Financial Optimization Problems in Life and Pension Insurance

Peter Holm Nielsen

Thesis submitted for the Ph.D. degree
Laboratory of Actuarial Mathematics
Department of Applied Mathematics and Statistics
Institute for Mathematical Sciences
Faculty of Science
University of Copenhagen

Supervisors:
Hanspeter Schmidli
Mogens Steffensen

Thesis committee:
Martin Jacobsen, University of Copenhagen
Ralf Korn, University of Kaiserslautern
Fred Espen Benth, University of Oslo

December 2005

Preface

This thesis has been prepared in partial fulfillment of the requirements for the Ph.D. degree at the Laboratory of Actuarial Mathematics in the Department of Applied Mathematics and Statistics, Institute for Mathematical Sciences, Faculty of Science, University of Copenhagen. The work has been carried out in the period from December 2002 to December 2005 under the supervision of Hanspeter Schmidli (University of Copenhagen until November 2004; now University of Cologne) and Mogens Steffensen (University of Copenhagen).

Acknowledgements

I would like to thank Hanspeter Schmidli for encouragement and support, and for comments on early versions of the published papers. I am deeply indebted to Mogens Steffensen for his encouragement, support, and general guidance, and for many valuable discussions in relation to my work as well as to other matters. He has played a role of great significance as a supervisor, colleague, and friend.

I also wish to thank Ragnar Norberg, former professor at the Laboratory of Actuarial Mathematics at the University of Copenhagen and supervisor of my Master's Thesis, who encouraged me to take the plunge and study for the Ph.D. degree. His continued interest, support, encouragement, and willingness to share and discuss ideas have been remarkable. He also made the necessary arrangements to ensure that a desk was provided for me at the Financial Markets Group at the London School of Economics during my stay in London from September 2003 to January 2004.

I highly appreciate the support and interest shown to me by many of my colleagues, academics as well as practitioners. I wish to mention in particular Mikkel Dahl, my fellow Ph.D. student and close colleague at the Laboratory of Actuarial Mathematics, with whom I have had a host of fruitful discussions. I also wish to thank Martin Jacobsen for his kindness and willingness to answer technical questions, Thomas Møller for his encouragement and guidance in the field of *quadratic hedging approaches*, and Jørgen Steen Larsen for his kindness and willingness to share his extensive knowledge about and insight in the insurance practice.

Furthermore, I want to thank my family and friends for general support, encouragement, understanding, and patience during the past three years. Special thoughts go to my girlfriend, Signe, and our little son, Josias.

Finally, I wish to thank for financial support from the Danish Society of Actuaries for my participation in "The Nordic Summer School in Insurance Mathematics 2003", and from Knud Højgaards Fond and Christian og Ottilia Brorsons Rejselegat for my stay in London.

Summary

This thesis deals with certain financial optimization problems connected to life and pension insurance. We take the perspective of a life insurance company or a pension fund that aims to optimize certain strategies related to the design and management of life and pension insurance policies, in particular investment and bonus strategies, on behalf of the policyholders and/or on behalf of the company itself (i.e., its owners).

Traditional life insurance mathematics is an actuarial discipline mainly concerned with qualitative and quantitative analyses of various aspects of risks connected to life and pension insurance policies, in particular modelling and assessment of mortality, calculation of present values and (equivalence) premiums, etc. The tools and principles for management of (portfolios of) life and pension insurance policies are to some extent also subjects of this field, but they are traditionally derived only implicitly from the abovementioned analyses or on the basis of more or less *ad hoc* considerations, that is, specific optimization criteria are rarely considered.

In this thesis we formulate and treat explicit optimization problems that formalize certain objectives of a life or pension insurance company arising from the natural desire to optimize the management, in a broad sense, of the portfolio of policies and the company as a whole. We apply and extend methods from mathematical finance, which is a mathematical/economical discipline concerned with various aspects of financial markets, in particular optimization of investment strategies in different respects.

The first main chapter of the thesis, Chapter 2, presents a survey of the literature on optimal investment theory. We attempt to cover most of the classical problems and approaches; we begin with the single-period analyses dating back to the beginning of the 1950's and plough our way through the literature chronologically. The amount of research conducted in this field has increased dramatically over the past 10-15 years, and we only go through a fraction of the modern literature. The theory of optimal investment obviously has substantial relevance in life and pension insurance (in particular pension saving), but we point out that the application of this theory in life and pension insurance by no means is trivial in general.

Chapters 3 and 4 deal with optimization problems concerning, in a broad sense, optimal redistribution of the systematic surplus generated by a generic so-called

participating life insurance policy, although other interpretations are possible in the model of Chapter 4. The problems are in some sense similar in nature, but the chapters represent quite different approaches. In Chapter 3 we work within a traditional setup of life insurance mathematics; in particular we do not consider different investment possibilities. In Chapter 4 we consider a fairly general model that explicitly includes a financial market model, and we propose an overall general approach to the optimal management of life and pension insurance policies. In both chapters we take the policy risks explicitly into account, but we treat them differently.

The last chapter addresses an important problem that arises in connection with the application of methods of mathematical finance for optimal investment in life and pension insurance. Life and pension insurance policies are often very long-term contracts, and from a practical point of view it is therefore quite unrealistic to assume the financial market to be *complete* over the full term of the policies since real-world markets in particular typically only offer bonds (and other interest rate derivatives) of a limited duration. The problem of optimal investment in life and pension insurance is further complicated by the fact that long-term minimum guarantees, for which the natural hedging derivatives are bonds, often are included in the policies. We consider a purely financial (i.e., no policy risk involved) optimization problem of a long-term investor in a fairly simple model of a financial market that exhibits the lack of long-term bonds (and other interest rate derivatives) in real-world markets.

Resumé

Denne afhandling behandler visse finansielle optimeringsproblemer knyttet til livs- og pensionsforsikring. Vi tager udgangspunkt i perspektivet for et livsforsikringsselskab eller en pensionskasse, der sigter mod at optimere visse strategier relateret til designet og styringen af livs- og pensionsforsikringspolicer, specielt investerings- og bonusstrategier, på vegne af policetagerne og/eller på vegne af selskabet selv (dvs. dets ejere).

Traditionel livsforsikringsmatematik er en aktuarmæssig disciplin, der hovedsageligt beskæftiger sig med kvalitative og kvantitative analyser af forskellige aspekter af risici knyttet til livs- og pensionsforsikringspolicer, specielt modellering og vurdering af dødelighed, beregning af nutidsværdier og (ækvivalens-)præmier, etc. Værktøjerne og principperne for styring af (porteføljer af) livs- og pensionsforsikringspolicer er til en vis grad også emner inden for dette felt, men de udledes traditionelt kun implicit fra ovennævnte analyser eller på basis af mere eller mindre *ad hoc* betragtninger, dvs. specificikke optimeringskriterier betragtes sjældent.

I denne afhandling formulerer og behandler vi eksplicitte optimeringsproblemer der formaliserer visse mål for et livsforsikringsselskab eller en pensionskasse, som opstår fra det naturlige ønske om at optimere styringen, i en bred forstand, af porteføljen af policer og selskabet som helhed. Vi anvender og udvider metoder fra matematisk finansiering, som er en matematisk/økonomisk disciplin, der beskæftiger sig med forskellige aspekter af finansielle markeder, specielt optimering af investeringsstrategier til forskellige formål.

Det første hovedkapitel i afhandlingen, Kapitel 2, præsenterer en oversigt over litteraturen om optimal investeringsteori. Vi forsøger at dække de fleste klassiske problemer og tilgange; vi begynder med en-periode analyserne, der daterer sig til begyndelsen af 1950'erne, og pløjer os gennem litteraturen kronologisk. Mængden af forskning der udføres inden for dette felt er vokset dramatisk de sidste 10-15 år, og vi gennemgår kun en del af den moderne litteratur. Teorien om optimal investering har naturligvis betydelig relevans for livs- og pensionsforsikring (specielt for pensionsopsparing), men vi gør opmærksom på, at anvendelsen af denne teori i livs- og pensionsforsikring på ingen måde er triviel generelt.

Kapitel 3 og 4 behandler optimeringsproblemer, der vedrører optimal tilbageførsel, i en bred forstand, af det systematiske overskud, der genereres af en livsforsikringspolice med bonusret, omend andre fortolkninger er mulige i modellen i Kapitel 4. Problemerne ligner i en vis forstand hinanden, men kapitlerne repræsenterer

ret forskellige tilgange. I Kapitel 3 arbejder vi inden for en traditionel opsætning i livsforsikringsmatematik; specielt betragter vi ikke forskellige investeringsmuligheder. I Kapitel 4 betragter vi en temmelig generel model, som eksplicit inkluderer en model af et finansielt marked, og vi foreslår en overordnet tilgang til optimal styring af livs- og pensionsforsikringspolicer. I begge kapitler tager vi policerisiciene eksplicit i betragtning, men vi behandler dem forskelligt.

Det sidste kapitel adresserer et vigtigt problem, der opstår i forbindelse med anvendelsen af metoder fra matematisk finansieringsteori til optimal investering i livs- og pensionsforsikring. Livs- og pensionsforsikringspolicer er ofte meget langsigtede kontrakter, og fra et praktisk synspunkt er det derfor ret urealistisk at antage, at det finansielle marked er *fuldstændigt* over policernes fulde løbetider, eftersom virkelighedens markeder specielt kun udbyder obligationer (og andre rentederivater) af begrænset varighed. Problemet om optimal investering i livs- og pensionsforsikring kompliceres yderligere af det forhold, at langsigtede minimumsgarantier, hvortil de naturlige afdækningsinstrumenter er obligationer, ofte er inkluderet i policerne. Vi betragter et rent finansielt (dvs. ingen policerisiko involveret) optimeringsproblem for en langsigtet investor i en ret simpel model for et finansielt marked, som udstiller manglen på lange obligationer (og andre rentederivater) på virkelighedens markeder.

Contents

Chapter 1

Introduction

This introductory chapter briefly relates the problems considered in this thesis to relevant fields of research in a fairly nontechnical manner. Furthermore, the nature and scope of the considered problems is clarified, without attention to details. Finally, an overview of the contents of the thesis is provided.

