jsp>). We take the stock returns provided by the DNB for 2000 scenarios over sixty years. For bond returns, we use the return on investment of the 10-year bond, rebalancing the portfolio every year. Per scenario, we calculate

$$\frac{1 + r_{t+1}}{1 + r_t} - 1$$

for $t \in \{1, 2000\}$ and $t \in \{1, 60\}$, where

$$r_t = [r_t^s, r_t^b, \pi_t^e, \pi_t^u] - 1$$

with $r_t^s, r_t^b, \pi_t^e, \pi_t^u$ interest parameters and $r_t^s, r_t^b, \pi_t^e, \pi_t^u$ state variables provided by the DNB.

The scenarios are based on the KNW-capital market model, estimated to reflect the Dutch market [8]. The model describes stock and bonds, both dependent on the inflation process. To ensure a realistic model, the inflation process is modelled alongside the stock and bond market. The real interest rate and the instantaneous expected inflation are modelled by two state variables, which in turn reflect the uncertainty and dynamics in these values. They are collected in the vector X . This thus determines the formulas of the interest rate and the instantaneous expected inflation, respectively

$$r_t^s = \alpha + \beta X_t$$

and

$$\pi_t^e = \gamma + \delta X_t$$

The state variables, collected in the vector X , follow a mean-reverting process around zero with the following dynamics

$$X_t = -KX_{t-1} + Z_t$$

where K is a 2x2 matrix and $Z_t = [Z_t^1, Z_t^2]$. The vector Z is a four dimensional vector of independent Brownian motions which drive the four sources of uncertainty in the market.

- uncertainty about the real interest rate
- uncertainty about the instantaneous expected inflation
- uncertainty about the unexpected inflation
- uncertainty about the stock return

Any correlation between the interest rate and inflation is modelled using ρ and σ . The price index is determined by expected inflation, resulting in the following dynamics

$$\frac{dP}{P} = \pi + \sigma \epsilon$$

with $\sigma^2 = 1$ and $\rho = 1$. The stock index S develops according to

$$\frac{dS}{S} = (r + \lambda) + \sigma \epsilon$$

with $\lambda = 1$, R the nominal instantaneous interest rate and λ the equity risk premium. Lastly, there is the nominal stochastic discount factor, modelled as

$$\frac{dM}{M} = -r - \lambda \epsilon$$

with the time-varying price of risk λ affine in the state variables X ,

$$\lambda = \alpha + \beta X$$

and α , β and X 4×2 . The stochastic discount factor gives the marginal utility ratio between consumption today and in the future. The assumption is that this rate is equal for everyone in complete markets. Furthermore, there is no risk premium for unexpected inflation, i.e. the third row of β contains only zeros, because the risk cannot be identified on the basis of data.

Returns over different time periods are available to estimate the discrete model. For this purpose they use a multivariate Ornstein-Uhlenbeck process of the form

$$dX = \kappa(\mu - X) + \Sigma Z$$

with

$$\Sigma = \begin{bmatrix} \sigma_1 & 0 & 0 & 0 \\ 0 & \sigma_2 & 0 & 0 \\ 0 & 0 & \sigma_3 & 0 \\ 0 & 0 & 0 & \sigma_4 \end{bmatrix} \quad (A.1)$$

Here, X is the vector with the two state variables, P the price index, S the stock index, C the cash wealth index, B the bond wealth index with duration τ , and Z the vector with the four independent Brownian motions extended with two zeros for cash and bond equations.

B

Regression Algorithms

In this appendix, we elaborate on the regression algorithms used for the conditional expectation in Chapter 4: ordinary least squares, ridge, lasso, support vector, and spline regression. We explain all algorithms using x as the independent variable and y as the dependent variable.

B.1. Ordinary Least Squares Regression

In linear ordinary least squares regression (OLS), we try to best approximate the linear relationship between (x_i, y_i) for $i = 1, 2, \dots, n$ and $\beta = (\beta_0, \beta_1, \dots, \beta_k)$. To do so, we estimate the parameters β_0 and β_1 in

$$y_i = \beta_0 + \beta_1 x_i + \epsilon_i,$$

where ϵ_i is an error term. In vector notation, we can absorb the constant term β_0 into the x vector so that it becomes $\tilde{x}_i = (1, x_i, \dots, x_i^k)$, for $i = 1, \dots, n$. Then the vector notation for y is

$$y = X\beta + \epsilon,$$

where X is an $(n \times (k+1))$ matrix with rows (\tilde{x}_i) . In OLS, we approximate β by minimizing over the squared of the error terms

$$\min_{\beta} \|y - X\beta\|^2 \tag{B.1}$$

$$= \min_{\beta} (y - X\beta)^T (y - X\beta), \tag{B.2}$$

with \tilde{x}_i representing row i of matrix X . We can find an explicit solution for this minimisation problem. The optimal approximation of β , denoted $\hat{\beta}$, is

$$\hat{\beta} = \frac{(X^T X)^{-1} X^T y}{(X^T X)^{-1} X^T X} = (X^T X)^{-1} X^T y.$$

Lemma B.1. *When $\sum_{i=1}^n \tilde{x}_i \tilde{x}_i^T = 0$, the estimated parameter $\hat{\beta}$ is unbiased.*

Proof. To show that $\hat{\beta}$ is an unbiased estimator, we plug the vector notation of y into the formula for $\hat{\beta}$ and take the expectation. Using properties of the expectation, we get:

$$\begin{aligned} E[\hat{\beta}] &= E[(X^T X)^{-1} X^T y] \\ &= E[(X^T X)^{-1} X^T (X\beta + \epsilon)] \\ &= (X^T X)^{-1} X^T E[X\beta + \epsilon] \\ &= (X^T X)^{-1} X^T [X\beta + E[\epsilon]] \\ &= \beta. \end{aligned}$$

□

For non-linear OLS, we also minimize over the squared error terms, for the non-linear expression of y :

$$S(\beta) = \sum_{i=1}^n (y_i - f(x_i; \beta))^2, \quad (\text{B.3})$$

with ϵ an error term. Non-linear OLS estimates the parameters β , such that the model function f has the best fit for the given data points.

OLS (linear and non-linear) is a good and frequently used technique to regress over data points. It has, however, the tendency to overfit, resulting in high variance of the coefficients. This causes the regression to be well-aligned with the original dataset, but affects how well it fits future observations. A way to counteract this behaviour is to add a penalty term to the objective function of the minimisation. Thus adding some bias to improve the prediction power of the regression. In this thesis we discuss two types: ridge and lasso regression.

B.2. Penalized Ordinary Least Squares Regression

Again we have data (x_i, y_i) for $i = 1, 2, \dots, n$ and $\beta = (\beta_0, \beta_1, \dots, \beta_p)$. Following the logic of OLS, we estimate β using minimization on the squared residuals. Differently to OLS, we introduce a penalty term, such that the original minimisation problem (B.1) becomes of the form

$$\min_{\beta} \sum_{i=1}^n (y_i - f(x_i; \beta))^2 + \lambda \|\beta\|^2. \quad (\text{B.4})$$

The penalty is placed on the size of the coefficients, which causes a bias to the estimator. The correct balance between bias and variability, however, improves the regression's predictability.

B.2.1. Ridge Regression

One version of least squares regression that incorporates a penalty on the size of the coefficients is ridge regression [15]. It imposes a squared 2-norm on the coefficients combined with a regularization factor λ . The minimisation problem (B.4) becomes

$$\min_{\beta} \sum_{i=1}^n (y_i - f(x_i; \beta))^2 + \lambda \|\beta\|^2.$$

Taking the derivative and performing some algebra, we get the optimal approximation of

$$\hat{\beta} = (X^T X + \lambda I)^{-1} X^T y,$$

where I is the identity matrix. The regularization parameter thus introduces bias, but reduces the variance of the coefficients.