1.1 A brief comparison with related fields of research

The topics of this thesis lie at the interface between life insurance mathematics and mathematical finance. These two fields of research have, until quite recently, evolved more or less independently of one another.

Classical topics of life insurance mathematics (see, e.g., Norberg (2000) or Gerber (1997)) are, in broad terms, modelling and estimation of *policy risks* (i.e., mortality risk and other risks associated with the random course of life of the insured under a life or pension insurance policy), the study of particular forms of policies, calculations and characterizations of reserves (and higher order moments of present values), premium calculations, etc.

Explicit modelling of the financial market in which a policy is issued is typically not an issue that is addressed in basic life insurance mathematics, where only a specific model of the policy risk is considered. As regards the financial aspects of life insurance, the traditional overall approach (to so-called *participating* or *with-profit* life insurance policies, at least) has been to simply base the calculation of premiums on a prudent (pessimistic) assessment of the future interest rates, the idea being that the actual yields from investments should be sufficient to cover the liabilities under virtually any conceivable economic development over the policy terms, see e.g. Norberg (1999) (the traditional approach regarding the aspects related to the policy risks is fairly similar, but we do not go into details since we mainly focus on the financial aspects). Thus, a formalization of investment strategies employed in relation to the (portfolio of) policies has traditionally not been studied as a topic in basic life insurance mathematics.

This approach raises several natural questions: How should life and pension insurance companies invest in the financial market? More generally, how should they manage their policies? How should the systematic surplus that typically emerges over the policy terms (due to the prudent assessment of future interest rates) be redistributed to the policyholder(s)? What criteria should be imposed to ensure that policies are *fair*? What happens if the actual yields from investments are insufficient to cover the liabilities? Can this be avoided with certainty? A more general and fundamental question along these lines is: How should life and pension insurance policies be designed, and how should they be managed? Let us interpose the remark that some of these questions are very much of current interest in the practical world of life and pension insurance.

Classical topics of mathematical finance (see, e.g., Karatzas and Shreve (1998) or Duffie (1996)) are, in broad terms, modelling and general studies of financial markets, optimal investment theory, pricing and hedging of various derivatives of the basic assets in the market, equilibrium studies, etc. In particular, specific models of financial markets, and investment strategies for various purposes, are studied explicitly.

Life and pension insurance policies are, of course, financial contracts, and several well-understood methods of mathematical finance are thus quite naturally very well suited for providing sound answers to many questions along the lines of the ones posed above. It is therefore comforting that the field of research concerned with various aspects of the integration of these methods in life and pension insurance is rapidly developing. This thesis contributes to this field by applying and extending methods of mathematical finance to optimization problems in life and pension insurance.

However, there are still many open problems in this field, but it should be stressed that the integration of mathematical finance and life insurance mathematics should be conducted with care. Although life and pension insurance policies are financial contracts, they are in general not just financial derivatives of the basic assets in the financial market. Moreover, they typically have certain features that make valuation, hedging, and optimization problems connected to them highly nontrivial. There is plenty of research left to be done.

It is worth mentioning that a parallel development currently takes place in the practical world of life and pension insurance. New types of policies, which to a larger and larger extent incorporate features of other financial products, are formed and put on the market. New legislative demands requiring the valuation of both assets and liabilities to become market based are issued, and so on.

From an accounting perspective it can be argued that the methods of life insurance mathematics primarily have addressed the issues pertaining to the liability side of a life and pension insurance company, whereas the methods of mathematical finance primarily have addressed the issues pertaining to the asset side. The integration of these two fields promotes the understanding of and development of tools for the management of the total economy of the company.

1.2 General scope and aims of the thesis

Life and pension insurance policies exist in numerous forms and shapes across (and within) different countries throughout the world. This variety reflects various differences in the traditional approaches to life and pension insurance and the development of these approaches in various directions.

There are, however, and quite naturally so, also similarities. This thesis addresses some financial optimization problems of a life insurance company or pension fund (in the following referred to as *the company*) in connection with a wide class of life and pension insurance policies. The (types of) policies we have in mind are characterized by the following basic features, stated here in fairly general terms: A certain stream of (nonnegative) premiums, to be paid by the policyholder, and a certain stream of basic (nonnegative) benefits, to be paid by the company, are laid down in agreement between the two parties upon issue of the contract. In addition to the basic benefits, the company has an obligation to pay a stream of (nonnegative) bonus benefits. The premium payment stream and the basic benefit stream may (and usually do) depend on the random course of the (life of the insured under the) policy, but they are laid down upon issue and cannot be changed at a later stage (unless some *intervention option* is exercised; see Remark 1.2.1 below). In particular, the basic benefits cannot be reduced once the contract has been issued, and they are thus often referred to as *guaranteed* or *contractual (minimum)* benefits. In contrast, the bonus benefits are not stipulated in the policy, and they may (and usually do) depend not only on the random course of the policy, but also in various ways on exogenous factors such as the economic development.

Obviously, some sort of constraint involving an assessment of the total payments connected to the policy needs to be imposed in order to ensure that it is a *fair* contract in a suitable sense, but we shall not discuss such constraints here.

This somewhat loose characterization involves the features of the policies rather than particular designations of such policies. This is deliberate, since the terminology of life and pension insurance constitutes one of several aspects of the above-mentioned variety and thus, unfortunately, is not uniform. To avoid confusion we therefore also make clear, already at this stage, that we use the (admittedly somewhat lengthy) term *life and pension insurance* in a rather broad sense, that is, we do not have any particular type of life insurance products or pension schemes in mind when we use this term.

Nevertheless, we shall briefly mention some types of policies belonging to this class, using the normal designations. Traditional *with-profit* or *participating* life insurance policies constitute perhaps the main example, and they are, indeed, the ones we typically have in mind. For such policies the guaranteed benefits and the premiums are usually settled according to the actuarial *equivalence principle* (which states that the expected present value of premiums and benefits should be equal) under technical, prudent (pessimistic) assumptions (see e.g. Norberg (1999) for further details). The so-called *defined contribution* policies constitute another example. In their purest form they are characterized by the fact that only a stream of premium payments is stipulated in the policy. This simply means that in the terminology used above, there are no guaranteed benefits, and all benefits are therefore bonus benefits. There are also *defined contribution policies with minimum guarantees*, which are quite similar to with-profit policies. Similarly, unit-linked policies (and in particular unit-linked policies with guarantees) are also included. Such policies are characterized by the fact that the benefits are linked explicitly to some financial index. In comparison, the benefits paid out from a traditional with-profit or defined contribution policy typically depend on the investment performance of the company, but are not linked to a specific financial index.

Various optimization problems related to the management of such policies can be imagined. We shall concern ourselves with problems related mainly to optimization, in a certain sense, of the retirement or pension benefits, which constitute the main part of the (expected) benefits in most pension schemes. Two types of strategies will be involved:

Since the bonus benefits are not stipulated in such policies, it is — more or less — up to the company to decide how the stream of bonus benefits should be formed, and the number of possibilities is virtually unlimited. Optimization of the stream of bonus benefits, i.e., the *bonus strategy*, constitutes a main issue of this thesis. The company of course has access to the financial market, but the policyholder typically has no (direct) influence on the investment strategy used by the company in connection with the policy. Optimization of the investment strategy of the company in relation to the policy, in a broad sense, constitutes another main issue of this thesis. However, as we shall argue (in particular in Chapter 4), the bonus and investment strategies, and the optimization problems involving them, are intimately related. Moreover, the problems are more complex in their nature than corresponding purely financial problems because of the policy risk. We shall in general take the fundamental view throughout this thesis that the company and the policyholder are two different economic agents and cannot in general be thought of as one. In other words, the company is not merely an investment fund for the policyholder in general.

As mentioned above we shall always take the stream of premiums as given. In particular, problems related to so-called *defined benefit* policies will not be considered (such policies are quite different from the abovementioned types of policies; they are characterized by the fact that only the guaranteed benefits are stipulated upon issue of the policy, the premiums are then adjusted as time passes, and there are no bonus benefits). Moreover, we do not discuss microeconomical problems such as the problem of determining the (optimal) amounts of premiums to be paid by an individual to an insurance company in order to save for retirement.

Remark 1.2.1 Most policies implicitly include certain natural *intervention options* that belong to the policyholder, in particular the *surrender* and *free policy* options. These options allow the policyholder to stop paying premiums at any given time during the policy term. In loose terms, if the surrender option is exercised the policyholder receives an amount corresponding to the value of the policy at the time in question, and the policy is cancelled. If the free policy option is exercised, the remaining benefits are adjusted to account for the fact that no further premiums are paid, but the policy remains in force. Some policies also include other intervention options allowing the policyholder to e.g. change the profile of the future benefits without stopping the premium payments. We shall in general ignore intervention options in this thesis (although we do provide some remarks on how the can be built into the model of Chapter 4 in Section 4.8). It can be argued that, when properly dealt with, they constitute an issue of minor importance. We refer to Steffensen (2002) for a general treatment of the surrender and free policy options.

1.3 Overview and contributions of the thesis

Apart from this introductory chapter, the thesis consists of four chapters and three appendices. Each chapter is self-contained and can be read independently of the other chapters. This has the consequence that certain basics of the models and problems studied in the chapters may appear more than once, but the subjects of the chapters vary enough to make this a minor issue.

In Chapter 2 we provide a survey of classical and recent results from the financial theory of optimal investment. We begin, for convenience, with a presentation of a fairly general model of a financial market and several important concepts of mathematical finance. We then discuss some general optimization objectives before moving on to the survey itself.

Chapter 3 deals with a bonus optimization problem based on the modelling framework from Norberg (1999). The company is assumed to invest exclusively in an asset that yields stochastic interest modelled by a finite-state continuous-time Markov chain that particularly allows the interest rate to become lower than the technical interest rate used in the calculation of premiums and technical reserves (some generalizations are considered in Section 3.6). The optimization problem concerns optimal surplus redistribution in the form of *dividends*, which, in short, represent internal reallocations of liabilities from the so-called *dividend reserve*, which can be viewed as a shared account owned by the portfolio of policyholders as a whole, to the individual reserve (account) of the policyholder. We address the problem by the method of *dynamic programming* and obtain an explicit solution. The underlying policy is a general multi-state policy, but we consider the so-called *mean surplus*, which does not depend on the policy state.

In Chapter 4 we consider a combined bonus and investment optimization problem in a fairly general modelling framework. The company can invest in a general financial market driven by a multidimensional Brownian motion. The valuation of both assets and liabilities is market based, but the policy risk is assumed to be unhedgeable. The company aims to optimize the bonus payment stream on behalf of the insured, according to a fairly general (envisaged) preference structure of the latter. We take a very general approach by optimizing over all strategies satisfying only the basic fairness constraint. Again, the actual bonus payments are formed through the distribution of dividends, and we obtain in particular a very general result on dividend optimization. The optimal bonus strategies are obtained in semi-explicit form. By the approach taken in this chapter, optimization of the bonus process alone does not lead to a satisfactory solution of the overall problem faced by the company, as they are left with an unhedgeable payment stream. To get a unified approach we therefore propose to deal with this payment stream by use of quadratic hedging strategies.

The financial market considered in Chapter 4 is, by construction, assumed to be complete (a discussion of a generalization to the incomplete-market case is considered in Section 4.8, though), which is quite a strong assumption, in particular since life and pension insurance policies typically have very long terms. In Chapter 5 we consider a purely financial investment problem of a long-term investor who invests in a financial market without long-term bonds (or other interest rate derivatives). Of course we have in mind an individual saving for retirement. We particularly consider optimization problems with terminal wealth constraints, but we only obtain general results in the unconstrained cases.

Appendix A contains various aspects of utility that are (considered to be) of interest in relation to this thesis. In particular, we focus on the foundations of the use of utility theory in decision making.

Appendix B contains some results from convex analysis, which is a mathematical discipline concerned with optimization (minimization) of convex functions under various constraints.

Finally, for the sake of completeness, Appendix C lists the definition and basic properties of elliptical distributions, which are mentioned in Chapter 2 but often do not appear in standard textbooks on probability theory.

Chapter 2

A Survey on Investment Theory

2.1 Introduction

Investment theory constitutes one of the main topics in finance. At its core, it deals with optimization of investment strategies for one or more investors or *agents*. Much advice on this issue can be (and is) provided on the basis of common sense rather than scientific studies. For example, it is a widely accepted popular advice nowadays that investments in risky assets should somehow be diversified so as to reduce risk. However, common sense and intuition can only provide (correct) answers to relatively simple questions, and often only in terms of general guidelines rather than specific unambiguous instructions.

Scientific studies on the subject aim at (more) precise statements based on sound economic principles. In mathematical finance, an investment problem is formulated in terms of some mathematical model of the "investment universe" that the investor faces, including a financial market comprising the available investment assets, a class of admissible investment strategies, a specification of exogenous income and/or expenditure, and a criterion for optimization. These ingredients constitute a setup, which, as we shall see in the following, can take various forms. Different setups may produce different answers. What is common, though, is that mathematical models can yield precise answers that, within the model, are indisputable.

This chapter presents a survey of some of the classical and also some of the more modern financial literature on optimal investment strategies. The research in this field has grown substantially and has been developed in various directions, in particular over the past few decades. With reference to the overall aims and scope of this thesis we make no attempt to cover all results and aspects. We take a certain route through the literature, which, admittedly, is chosen somewhat subjectively and perhaps to some extent reflects the personal interests of the author. However, one aim has been to cover at least the "classical" and most fundamental results, and to explore some generalizations, with a particular focus on aspects that are considered relevant in connection with life and pension insurance. Many results and aspects will be discussed only scarcely or not at all, even though they may be very important in other respects. One such aspect is the entire issue of model selection and (parameter) estimation from available financial data. In spite of its tremendous importance in theory as well as in practice it is considered to be outside the scope of this thesis. The survey is chronological, and comments and remarks are provided along the way.

Before turning to the actual survey we offer a few technical remarks on the contents in this chapter. There are, of course, many different styles of notation in the literature. For the sake of convenience and legibility of this chapter we choose to work with a common terminology and notation (introduced below) throughout this chapter, although this has the obvious drawback that our notation will often differ — in some cases substantially — from the notation in the original works. To ease the presentation we furthermore take the liberty of sometimes leaving out certain mathematical technicalities, and we do not always bring up underlying implicit assumptions unless they need to be discussed explicitly in order to stress a particular point. For basic terminology from the theory of stochastic processes, stochastic integrals, etc., we refer to, e.g. Protter (1990) or Jacod and Shiryaev (2003).

Vectors are always interpreted as column vectors. The transpose of a vector or matrix A is denoted by A'. The d-dimensional unit vector $(1, \ldots, 1)' \in \mathbb{R}^d$ is denoted by 1_d. For any set of d quantities denoted by a common base symbol equipped with indices $1, \ldots, d$, the base symbol with no index denotes the d-dimensional (column) vector composed of the individual quantities when no confusion can arise, e.g., for given $x^1, \ldots, x^d$, x denotes the vector $(x^1, \ldots, x^d)'$. For a d-dimensional vector x, $D(x)$ denotes the $d \times d$ diagonal matrix with the components of x in the diagonal.

We refer to Appendix A for unexplained notions related to utility theory.

2.2 A general financial market model

A. Prefatory remarks.

We shall deal exclusively with theoretical studies of (optimal) investment strategies in given mathematical models of financial markets, and the general theory of mathematical finance therefore plays an important role. In order to make the survey self-contained, and for convenience later on, we shall here present a fairly general financial market model and the most fundamental concepts associated with it. Moreover, we discuss the model and some of its properties in Paragraph D below. For proofs and further details we refer to the references that appear below and to the ever growing financial literature in general.

B. The financial market model.

In general, we consider a model of a financial market with a fixed finite time horizon $T > 0$. There is a *(locally) risk-free* asset, and there are $n \in \mathbb{N}$ basic *risky* assets (see Remark 2.2.2 below for an elaboration of this terminology). The future development of the market is uncertain. Therefore, the prices-per-share of the assets are modelled by stochastic processes defined on an underlying filtered probability space $(\Omega, \mathcal{F}, \mathbb{F}, \mathrm{P})$. They are denoted by S_0 and $S_1, \dots, S_n$, respectively, and assumed to be optional and càdlàg. The filtration $\mathbb{F} = (\mathcal{F}_t)_{t \in [0,T]}$ represents the flow of information from the market. In general the price processes are not the only processes related to the market, there may be other, so-called *factor* processes as well (they need not be visualized here, though).

On a technical note, $\mathbb{F}$ is assumed to satisfy the *usual conditions*, i.e., it is right-continuous, and $\mathcal{F}_0$ contains all P-null sets of $\mathcal{F}$ (see Paragraph D below for a comment). We assume, furthermore, that $\mathcal{F}_0$ is trivial, i.e., it only contains sets of measure 0 or 1, meaning that the initial states of all processes under consideration are known.

Prices are denoted in some common unit; usually (but not necessarily) a monetary currency. We use i and j as indices for the assets, and it is thus always implicitly understood that $i, j \in \{0, \dots, n\}$. For simplicity we shall in general take dividends (on stocks etc.) to be included in the price processes so that only the characteristics of the actual price processes matter. In particular, the *return* (or *rate of return*) of the i'th asset over the period $(s, t]$ is $(S_i(t) - S_i(s))/S_i(s)$.

Remark 2.2.1 We shall also report a few results from models with an infinite time horizon and from models, where no (locally) risk-free asset exists. However, to keep things fairly simple in this section, we shall stick to the model introduced above. The extension to models with the mentioned properties is fairly straightforward, although not trivial in all respects.

Remark 2.2.2 Although the usage of the descriptive terms "risky" and "(locally) risk-free" is fairly standardized in the financial literature, the exact implications are by no means obvious from their linguistic meanings: Usually, each asset is said to be either risky or (locally) risk-free, which would seem to indicate that these properties be of a global nature. However, what they really refer to are the *local* characteristics of the asset price process in question at any point in time as long as the asset is still on the market (loosely speaking, an asset is locally risky (resp. locally risk-free) if, at the point in time under consideration, the instantaneous return on the asset over the following (infinitesimally) short time interval is uncertain (resp. known with certainty)). An asset that is locally risky is thus usually called "risky" even if it is "terminally risk-free" in the sense that it has a deterministic terminal value at the time of maturity (the obvious example is a zero coupon bond in a market with a stochastic interest rate), whereas an asset that is (locally) risk-free is sometimes simply called "risk-free", i.e., without the parenthesized qualifier "locally", even if it is "terminally risky".

We shall use the following convention: All assets that are locally risky as long as they are on the market will commonly be called "risky assets". When appropriate, we distinguish between those that are terminally risk-free and terminally risky in the above sense; the former will be called "zero coupon bonds" whereas the latter will often be called "stocks" unless specific derivative names are appropriate. The (locally) risk-free asset will sometimes be called the "bank account" or the "money market account". Unless otherwise is stated, its price process S_0 is driven by an interest rate process, which we denote by r (exceptions occur only in the most general continuous-time models where S_0 need not be absolutely continuous wrt. Lebesgue measure).

We assume that all price processes are nonnegative (this is not crucial, but it eases the presentation). Moreover, we assume that S_0 is continuous, of finite variation, and strictly positive. We can then define the *discounted* price processes $\widetilde{S}_0, \ldots, \widetilde{S}_n$ by $\widetilde{S}_i = S_i/S_0$, $i = 0, \ldots, n$; in particular, $\widetilde{S}_0 \equiv 1$. As is well known (and explained below), the discounted price processes play an important role in financial mathematics. We shall work under the following fairly general modelling assumption.

Assumption 2.2.3 *The discounted price processes $\widetilde{S}_0, \ldots, \widetilde{S}_n$ are locally bounded semimartingales.*

Note that, due to the regularity assumptions imposed on S_0, the undiscounted price processes $S_0, \ldots, S_n$ are also locally bounded semimartingales.