B.2.2. Lasso Regression

Another regression that places a penalty on the size of the coefficients is the lasso regression. Instead of introducing a squared 2-norm on the coefficients, it imposes a 1-norm. Again, a regularization factor is involved. The minimisation problem (B.4) becomes

$$\min_{\beta} \left(\frac{1}{2} \|y - X\beta\|_2^2 + \lambda \|\beta\|_1 \right).$$

Unlike the original OLS and the ridge regression, the minimisation problem does not have an explicit solution for β . Although, by rewriting the lasso minimisation problem as

$$\begin{aligned} \beta &= \operatorname{argmin} \left\{ \frac{1}{2} \|y - X\beta\|_2^2 \right. \\ &\quad \left. \dots \|\beta\|_1 \right\} \end{aligned}$$

we end up with a quadratic programming problem with linear inequality constraints. For this problem efficient and stable algorithms exist [24]. Similarly to the ridge regression, the lasso regression shrinks the size of the coefficients, thus resulting in better predictability of the regression. Additionally, it provides better interpretation as it sets some coefficients to zero [24]. It performs feature selection. It determines which coefficients strongly define the characteristics of the regression and sets the non-important features to zero, resulting in sparsity of the solution. Although the advantage of the lasso regression is two-fold, it does give up computational efficiency since it has no explicit solution.

B.2.3. Regularization Parameter

The difficulty that comes with these penalized OLS regressions is choosing the value of the regularization parameter λ . Although introducing bias to the coefficients can decrease the variance and improve predictability, it can also oversimplify the regression. The extreme cases are clear:

- $\inf_{\beta} \frac{1}{2} \|y - X\beta\|_2^2 = 0$
- $0 \leq \lambda \rightarrow \beta = 0$.

We thus need to fine tune the value of λ such that we achieve the right balance between bias and variance. The standard is to use cross-validation. We first divide the data set into a training and a validation set. We then run the regression on the training set for different values of λ , testing the accuracy with the validation set. We then calculate the mean squared error for each value λ . Thus we choose λ by finding

$$\lambda = \operatorname{argmin}_{\lambda} \left(\frac{1}{n} \|y - X\beta\|_2^2 \right).$$

When there is only a small data set, we use k-fold cross-validation. Instead of determining one selection of the data set as validation set, we rotate the data. We first divide the data set in k partitions. For each value of λ , we select a partition to serve as validation set. The rest of the data serves as training set. We then calculate the mean squared error for the validation set. We repeat the process until all k partitions have been the validation set. The λ which results in the lowest average mean squared error is chosen as regularization parameter.

Although these regression algorithms are well-used and successful, they require knowledge of the functional form underlying the data. When the functional form is not linear, polynomial or other known shape, these algorithms are not useful. Therefore we introduce two regression algorithms that do not require a functional form supplied by the user: support vector and spline regression.

B.3. Support Vector Regression

In this section we discuss support vector regression as denoted by [23]. We first discuss support vector regression (SVR) for linear regression. Afterwards, we shortly introduce the case of non-linear SVR.

Suppose we have observations $\{(x_1, y_1), (x_2, y_2), \dots, (x_n, y_n)\}$ in $\mathbb{R}^d \times \mathbb{R}$. The main idea behind SVR is to find a function $f(x)$ such that for a chosen ϵ , $|y_i - f(x_i)| \leq \epsilon$ for every i . Additionally, we want $f(x)$ to be as flat as possible. The ϵ stands for the error we “tolerate”; only deviations greater than this are considered unacceptable. We define the linear function f as

$$f(x) = w \cdot x + b, \tag{B.5}$$

with $w \in \mathbb{R}^d$, $b \in \mathbb{R}$ and $\langle \cdot, \cdot \rangle$ the dot product in \mathbb{R}^d . To ensure flatness of the function we seek a small $\|w\|$. We can thus solve this problem by minimizing over the norm of w . To ensure this minimisation is feasible, we introduce slack variables ξ_i to deal with infeasible constraints. Thus arriving at the convex minimisation problem for linear SVR.

Definition B.1 (Linear SVR). *For linear SVR, we estimate the parameters w, b solving the following convex minimisation problem:*

$$\begin{aligned} \min \quad & \frac{1}{2} \|w\|^2 + C \sum_{i=1}^n (\xi_i + \eta_i) \\ \text{s.t.} \quad & y_i - \epsilon \leq w \cdot x_i + b + \xi_i \\ & y_i + \epsilon \leq w \cdot x_i + b - \eta_i \\ & \xi_i, \eta_i \geq 0 \end{aligned}$$

The constant $C > 0$ determines the trade-off between the flatness of the function f and the tolerance for deviations larger than ϵ . In Figure B.1, we see a graphical interpretation of the linear SVR.