C. Investment strategies and related concepts.
We assume that the market is frictionless, i.e., there are no trading costs, no short-sale constraints, the supply of all assets is unlimited, etc. (see Paragraph D below for some comments). An investment strategy (of some investor who trades in the market) is represented by a predictable stochastic portfolio process $(\eta_0, \eta) = (\eta_0(t), \eta(t))_{t\in[0,T]}$, which is integrable with respect to $(\widetilde{S}_0, \ldots, \widetilde{S}_n)'$; $\eta_0(t) \in \mathbb{R}$ and $\eta(t) = (\eta_1(t), \ldots, \eta_n(t))' \in \mathbb{R}^n$ denote the *number of shares* of the money market account and the risky assets, respectively, held by the investor at time t. The *value process* or *wealth process* associated with a portfolio process (η_0, η) is naturally given by

$$X^{(\eta_0,\eta)}(t) = \eta_0(t)S_0(t) + \eta(t)'S(t),\ t \in [0, T],$$

where $S = (S_1, \ldots, S_n)'$. Its discounted counterpart $\widetilde{X}^{(\eta_0,\eta)}$ is defined similarly, with S_i replaced by $\widetilde{S}_i$, $i = 0, \ldots, n$. The *discounted gains process* $\widetilde{G}^{(\eta_0,\eta)}$ is defined by

$$\widetilde{G}^{(\eta_0,\eta)}(t) = S_0^{-1}(t) \int_{(0,t]} \left(\eta_0(s)\, dS_0(s) + \eta(s)'\, dS(s)\right),\ t \in [0, T];$$

it measures the discounted value of the accumulated investment gains from the portfolio over the interval $(0, t]$.

Definition 2.2.4 A portfolio process (η_0, η) is said to be *financed* by the semimartingale $\Gamma^{(\eta_0,\eta)}$ given by

$$\Gamma^{(\eta_0,\eta)}(t) = S_0(t)\left(\widetilde{X}^{(\eta_0,\eta)}(t) - \widetilde{G}^{(\eta_0,\eta)}(t)\right),\ t \in [0, T],$$

and it is called *self-financing* if

$$\Gamma^{(\eta_0,\eta)}(t) = 0,\ \forall t \in [0,T].$$

A portfolio process (η_0, η) is said to be *admissible* if the discounted gains process $\widetilde{G}^{(\eta_0,\eta)}$ is almost surely bounded from below.

Remark 2.2.5 Points of increase (decrease) of the financing process $\Gamma^{(\eta_0,\eta)}$ of Definition 2.2.4 correspond to points of capital injection (withdrawal) into (from) the wealth process. A portfolio process (η_0, η) with a given initial value, say x, and with no further exogenous withdrawals or injections of capital, such that the value function develops as

$$dX^{(\eta_0,\eta)}(t) = \eta_0(t)\, dS_0(t) + \eta(t)'\, dS(t),\ t \in [0,T],$$

is often called self-financing in the financial literature. However, according to Definition 2.2.4, such a portfolio process is said to be financed by the process $\Gamma^{(\eta_0,\eta)} \equiv x$, and we shall simply call it x-financed. We thus use the term "self-financing" only if $x = 0$.

For a *given* financing process $\Gamma(\cdot)$, any admissible portfolio process (η_0, η) is completely specified by η, (and the price processes), since it can be shown that the discounted gains process satisfies

$$\widetilde{G}^{(\eta_0,\eta)}(t) = \int_{(0,t]} \left(\Gamma(s)\, dS_0^{-1}(s) + \eta(s)'\, d\widetilde{S}(s) \right),\ \forall t \in [0,T],$$

so that η_0 is determined residually for every $t \in [0,T]$ by

$$\begin{aligned} \eta_0(t) &= \widetilde{\Gamma}(t) + \widetilde{G}^{(\eta_0,\eta)}(t) - \eta(t)'\widetilde{S}(t) \\ &= \widetilde{\Gamma}(t) + \int_{(0,t]} \left(\Gamma(s)\, dS_0^{-1}(s) + \eta(s)'\, d\widetilde{S}(s) \right) - \eta(t)'\widetilde{S}(t), \end{aligned} \tag{2.1}$$

where, of course, $\widetilde{\Gamma} = S_0^{-1}\Gamma$. It is seen that the right-hand side of (2.1) does not depend on η_0.

Furthermore, whenever $X^{(\eta_0,\eta)}(t) \neq 0$, the portfolio process is also uniquely determined by its corresponding value $X^{(\eta_0,\eta)}(t)$ and the *relative portfolio weights* or *proportions* $w_1(t), \ldots, w_n(t)$ in the risky assets, defined by

$$w_i(t) = \frac{\eta_i(t) S_i(t)}{X^{(\eta_0,\eta)}(t)}.$$

The relative portfolio weight in the money market account is then $1 - \sum_{i=1}^{n} w_i(t)$. These are often easier to work with, and we shall do so quite a bit in the following (note that since short-sales will be allowed, the relative portfolio weights are not restricted to be in the interval $[0, 1]$). However, the condition $X^{(\eta_0,\eta)}(t) \neq 0$ should not be ignored.

For an admissible, self-financing portfolio process (η_0, η), the (discounted) value process always equals the (discounted) gains process, and we have

$$\widetilde{X}^{(\eta_0,\eta)}(t) = \widetilde{G}^{(\eta_0,\eta)}(t) = \int_{(0,t]} \eta(s)'\, d\widetilde{S}(s), \ \forall t \in [0, T].$$

A fundamental concept in modern mathematical finance is the notion of *arbitrage.*

Definition 2.2.6 An *arbitrage opportunity* is an admissible, self-financing portfolio process (η_0, η), such that the discounted terminal value $\widetilde{X}^{(\eta_0,\eta)}(T)$ satisfies

$$\widetilde{X}^{(\eta_0,\eta)}(T) \geq 0, \text{ a.s., and } \mathrm{P}\left(\widetilde{X}^{(\eta_0,\eta)}(T) > 0\right) > 0. \tag{2.2}$$

Thus, an arbitrage opportunity is an investment strategy that requires no initial capital but generates a nonnegative payoff at time T, which is strictly positive with strictly positive probability, i.e., a risk-free “money machine”. We shall assume throughout, as is standard, that the market does not allow arbitrage opportunities (we use the abbreviation “NA” for this assumption). Real-world markets may at times allow for small-scale arbitrage opportunities, but a theoretical optimization problem cannot be well posed if arbitrage is allowed.

As is well known, this basic, purely economical assumption is deeply related to certain mathematical properties of the financial market model, as it is “essentially” equivalent to the existence of an *equivalent local martingale measure*, which is a probability measure Q that is equivalent to P and turns all the discounted price processes $\widetilde{S}_0, \ldots, \widetilde{S}_n$ into local martingales (see Paragraph D below for more precise statements). In particular, this allows for “pricing by (no) arbitrage” of *derivatives* (options, futures, etc.) of the basic assets (at least if there is only one equivalent local martingale measure), which constitutes one of the most important tools of the theory. We denote by $\mathbb{P}$ the set of equivalent local martingale measures as defined above.

Other fundamental concepts are the notions of *contingent claims*, *hedging strategies*, and *completeness.*

Definition 2.2.7 A *simple contingent claim* is a nonnegative $\mathcal{F}_T$-measurable random variable. A given simple contingent claim B is said to be *attainable* or *hedgeable*, if there exists an $x \in \mathbb{R}$ and an admissible x-financed strategy (η_0, η) such that

$$X^{(\eta_0,\eta)}(T) = B, \text{ a.s.}$$

Such a strategy is called a *hedging strategy*, and it is said to *hedge* or *generate* B. Finally, if $\mathbb{P} \neq \emptyset$, then the market is said to be *complete* if every simple contingent claim B satisfying the technical condition

$$\sup_{\mathrm{Q}\in\mathbb{P}} \mathrm{E}^{\mathrm{Q}}\left(B/S_0(T)\right) < \infty \tag{2.3}$$

is attainable.

Under the technical assumption that $\mathcal{F} = \mathcal{F}_T$ it can be shown (see Paragraph D below) that the market is complete if and only if the set $\mathbb{P}$ is a singleton, i.e., if and only if there is a *unique* equivalent local martingale measure Q. In particular, this means that derivatives can be priced uniquely.

Finally, an admissible portfolio process (not necessarily self-financing) (η_0, η) is called *simple* if each η_i has the form

$$\eta_i(t) = \sum_{k=1}^{K} \xi_i(k) 1_{(T(k),T'(k)]}(t),\ t \in [0,T],\ i = 0, \ldots, n,$$

where $K \geq 1$ and, for $k = 1, \ldots, K$, $T(k)$ and $T'(k)$ are stopping times with values in $[0, T]$, and $\xi_i(k)$ is an $\mathcal{F}_{T(k)}$-measurable random variable. Loosely speaking, a simple strategy as in Definition 2.2.4 consists in buying $\xi_i(k)$ units of the i'th asset at time $T_i(k)$ and selling them at time $T_i'(k)$, $i = 1, \ldots, n$, $k = 1, \ldots, K$.

D. Discussion of the model.

This paragraph contains a discussion of some of the fundamental assumptions and features of the financial market model presented above. Some of the issues that we broach below are often not discussed in textbooks on mathematical finance, at least not introductory ones, but they may be of significant importance and should therefore not be neglected.

Starting with the basics, let us note that the probability measure P represents (the investor's beliefs of) the real probabilities of all events under consideration, i.e., all events in $\mathcal{F}_T$. As is the case in most of the financial literature, it will be taken as given throughout this chapter. However, as we shall see below, P need not always be fully specified; in certain models (in particular discrete-time models) it may be the case that P is just assumed to be in a certain (equivalence) class of probability measures. In this case it is this class of probability measures that is taken as given.

The filtration representing the flow of information has been assumed to satisfy the *usual conditions*. This is a standard assumption in the literature (hence the designation), and its motivation is purely technical. However, since filtrations are typically generated by basic stochastic processes, and since natural filtrations of many types of processes (e.g. Brownian motions) do not satisfy the usual conditions, one typically needs to work with an augmentation of the natural filtration, which thus contains more information than that obtained by observing the basic processes. This seems unnatural, but it does not cause serious modelling concerns, because the added "unnatural" information essentially only concerns null sets and therefore cannot be exploited. We refer to the literature on stochastic processes (e.g. Rogers and Williams (1987)) for more details.

We have assumed that the (discounted) price processes are locally bounded semimartingales. In general, any stochastic process representing the price process of an asset that can be traded continuously must be a semimartingale. As such, this is a consequence of pure mathematics: Any càdlàg, adapted process is an *integrator* if and only if it is a semimartingale (see, e.g. Rogers and Williams (1987), Theorem IV.16.4), and the only way to represent gains processes in general is through integrals. If we stick to simple portfolio processes, however, then gains processes are well defined for *any* càdlàg, adapted price processes (not necessarily semimartingales), since they are just finite sums of random variables. Mathematically, it thus poses no problem to work with more general processes as price processes. Economically, however, it is a different story. To see this, we need the following definition, slightly extending the notion of arbitrage.

Definition 2.2.8 A *free lunch with vanishing risk* is a sequence of admissible, self-financing portfolio processes $(\eta_0^{(k)}, \eta^{(k)})_{k\geq 1}$ such that

$$\widetilde{X} := \lim_{k\to\infty} \widetilde{X}^{(\eta_0^{(k)},\eta^{(k)})}(T)$$

is well defined as an almost sure limit and satisfies (2.2), and such that

$$\left|\widetilde{X} - \widetilde{X}^{(\eta_0^{(k)},\eta^{(k)})}(T)\right| < 1/k, \text{ a.s., } \forall k \geq 1,$$

in particular $\widetilde{X}^{(\eta_0^{(k)},\eta^{(k)})}(T) > -1/k$, a.s., $\forall k \geq 1$.

A free lunch with vanishing risk is almost an arbitrage, as one can get arbitrarily close to the (discounted) "arbitrage payoff" $\widetilde{X}$ uniformly. Assuming that there is *no free lunch with vanishing risk* (NFLVR) in the market is therefore only slightly stronger than assuming NA. Now, if the discounted price processes are merely càdlàg, adapted, locally bounded processes, and if we just assume that there is no free lunch with vanishing risk consisting of simple portfolio processes, then it turns out that the discounted price processes must in fact be semimartingales (see Delbaen and Schachermayer (1994)). Unfortunately, the condition of local boundedness is in general necessary, but we shall not discuss this further.

The restriction that a portfolio process must be predictable is basically a mathematical way of expressing that the agent cannot be allowed to use "future information" to form his decisions, i.e., no insider trading is allowed. However, it is necessary only if the price processes are discontinuous.

The assumption that the market is frictionless, which will actually be imposed throughout this thesis, is in many ways a standard assumption in much of the financial literature (although there also exist many papers adressing problems that arise when it is relaxed in some way). It is a quite strong assumption, but the severity of it depends on the investor, and it can be argued that it is somewhat milder for a pension fund (or a mutual fund) investing on behalf of its members than it is for an individual.

The definition of admissibility for a portfolio process includes the condition that the discounted gains process must be bounded from below. With self-financing portfolio processes that do not satisfy this condition it is possible in many specific models — even simple ones such as the Black-Scholes model (see Karatzas and Shreve (1998)) — to generate payoffs that are arbitrarily large (almost surely), by use of, e.g., "doubling strategies". Clearly, allowing for this is absurd, so such portfolio processes must be excluded from consideration. This was recognized (in an informal manner) already by Harrison and Kreps (1979).

The exact connection between *absence of arbitrage* and *existence of an equivalent local martingale measure*, which we alluded to in the previous paragraph, was established by Delbaen and Schachermayer (1994): There exists an equivalent local martingale measure Q (for $(\widetilde{S}_0, \ldots, \widetilde{S}_n)$), i.e., $\mathbb{P} \neq \emptyset$, if and only if the market satisfies the condition NFLVR. Thus, the purely economical assumption NFLVR implies that there exists a probability measure Q, equivalent to P, such that $\widetilde{S}_0, \ldots, \widetilde{S}_n$ are local martingales under Q. It can then be shown that also the discounted gains process corresponding to any self-financing admissible portfolio processes is a local martingale, and therefore in particular a supermartingale (Delbaen and Schachermayer (1994), Theorem 2.9).

The equivalence of $\mathbb{P} \neq \emptyset$ and NFLVR is based on the assumption of local boundedness of the discounted price processes. Under other conditions the connection between absence of arbitrage and existence of equivalent martingale measures is slightly different, and for completeness we list the exact relationships here. In the general case, where the discounted price processes $\widetilde{S}_0, \ldots, \widetilde{S}_n$ are allowed to be general semimartingales (not necessarily locally bounded), the NFLVR condition is equivalent to the existence of an equivalent *sigma-martingale measure* Q (for $(\widetilde{S}_0, \ldots, \widetilde{S}_n)$), i.e., a probability measure equivalent to P, such that $\widetilde{S}$ has the form $\widetilde{S}(t) = \widetilde{S}(0) + \int_{(0,T]} \phi(s)\, dM(s)$ for some $[0, \infty)$-valued predictable process ϕ and some $\mathbb{R}^n$-valued Q-martingale M; see Delbaen and Schachermayer (1998) for a proof and further details.

If the discounted price processes are continuous, then they are in particular locally bounded, so the above statement on the NFLVR condition applies. In this case, however, the milder assumption NA in itself implies that there exists an absolutely continuous, but not necessarily equivalent, local martingale measure, i.e., a probability measure Q absolutely continuous with respect to P under which $\widetilde{S}_0, \ldots, \widetilde{S}_n$ are local martingales (see Delbaen and Schachermayer (1995a)).

Finally, if the discounted price processes are pure jump processes (possibly unbounded) with fixed jump times $0 < t_1 < \ldots < t_K = T$, i.e.,

$$\widetilde{S}_i(t) = \widetilde{S}_i(0) + \sum_{k: 0 < t_k \leq t} \Delta \widetilde{S}_i(t_k), \ \forall t \in [0, T], \ \forall i = 1, \ldots, n,$$

then NA is equivalent to the existence of an equivalent martingale measure, i.e., a probability measure Q equivalent to P under which $\widetilde{S}_0, \ldots, \widetilde{S}_n$ are martingales (see Schachermayer (1992) and references therein). Note that in this case the discounted price processes are automatically semimartingales, but they need not be locally bounded, in particular they are not if the jumps are i.i.d. with a distribution that has unbounded support.

Remark 2.2.9 The fact that local boundedness is not required when the discounted price processes are pure jump processes with fixed jump times $0 < t_1 < \ldots < t_K = T$ is quite important, in particular since we have, rather surprisingly, not been able to find conditions for local boundedness of such processes in the literature. A necessary condition for local boundedness of a càdlàg pure jump process Y with a fixed, finite number of jump times $0 < t_1 < \ldots < t_K = T$ is that the conditional distribution of $Y(t_k)$ given $\mathcal{F}_{t_k-}$ has bounded support almost surely, $\forall k = 1, \ldots, K$ (note that $\mathcal{F}_{t_k-}$ may be strictly larger than $\mathcal{F}_{t_{k-1}}$). To see this, let us first note that if τ is a stopping time such that the stopped process Y^τ is bounded, say $|Y(t \wedge \tau)| \leq L$, $\forall t \in [0, T]$, a.s., for some $L \in \mathbb{N}$, then, necessarily, we must have $\tau < t_k$ on the set $(|Y(t_k)| > L)$, i.e.,

$$\mathrm{E}\left(1_{(|Y(t_k)|>L,\, \tau \geq t_k)}\right) = 0,$$

for each $k = 1, \ldots, K$. We have $(\tau \geq t_k) = (\tau < t_k)^c \in \mathcal{F}_{t_k-}$, so

$$\begin{aligned} 0 &= \mathrm{E}\left(1_{(|Y(t_k)|>L,\, \tau \geq t_k)}\right) \\ &= \mathrm{E}\left(\mathrm{E}\left(1_{(|Y(t_k)|>L,\, \tau \geq t_k)} \,\middle|\, \mathcal{F}_{t_k-}\right)\right) \\ &= \mathrm{E}\left(1_{(\tau \geq t_k)} \mathrm{E}\left(1_{(|Y(t_k)|>L)} \,\middle|\, \mathcal{F}_{t_k-}\right)\right). \end{aligned}$$

This means that we must have $\mathrm{E}\left(1_{(|Y(t_k)|>L)} \,\middle|\, \mathcal{F}_{t_k-}\right) = 0$ on the set $(\tau \geq t_k)$. Now, if the conditional distribution of $Y(t_k)$ given $\mathcal{F}_{t_k-}$ does not have bounded support almost surely, then the set

$$A := \bigcap_{L=1}^{\infty} \left(\mathrm{E}\left(1_{(|Y(t_k)|>L)} \,\middle|\, \mathcal{F}_{t_k-}\right) > 0\right)$$

has strictly positive probability. Thus, if $(\tau_i)_{i \geq 1}$ were a localizing sequence of stopping times, then we would have a contradiction, because then

$$\mathrm{P}(A) = \mathrm{P}\left(A \cap (\cup_{i=1}^{\infty} (\tau_i \geq t_k))\right) = \mathrm{P}\left(\cup_{i=1}^{\infty} (A \cap (\tau_i \geq t_k))\right) = 0.$$

This proves necessity. We believe that the condition is also sufficient, but no proof has been constructed for the general case (this seems to require quite a bit of technical fuss). However, if $\mathcal{F}_{t_k-} = \mathcal{F}_{t_{k-1}}$, i.e., all information is contained in the price processes, then the condition certainly is sufficient: Put

$$\tau_i = \inf\left\{t_k \in \{t_0, \ldots, t_{K-1}\} : \operatorname{ess\,sup}\left(Y(t_{k+1}) \,|\, \mathcal{F}_{t_k}\right) \geq i\right\} \wedge t_K,\ i = 1, 2, \ldots,$$

where

$$\operatorname{ess\,sup}\left(Y(t_{k+1}) \,|\, \mathcal{F}_{t_k}\right) = \inf\{L > 0 : \mathrm{P}\left(Y(t_{k+1}) \leq L \,|\, \mathcal{F}_{t_k}\right) = 1\}.$$

Then $(\tau_i)_{i \geq 1}$ is a localizing sequence of stopping times.