Figure B.1: The graphical interpretation of linear SVR. The left figure represents a regression problem with points outside of the tolerated barrier. The right figure shows the interpretation of the error.

For the convex minimisation problem described in Definition B.1, we have a dual problem. From this dual, using the saddle point condition and substitutions, we derive the formula for w , called the support vector expansion.

Definition B.2 (Support Vector Expansion). *The support vector expansion of w for linear SVR denotes that w can be entirely written in terms of the observations (x_i, y_i) :*

$$w = \sum_{i=1}^n (\alpha_i - \beta_i) x_i.$$

The constants α_i, β_i are Lagrange multipliers with the following constraints:

$$(\alpha_i - \beta_i) = 0 \text{ and } \alpha_i, \beta_i \in [0, C].$$

The importance of the support vector expansion is that $f(x)$ can be written in terms of only the observations y_i :

$$f(x) = \sum_i (\alpha_i - \beta_i) y_i + \gamma$$

To evaluate $f(x)$ we thus do not have to compute $\phi(x)$ explicitly. This information is necessary for the non-linear SVR. As for γ , we use the Karush-Kuhn-Tucker (KKT) conditions to derive the following bounds:

$$\max\{-\alpha + \beta, \alpha - \beta, \alpha < 0 \text{ or } \beta > 0\} \quad \min\{-\alpha + \beta, \alpha - \beta, \alpha > 0 \text{ or } \beta < 0\}.$$

The KKT conditions also imply that the Lagrange multipliers are only non-zero for the function values that deviate more than ϵ from the original target observation. This means that for all samples inside the ϵ boundary, i.e. for $|\phi(x) - y| \leq \epsilon$, α and β are zero. The observations for which the Lagrange multipliers are non-zero are called support vectors.

Definition B.3 (Support Vectors). *Suppose we have observations $\{(x_1, y_1), (x_2, y_2), \dots, (x_n, y_n)\}$ in $\mathbb{R}^d \times \mathbb{R}$. The support vectors of the SVR are the observations such that $|\phi(x) - y| = \epsilon$. For these observations, the Lagrange multipliers of the dual optimisation problem are non-zero.*

We can also apply support vector regression on non-linear problems. For non-linear relations between x and y , we try to find the best function $f(x) = \sum_i \alpha_i \phi(x) + \gamma$ where $\phi(x)$ is a non-linear map. For the map $\phi(x)$, \mathbb{R}^d is called the feature space.

Example

For $\phi(x) = (x, \bar{x}^2, x^2)$, we can apply the map $f(x) = \sum_i \alpha_i \phi(x) + \gamma$. Training a linear SV machine on the mapped observations $\phi(x)$ would result in a quadratic function.

This approach quickly becomes computationally infeasible. As noted above, the model function can be computed using only the dot product of the observations. As such, we do not need to know the map $\phi(x)$ explicitly. Instead, we can use the kernel, $k(x, x') = \phi(x) \cdot \phi(x')$. For instance, for the example above we get the following kernel $k(x, x') = (x, \bar{x}^2, x^2) \cdot (x', \bar{x}'^2, x'^2) = xx' + \bar{x}\bar{x}'^2 + x^2x'^2$. Computing the kernel is more efficient than mapping the original observations and subsequently calculating the dot product. Using the kernel and the vector expansion of $f(x)$, the formula for $f(x)$ becomes:

$$f(x) = \sum_i (\alpha_i - \beta_i) k(x, x_i) + \gamma$$

The only difference is that $\phi(x)$ can no longer be calculated explicitly.

There are many commonly used kernels, but in this thesis we use the radial basis function (RBF) kernel. We use the Gaussian RBF kernel, denoted $k(x, x') = \exp(-\gamma \|x - x'\|^2)$, where $\|x - x'\|^2$ is the squared Euclidean distance. As seen in Figure B.2, parameter γ determines the width of the kernel and is an important feature. The γ determines how far the influence of an observation reaches and thus is a major factor in the over or underfitting of the regression. It is interpreted as the inverse of the radius of influence of the support vectors.