Let us end this remark by noting that the condition is far from necessary for a left-continuous pure jump process to be locally bounded, since this is the case for *any* left-continuous process with right-hand limits.

There seems to be a fairly (but not totally) standardized way to define contingent claims and hedging strategies in the financial literature. However, there seems to be no universally agreed-upon definition of completeness. Our definition, which has not been encountered in the literature, is inspired by Delbaen and Schachermayer (1995b), Theorem 16, according to which a simple contingent claim is hedgeable if and only if there exists an equivalent local martingale measure $\mathrm{Q}' \in \mathbb{P}$ such that

$$\mathrm{E}^{\mathrm{Q}'}\left(X/S_0(T)\right) = \sup_{\mathrm{Q}\in\mathbb{P}} \mathrm{E}^{\mathrm{Q}}\left(X/S_0(T)\right) < \infty,$$

provided that $\widetilde{S}_0, \dots, \widetilde{S}_n$ are locally bounded semimartingales, as in the model above. This means that a contingent claim that does not satisfy (2.3) cannot be hedged under any circumstances.

For comparison, the definition of completeness in Harrison and Pliska (1981, 1983) is based on a specific equivalent martingale measure $\mathrm{Q}^* \in \mathbb{P}$, and the market is said to be complete (with respect to Q^*) if every contingent claim X satisfying $\mathrm{E}^{\mathrm{Q}^*}\left(X/S_0(T)\right) < \infty$ is hedgeable. This is somewhat unsatisfactory, because the completeness property should, ideally, only depend on the original model (under P), as is the case with our definition above. Nonetheless, the proof of the theorem in Harrison and Pliska (1983), which states that the market is complete with respect to Q^* if and only if the set of equivalent martingale measures is a singleton (i.e., Q^* is uniquely determined), can easily be modified to yield the desired property: The market is complete (according to our definition) if and only if the set $\mathbb{P}$ of equivalent *local* martingale measures is a singleton. The technical assumption $\mathcal{F} = \mathcal{F}_T$ was imposed in Paragraph C above. It is indeed necessary for this result to hold, but the essence of the result is valid also if $\mathcal{F} \neq \mathcal{F}_T$: Then the market is complete if and only if the restriction of any equivalent local martingale measure to $\mathcal{F}_T$ is unique, and since contingent claims are $\mathcal{F}_T$-measurable, the condition $\mathcal{F} = \mathcal{F}_T$ is just technical.

Let us end this discussion with some bibliographic remarks. The study of the general theory of trading in financial markets modelled by semimartingales was initiated by Harrison and Kreps (1979) and Harrison and Pliska (1981, 1983). Harrison and Kreps (1979) work with general price processes but only simple trading strategies and show, under certain square-integrability conditions, that NA is equivalent to the existence of an equivalent martingale measure, implying in particular that the price processes are semimartingales. Harrison and Pliska (1981, 1983) extend the theory to general trading strategies and discuss contingent claims pricing, completeness, etc., under the assumption that an equivalent martingale measure exists. The general theory has been refined since then (as demonstrated above); we refer to the abovementioned papers by Delbaen and Schachermayer and their references for further details.

2.3 Optimization objectives and solution techniques

A. Objectives.
In order to analyze and solve an optimization problem mathematically it is necessary to consider an optimization *objective*. Most modern portfolio problems assume objectives that can be written in the form

$$\sup_{((\eta_0,\eta),C)\in\mathcal{A}} \mathrm{E}\left(\mathcal{U}(C, X(T))\right), \tag{3.1}$$

where C and X denote, respectively, the *consumption stream* and the *wealth process* of the agent, $\mathcal{A}$ denotes some set of *admissible* pairs of consumption and investment strategies, and $\mathcal{U}$ is a (generalized) *utility* functional (taking values in $[-\infty,\infty)$) representing the agent's *preference structure* over consumption and terminal wealth. Here and throughout, a consumption stream $C = (C(t))_{t\in[0,T]}$ is a positive, increasing, and càdlàg stochastic process. As the name indicates we typically interpret $C(t)$ as the accumulated consumption of the agent until time t, but it actually just denotes the accumulated amount that has been *withdrawn* from the investment portfolio until time t.

The objective (3.1) is in accordance with the *expected utility maxim* or *hypothesis*, by which the aim is to maximize, over all *admissible* investment and consumption strategies, the expected utility of consumption and/or terminal wealth of the investor. For a discussion of the *expected utility maxim* we refer to Appendix A.

As we shall see, certain restrictions (at least technical ones) on (η_0, η) and C are imposed in almost all particular problems studied in the literature (hence the term "admissible" and the introduction of the set $\mathcal{A}$ above).

It is of course necessary to specify $\mathcal{U}$ in order to get any specific results. As a first step in this direction, $\mathcal{U}$ is typically assumed to be of the form

$$\mathcal{U}(C, X(T)) = \mathcal{V}(C) + u(X(T)), \tag{3.2}$$

where $\mathcal{V}$ (resp. u) is a functional (resp. function) measuring utility of consumption (resp. terminal wealth). Thus, utility of consumption and utility of terminal wealth are assumed to be additive.

It is a highly delicate issue to specify $\mathcal{V}$ further so as to obtain a functional that is economically reasonable as well as mathematically tractable. From an economic point of view it should, for example, be taken into consideration that, for most agents, the actual utility of consumption at some given time is, in some way or another, related to the (utility of) consumption in the past (and thus also in the future). In economic terms there is, e.g., a degree of *local substitution*, i.e., consumptions at nearby points in time are to some extent substitutes of one another. Further along these lines, some goods are durable and one thus enjoys a certain amount of utility from past consumption. It is difficult to capture such effects in a functional that at the same time is simple enough to make sufficient room for mathematical tractability. The "classical" and most widely used approach is to assume that utility of consumption is time-additive, i.e., that $\mathcal{V}$ has the form

$$\mathcal{V}(C) = \sum_{k=0}^{N} v(t_k, c_k)$$

if consumption must take place in discrete time, or, if consumption can take place continuously (in which case C is typically constrained to being absolutely continuous with respect to the Lebesgue measure with $dC(t) = c(t)\,dt$),

$$\mathcal{V}(C) = \int_0^T v(t, c(t))\,dt,$$

where, for each $t \in [0, T]$, $v(t, \cdot)$ is some utility function. This specification has the advantage that it can be highly tractable and yield explicit results. However, the assumption of time additivity implies that the utility of consumption at any given time is independent of past consumption and does not affect the utility of future consumption, which is somewhat implausible.

In the survey below we shall in general avoid dealing further with this issue by simply focusing on the results concerning investment strategies. This is partly due to the fact that, as indicated above, the form of $\mathcal{V}$ typically has certain somewhat undesirable — or at least questionable — features, but the main reason is that, as the subsequent chapters will show, this thesis deals with a certain kind of optimization problems where a "consumption stream" has a different interpretation than it has for a typical economic agent and thus motivates a particular type of preference structures.

A pure investment problem, where the objective is to maximize the expected utility of terminal wealth at time T only, is obtained as a special case of the objective (3.2) by setting $\mathcal{V} \equiv 0$. It could be argued that the objective (3.2), regardless of the specification of $\mathcal{V}$, has drawbacks similar to those imposed by the assumption of time additivity of utility of consumption as discussed above. However, a pure investment problem makes perfect economic sense, as it can be viewed, e.g., as that of an agent who has exogenously divided his wealth and future income into two parts in such a way that one part is used for consumption until time T while the other is used solely for investment purposes, e.g., so as to provide the means for consumption from time T and onwards. It is in fact particularly apt for life and pension insurance companies with policies as the ones we deal with in this thesis, where a certain stream of premium payments from the insured to the company is prespecified, and where pension benefits are to be paid only from the time of retirement.

We round off this paragraph by pointing out for the interested reader that preference structures for consumption aiming to capture some of the economic features that are ignored by the "classical" one have indeed been studied by several authors, see, e.g., Epstein and Zin (1989), Hindy and Huang (1993), and references therein.

B. Optimization techniques.
There are two main approaches to portfolio optimization in models where the investor is allowed to change his portfolio over time. The first of these is *dynamic programming*, which, for convenience and later reference, will be introduced briefly in this paragraph. The second is the *martingale method*, which will be treated in Section 2.7.

The dynamic programming approach can be used if (and only if) the price and factor processes constitute a Markov process (denoted here by Z), so that, at any time, the current state of the Z is all that matters for its conditional future distribution. Apart from this constraint the approach is (in principle) applicable to fairly general optimal investment and consumption problems in discrete as well as continuous time, but for simplicity we focus here on the case of a pure investment problem in continuous time. The basic idea is to view the initial optimization problem as part of a much larger family of corresponding optimization problems obtained by considering every point in the state space of Z (denoted here by E) as a (hypothetical) starting point of Z. The *value function* $V : [0,T] \times E \to \mathbb{R}$ is then defined as

$$V(t,z) = \sup_{(\eta_0,\eta)\in\mathcal{A}(z)} \mathrm{E}\left(u(X^{(x,t,\eta_0,\eta)}(T)) \,\middle|\, Z(t) = z \right), \ (t,z) \in [0,T] \times E,$$

where $\mathcal{A}(z)$ is the set of admissible investment strategies as viewed from (t,z), and where the wealth process has been equipped with a superscript indicating the (hypothetical) initial point (x is a component in z).

The *dynamic programming principle (DPP)* states that for any $0 \le t_1 < t_2 \le T$ and any $z \in E$,

$$V(t_1,z) = \sup_{(\eta_0,\eta)\in\mathcal{A}(z)} \mathrm{E}\left(V\left(t_2, X^{(x,t_1,\eta_0,\eta)}(t_2)\right) \,\middle|\, Z(t_1) = z \right). \tag{3.3}$$

This equation has the interpretation that a strategy is "globally optimal" if and only if it is "locally optimal" in the sense that it is optimal for the sub-problem of maximizing the conditional expected value of the value function at time t_2 from time t_1, given the state at time t_1. Thus, the original problem is decomposed into (a continuum of) "local problems" (in the discrete-time case this decomposition consists of a finite number of problems).

Solving the "local problems" requires that the value function is known, and it is virtually impossible to calculate the value of the value function from its definition. However, the DPP leads to the so-called *Hamilton-Jacobi-Bellman (HJB) equation(s)*, which is a (system of) partial differential equation(s) for the value function containing a supremum-expression involving the values of (η_0, η) (at the different points in $[0,T] \times E$). The value function (in closed form) and an optimal strategy (as a function of the derivatives of the value function) can thus (in principle) be obtained from the HJB equation. It is typically impossible to actually prove that the value function satisfies the HJB equation, however. Fortunately, one has quite strong *verification theorems*, which basically say that *if* a solution

ϕ (possibly in some generalized, non-smooth sense) to the HJB equation can be found, and if (η_0, η) is an admissible strategy that maximizes the abovementioned supremum-expression, *then*, under certain (mild) regularity conditions, the value function equals ϕ, and (η_0, η) is optimal. Note, though, that the HJB equation typically is non-linear and often impossible to solve, so the approach does not always lead to concrete results.

We shall not go further into details on the subject of dynamic programming, for which there is a vast literature. We refer the interested reader to the textbooks by Fleming and Rishel (1975), Krylov (1980), Davis (1993), Fleming and Soner (1993), Yong and Zhou (1999), and their references.

However, from an actuarial perspective it is interesting to note the similarities, as regards the value function, with the classical notion of prospective reserves in (multi-state) life and pension insurance (see, e.g. Norberg (1991)). Under the Markov assumption, which is standard, the prospective reserve is, at any time during the policy term, defined as the conditional expected value of the future payments pertaining to the contract (benefits less premiums, properly discounted), given the present state of the policy, and thus corresponds to the value function above (although, of course, the classical definition of a reserve does not involve an optimization problem). In general it is impossible to calculate a prospective reserve directly from the definition. The well-known actuarial approach to calculating them is based on the fact that the reserves corresponding to different times and different states are related in a way that is similar to the relationship given by (3.3). As in the theory of dynamic programming, this relationship leads to a (system of) differential equations from which the reserves can be obtained. In the life and pension insurance context these are the classical Thiele differential equations.

2.4 Single-period portfolio selection

In a single-period portfolio selection problem the investor chooses his portfolio at time 0 and cannot change it until time T, where the new asset prices materialize and the return on the portfolio is revealed.

A cornerstone of modern investment theory was laid by Markowitz (1952, 1959) and Tobin (1958). They consider a single-period problem and argue, with no particular assumptions about the distribution of returns of the individual assets, that one should only consider investing in portfolios that are *efficient*. An efficient portfolio is defined as a portfolio whose return has smaller variance than any other portfolio with a greater mean return, i.e., a portfolio with a greater mean return than an efficient portfolio also entails greater risk as measured by its variance. Thus, the choice of portfolio should only depend on return distributions through their means and variances.

There are infinitely many efficient portfolios (each feasible value of the variance of return corresponds to an efficient portfolio), and they constitute what is called the *efficient set* or *efficient frontier*; the latter designation is due to the fact that the efficient set constitutes a frontier (a part of the boundary) in the *feasible set* of combinations of means and standard deviations of portfolio returns. Within the efficient set, the investor should choose a portfolio that somehow reflects his risk aversion. The abovementioned boundary in the set of feasible combinations consists of all so-called *minimum-variance* portfolios, i.e., portfolios with minimum variance for a given mean. Note that a minimum-variance portfolio is not necessarily efficient.

By elimination of redundant assets (assets that are perfectly correlated with (a portfolio of) other assets) it can be assumed that the $n \times n$ variance matrix of returns of the risky assets, Σ, is non-singular. It can then be shown that any efficient portfolio in the general case can be formed as a combination of the risk-free asset and a single efficient portfolio of risky assets (a "mutual fund") with asset proportions determined by the vector

$$\tilde{w} = \Sigma^{-1}(\alpha - 1_n r), \tag{4.1}$$

where α and r are, respectively, the vector of mean returns of the risky assets and the (known) return of the risk-free asset. This result, sometimes called the *one-fund theorem* or the *separation theorem*, was obtained by Tobin (1958) (in the case $r = 0$).

An important point is that all investors who optimize according to the mean-variance criterion can use the same mutual fund (provided, of course, that they agree on α and Σ) although the fractions of wealth that each investor places in the fund may differ. In particular, this has tremendous importance for pension funds because it means that all their members can share the same risky-asset fund (if they are assumed to be mean-variance optimizers).

Remark 2.4.1 For $\tilde{w}$ to be a genuine vector of portfolio fractions for the mutual fund it is required that $\sum_{i=1}^n \tilde{w}_i = 1$, but this need not be fulfilled. However, in realistic situations (particularly if market equilibrium conditions based on the mean-variance criterion are imposed) one has $\sum_{i=1}^n \tilde{w}_i > 0$, and it is thus possible to scale $\tilde{w}$ so as to obtain a genuine mutual fund with portfolio fractions adding up to 1. For the sake of completeness we briefly comment on the case where $\sum_{i=1}^n \tilde{w}_i \leq 0$ (albeit unrealistic). If $\sum_{i=1}^n \tilde{w}_i = 0$, the net balance of the risky assets is 0. In other words, the short and long positions in risky assets balance out each other, and therefore they do not form a genuine mutual fund. If $\sum_{i=1}^n \tilde{w}_i < 0$, the net balance is negative, i.e., the short positions in risky assets dominate the long positions. In this case a mutual fund determined by $-\tilde{w}$ could be formed, and the investor would then go short in this fund (which would give him a positive initial payment allowing him to buy more of the risk-free asset). Note that in all cases the expected excess return of the fund over the risk-free interest rate, $\tilde{w}'(\alpha - 1_n r) = (\alpha - 1_n r)'\Sigma^{-1}(\alpha - 1_n r)$, is strictly positive (as it should be) because Σ is positive definite.

Remark 2.4.2 In the case with no risk-free asset one can still obtain efficient portfolios from the expression (4.1) by letting r vary as a free parameter that determines the mean return of the portfolios. To ensure that all such portfolios are genuine it is necessary (and sufficient) that $r < \bar{r} := (1_n'\Sigma^{-1}\alpha)/(1_n'\Sigma^{-1}1_n)$. It can be shown that these relative portfolios, given by

$$w(r) = \frac{\Sigma^{-1}(\alpha - 1_n r)}{1_n'\Sigma^{-1}(\alpha - 1_n r)}, \; r < \bar{r},$$

lie on a straight line in the $(n-1)$-dimensional hyperplane $\{x \in \mathbb{R}^n : 1_n'x = 1\}$. Thus, any two efficient portfolios $w(r_1)$ and $w(r_2)$ with $r_1 < r_2 < \bar{r}$ can be used as "mutual funds" in the sense that any other efficient portfolio can be written as $\xi w(r_1) + (1-\xi)w(r_2)$ for some $\xi \in \mathbb{R}$. This result is sometimes called the *two-fund theorem*, and it means that all mean-variance investors can used the same two funds. Note, however, that not all portfolios of the form $\xi w(r_1) + (1-\xi)w(r_2)$ are efficient, some are merely minimum-variance portfolios. The efficient portfolios therefore only constitute a "half-line".

Alternatively, all efficient portfolios can be written in the form $w^* + \xi\tilde{w}$ for some $\xi \geq 0$, where w^* and $\tilde{w}$ are given by

$$\begin{aligned} w^* &= \left(1_n'\Sigma^{-1}1_n\right)^{-1}\Sigma^{-1}1_n, \\ \tilde{w} &= \Sigma^{-1}\left(1_n'\Sigma^{-1}1_n\alpha - 1_n'\Sigma^{-1}\alpha 1_n\right). \end{aligned}$$

The portfolio w^* is the one with minimum variance (of all portfolios), and $\tilde{w}$ is a zero-sum portfolio (i.e., $1_n'\tilde{w} = 0$). With this parametrization the factor ξ is a direct measure of the risk of a given efficient portfolio.

Under the assumption that all investors are mean-variance optimizers and agree on the return parameters, Sharpe (1964), Lintner (1965), and Mossin (1966) developed a single-period equilibrium model of the financial market, which is now widely known as the CAPM (Capital Asset Pricing Model). A key property of this model is that any investor's portfolio of risky assets is a fraction of the market portfolio, i.e., the proportions of the risky assets held by any investor are the same as the proportions of the risky assets in the market. This is a consequence of the one-fund theorem and the simple fact that the portfolios held by all investors must constitute the market portfolio. This property, in turn, implies that any investor who believes that there is equilibrium in the market can invest optimally by simply holding a share of the market portfolio, i.e., there is no need for him to perform any estimation or calculation.

Even though it is not directly based on the expected utility maxim, the mean-variance approach is actually consistent with it in the single-period problem under certain conditions:

For an agent who acts according to the expected utility maxim, the optimal portfolio is mean-variance efficient regardless of the joint distribution of asset returns if and only if the agent has a quadratic utility function of the form

$$u(x) = x - \frac{k}{2}x^2,\ x \in \mathbb{R}, \tag{4.2}$$

for some $k > 0$. This was in fact shown already in Markowitz (1959). Furthermore, in a *given* model of asset returns, any mean-variance efficient portfolio is the optimal portfolio corresponding to some quadratic utility function of the form (4.2), so any mean-variance optimizer can be assumed to have a suitable quadratic utility function as long as one works with the given model. However, as is well known, this utility function has certain properties that are economically implausible, in particular it is increasing only for $x \leq 1/k$ and thus only makes sense if the asset returns are bounded from above so that the wealth stays below $1/k$ with certainty.

For an agent who acts according to the expected utility maxim the optimal portfolio is mean-variance efficient if

(i) the risky asset returns all have finite second order moments, and the distribution of any linear combination of the returns is uniquely determined by its mean and variance, and

(ii) his utility function is increasing, concave, and does not assume the value $-\infty$ within the range of possible return values.