The main advantage of SVR is that, when using the RBF kernel, it can be used without knowing the underlying relationship between the observations. This is a major improvement on non-linear OLS, where a model function needs to be given. The effectiveness of the SVR, however, is highly dependent on the values of ϵ and γ . These factors contribute to the level the regression overfits.

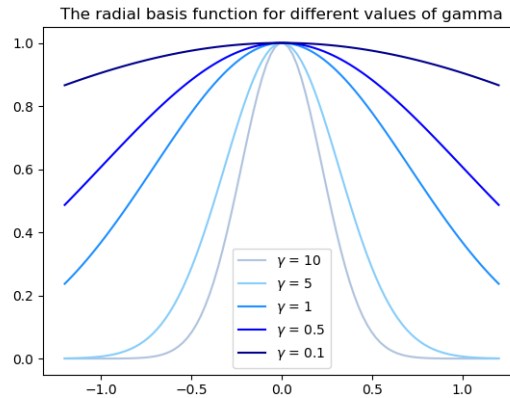


Figure B.2: The Gaussian radial basis function for different values of γ . The parameter determines the width of the peak.

B.4. Spline Regression

The idea behind spline regression is to divide the range of x into contiguous intervals and fit a specific function on each interval [13]. This way, you circumvent fitting one particular functional form on the data. The spline is made up of M basis functions ϕ_1, \dots, ϕ_M so that

$$f(x) = \sum_{j=1}^M \beta_j \phi_j(x).$$

A common type of spline basis functions are piecewise polynomials. The simplest basis function is the indicator function $\phi_j(x) = \begin{cases} 1 & \text{if } x < x_j \\ 0 & \text{else} \end{cases}$. This basis function is a constant representing the mean of $[x_j, \infty)$. To create a piecewise linear fit, we multiply the indicator function by x , such that $\phi_j(x) = \begin{cases} x & \text{if } x < x_j \\ 0 & \text{else} \end{cases}$. These basis functions are unsuited for continuous functions, since the spline $f(x)$ is not continuous at the knots. The knots are the points $(x_j, f(x_j))$, where the x_j 's are the edge points of the intervals. We thus have to determine an extra condition so that the polynomials on either side of a knot are continuous at the knot. Many other family of basis functions are possible, chosen depending on the shape and characteristics of the original data. Most importantly, the family of basis functions should be general enough that they can represent various shapes.

In this thesis we will be using a set of triangular basis functions. If we have breakpoints x_1, \dots, x_M , then the spline is $f(x) = \sum_{j=1}^M \beta_j \phi_j(x)$ for basis functions $\phi_1(x), \dots, \phi_M(x)$.

As shown in Figure B.3, the basis functions are

- $\phi_1(x) = \begin{cases} x & \text{for } x < x_1 \\ 0 & \text{else} \end{cases}$,
- $\phi_j(x) = \begin{cases} x - x_{j-1} & \text{if } x_{j-1} < x < x_j \\ 0 & \text{else} \end{cases}$,
- $\phi_M(x) = \begin{cases} 1 & \text{if } x > x_M \\ 0 & \text{else} \end{cases}$.

The intuition is that $\phi_1(x) = 0$ at x_1 and sets the slope for $x < x_1$. Then $\phi_2(x)$ compensates and sets the slope for $x_1 < x < x_2$. Right of x_2 , the spline equals a constant. The resulting spline is a piecewise, linear, continuous function for $x < x_M$. At x_M , there is a potential discontinuity for which $\phi_M(x)$ is a constant. At x_M the spline equals zero and for $x > x_M$ we have a linear extrapolation.

The difficulty with spline regression is using the correct intervals. The breakpoints and knots should be chosen in areas of high variability, such that the spline captures the complexity of the data. If the intervals are not placed correctly, the spline regression could misrepresent the relation between x and

y. Just as for the regularization parameter, cross-validation is used to determine the placement and number of breakpoints for spline regression.

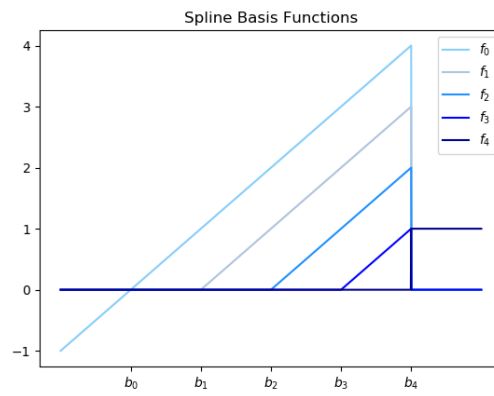


Figure B.3: The basis functions for the spline regression

Bibliography

- [1] N. C. Barberis. Thirty years of prospect theory in economics: A review and assessment. *Journal of Economic Perspectives*, 27(1):173–96, 2013.
- [2] D. Barton, J. Manyika, T. Koller, P. Robert, J. Godsall, and J. Zoffer. Measuring the economic impact of short-termism. 2017.
- [3] E. Barucci and C. Fontana. *Financial Markets Theory*. Springer, 2017.
- [4] D. Blanchett. The value of goals-based financial planning. *Journal of Financial Planning*, 28(6), 2015.
- [5] R. P. Brent. *Algorithms for Minimization Without Derivatives*. Prentice-Hall, 1973.
- [6] M.A.H. Dempster and E.A. Medova. Asset liability management for individual households. *British Actuarial Journal*, 16:405 – 439, 07 2011.
- [7] U. Dörr. Valuation of swing options and examination of exercise strategies by monte carlo techniques. *Mathematical Finance*, 10:27, 2003.
- [8] N. Draper. A financial market model for the netherlands. Cpb background document, CPB Netherlands Bureau for Economic Policy Analysis, 03 2014.
- [9] P. A. Forsyth. Optimal dynamic asset allocation for dc plan accumulation/decumulation: Ambition-cvar. *Insurance: Mathematics and Economics*, 93:230 – 245, 2020.
- [10] G. B. Fowler and V. de Vassal. Holistic asset allocation for private individuals: Simultaneously solving for an optimal asset allocation given the multiple goals and multiple locations of private clients. *The Journal of Wealth Management*, 9(1):18–30, 2006.
- [11] P. Glasserman. *Monte Carlo methods in financial engineering*, volume 53. Springer Science & Business Media, 2013.
- [12] P. Haddawy and S. Hanks. Issues in decision-theoretic planning: Symbolic goals and numeric utilities. In *Proceedings of the DARPA Workshop on Innovative Approaches to Planning, Scheduling, and Control*, pages 48–58, 1990.
- [13] T. Hastie, R. Tibshirani, and J. Friedman. *The elements of statistical learning: data mining, inference, and prediction*. Springer Science & Business Media, 2009.
- [14] C. Heath, R. Larrick, and G. Wu. Goals as reference points. *Cognitive Psychology*, 38:79–109, 1999.
- [15] A. E. Hoerl. Application of ridge analysis to regression problems. *Chemical Engineering Progress*, 58:54–59, 1962.
- [16] P. Jaillet, E. I. Ronn, and S. Tompaidis. Valuation of commodity-based swing options. *Management science*, 50(7):909 –921, 2004.
- [17] D. Kahneman and A. Tversky. Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2):263–292, 03 1979.
- [18] D. Kahneman and A. Tversky. Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, 5:292–323, 1992.
- [19] W. C. Kim, D. Kwon, Y. Lee, J. H. Kim, and C. Lin. Personalized goal-based investing via multi-stage stochastic goal programming. *Quantitative Finance*, 20(3):515–526, 2020.

-
- [20] F. A. Longstaff and E. S. Schwartz. Valuing american options by simulation: a simple least-squares approach. *The review of financial studies*, 14(1):113–147, 2001.
- [21] W. Mischel, B. E. Ebbsen, and A. Raskoff Zeiss. Cognitive and attentional mechanisms in delay of gratification. *Journal of personality and social psychology*, 21 2:204–18, 1972.
- [22] C. Oosterlee and L. Grzelak. *Mathematical modeling and computational finance*. World Scientific Publishing Europe Ltd., 2020.
- [23] A. J. Smola and B. Schölkopf. A tutorial on support vector regression. *Statistics and computing*, 14(3):199–222, 2004.
- [24] R. Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)*, 58(1):267–288, 1996.