Indeed, if (i) holds then the expected utility is effectively a function, say f, of the mean and the variance of the portfolio return only. If (ii) also holds then the monotonicity and concavity of the utility function imply that f is increasing (resp. decreasing) with respect to the mean (resp. the variance). Thus, if an optimal portfolio (for which f is maximized) exists, it must be mean-variance efficient. Chamberlain (1983) provides necessary and sufficient conditions on the joint distribution of the risky asset returns for (i) to hold. In particular, it holds in the case with a risk-free asset if and only if the joint distribution of the risky asset returns is a multivariate elliptical distribution with finite second order moment (see Appendix C), the main example of which is the multivariate normal distribution. It is important in this regard to note that the aptitude of multivariate elliptical distributions to model return distributions is somewhat questionable. In particular, the multivariate normal distribution has the apparent flaw (as a model for asset returns) that it implies that all risky assets have unlimited liability, i.e., the returns are not bounded from below, which is not the case for stocks, say, in reality. Furthermore, as indicated above, a necessary condition for the existence of an optimal portfolio that includes risky assets is that the utility function does not assume the value $-\infty$, which rules out, e.g., the CRRA utility functions.

Remark 2.4.3 As in the literature quoted above we have discussed the mean-variance approach in terms of the *return* of a portfolio rather than the *terminal wealth* corresponding to a portfolio, and this may seem to be inconsistent with the expected utility maxim. However, since the terminal wealth $X(T)$ in a single-period model corresponds one-to-one to the return R by the formula

$$X(T) = X(0)(1 + R),$$

where $X(0)$, of course, is the (known) initial wealth, it is easily seen that any quadratic utility function of the form (4.2) corresponds to a quadratic utility function in R (or, equivalently, in $1 + R$). Also note that if the distribution of the risky asset returns is multivariate elliptical, then so is of course the distribution of the terminal values of the risky assets, and vice versa. Qualitatively, there is therefore no inconsistency. Quantitavely there is, however: Although the utility functions of return and terminal wealth are both quadratic, their coefficients (corresponding to k in (4.2)) are different. Thus, if an agent has a given quadratic utility function of terminal wealth, the corresponding utility function of return depends on his initial wealth, and it would therefore in general be inconsistent with the expected utility maxim for the agent to choose the same relative portfolio at different wealth levels.

Without particular assumptions about the joint distribution of the risky asset returns, little can be said in general about the structure of the optimal portfolio (if it exists) for an investor with a general non-quadratic utility function. Cass and Stiglitz (1970) derive conditions on the utility function for *separability*, i.e., for the existence of two mutual funds such that the investor can obtain his maximum expected utility by investing only in the mutual funds regardless of his initial wealth (the more general case where more than two funds are required is also studied). As such, the results are of minor interest because the asset allocation in each of the mutual funds may depend on the utility function of the investor, implying that the funds are not (necessarily) actually mutual, i.e., an investor with a different utility function may require a different set of funds. However, by inspection of the results it can be concluded (as was done in Merton (1992), Ch. 2) that in a general market with a risk-free asset there exists a single genuine mutual fund of risky assets such that any number, say K, of investors can obtain their maximum expected utility by investing only in the risk-free asset and the mutual fund if and only if the k'th investor has HARA utility with a risk aversion function in the form

$$a(x) = (cx + d_k)^{-1},$$

$k = 1, \dots, K$. Note that c must be the same for all investors, implying that the utility functions are very similar. Furthermore, in a general market *without* a risk-free asset there exist two (genuine) mutual funds providing sufficient investment possibilities if and only if all investors have quadratic utility functions or all investors have CRRA utility functions with the same CRRA coefficient. Of course, these results also require that the investors agree on the joint return distribution.

Various conditions on the joint distribution of asset returns for the existence of mutual funds can be found in Merton (1992) (Ch. 2), where many of the abovementioned results in the single-period setup are collected, and in some cases generalized.

2.5 Multi-period portfolio selection

In a multi-period portfolio selection problem the time set is given by $\mathcal{T} = \{0 = t_0 < \ldots < t_N = T\}$. At each time $t_k < T$ the investor chooses a portfolio that he holds until t_{k+1}, where the new asset prices materialize and the return on the portfolio over the period concerned is thus revealed.

Tobin (1965) appears to be (one of) the first to study optimal multi-period portfolio selection. He extends the Markowitz-Tobin single-period static mean-variance approach to more than one period and thus argues that the investor's choice in the multi-period case is described in the same way as in the single-period case: The investor should only consider portfolio sequences (i.e. strategies) that are efficient in terms of the overall return R given by $(1 + R) = (1 + R(1)) \cdot \ldots \cdot (1 + R(N))$, where, of course, $R(1), \ldots, R(N)$ are the N single-period returns. In other words, the expected value $\mathrm{E}(1 + R)$ should be maximized for a given value of the total risk σ, given by $\sigma^2 = \mathrm{V}(1 + R)$. In the general case, where $R(1), \ldots, R(N)$ can be mutually dependent and have different marginal distributions, an efficient portfolio sequence cannot be derived, but it is claimed that if they are taken to be i.i.d., all efficient portfolio sequences should be "stationary", that is, the expected returns (and thereby the risks) of all single-period portfolios in an efficient portfolio sequence should be equal. However, as pointed out by Mossin (1968), this is not correct.

Mossin (1968) begins with a one-period analysis and focuses on the quantitative difference between the utility functions of return and terminal wealth as discussed in Remark 2.4.3. It is thus made clear that the single-period mean-variance approach must be used with caution. He then shows for general asset return distributions that the only utility functions implying the same preference orderings regardless of whether the return or the terminal wealth is used as argument are the CRRA ones. Thus, these are the only utility functions leading to relative portfolio choices that are independent of initial wealth.

He then turns to the multi-period case and considers the general problem of maximizing expected utility of terminal wealth. He argues that dynamic programming is appropriate as an approach to the problem. The claim made by Tobin (1965) about stationarity of mean-variance efficient portfolio sequences is rejected — Tobin's analysis is flawed because he does not allow the single-period relative portfolios to depend on the changing wealth levels over time, which is necessary in the case of quadratic utility (however, although it is of minor importance we mention that within the *restricted* class of portfolio sequences that are not allowed to depend on the development of the prices, the optimal ones (in terms of the mean-variance criterion) are indeed the stationary ones).

In a setup with two independent assets and independence between periods Mossin characterizes the class of utility functions for which optimality is obtained by optimizing in each period without looking ahead. Such behaviour is termed "myopic" (i.e., short-sighted). In general, *complete myopia*, where the agent behaves in each period as if it were the last one and thus simply maximizes utility of the end-of-period wealth, is optimal if and only if the utility function is in the CRRA class. It is further shown that if the utility function is in the larger HARA class and one of the assets is risk-free, then complete myopia is still optimal if the interest rate is 0, whereas so-called *partial myopia*, where the agent behaves in each period *as if* the entire end-of-period wealth were to be invested in the risk-free asset for the remaining periods, is optimal if the rate is non-zero. Finally, Mossin shows that in general, a stationary portfolio policy is optimal if and only if the utility function is in the CRRA class and the yield distributions in the different periods are i.i.d.

Samuelson (1969) is apparently the first to study the combined problem of optimal investment *and* consumption. He works with a discrete-time model with $t_k = k$, $k = 0, \ldots, N$ (in particular $T = t_N = N$). There is a risk-free asset with rate r and a risky asset with dynamics

$$S(t+1) = S(t)Z(t),\ t = 0, \ldots, N,$$

where $Z(0), \ldots, Z(N)$ are i.i.d. random variables with values in $[0, \infty)$ and an arbitrary cdf. P. The objective can be written as

$$\sup_{C,w} \mathrm{E}\left[\sum_{t=0}^{N-1} (1+\rho)^{-t} v(\Delta C(t)) + (1+\rho)^{-T} v(X(T))\right]$$

subject to the wealth dynamics

$$X(t+1) = (X(t) - \Delta C(t))[(1+r)(1-w(t)) + w(t)Z(t)],$$

with $X(0)$ given. The parameter ρ is a subjective time preference parameter, $\Delta C(t)$ is a consumption lump at time t (consumption thus takes place at the beginning of each period), v is the utility function, and the $w(t)$ are the fractions of wealth put into the risky asset. Samuelson derives a solution algorithm for the general problem.

The important part of the main theorem of Samuelson (1969) states that for CRRA utility functions the optimal portfolio decision is independent of time, wealth, and all consumption decisions, leading to a constant optimal fraction w^*, which is characterized as the solution to

$$0 = \mathrm{E}\left(v'((1-w)(1+r) + wZ)(Z - 1 - r)\right).$$

This is in perfect accordance with the results of Mossin (1968). The result can easily be extended to the case with several risky assets, but, of course, no explicit portfolio expressions can be derived in general.

Hakansson (1970) studies a problem very similar to that of Samuelson (1969) but works with an infinite time horizon and includes fixed noncapital (e.g., wage) income $y \geq 0$ at the end of each period. He obtains explicit results for CRRA and

CARA utility functions similar to those obtained by Samuelson (1969). With the inclusion of fixed noncapital income one should consume and invest in the risky assets exactly as if the future noncapital income were replaced by its present value and added to the wealth. Thus, the sum S of the current wealth X and the present value y of future noncapital income should be thought of as the state variable. The amount to invest in the risk-free asset is then given as the remaining wealth, which equals the amount to invest if S were the wealth (and there were no future noncapital income) minus y. With S as the state variable the results regarding both investment and consumption are as in Samuelson (1969).

Hakansson (1969) extends the setup in Hakansson (1970) by considering the agent's lifetime as being uncertain (although the publication years of the two papers would seem to indicate otherwise, the 1970 paper does precede the 1969 paper). Hakansson (1969) works with CRRA utility functions and a finite time horizon, within which death is assumed to occur (this does not appear to be a crucial assumption, however). He first considers the case where no utility is assigned to the wealth at the time of death. The optimal investment strategy is then exactly as the one found in Hakansson (1970), because the uncertainty regarding the lifetime basically just changes the preference structure with respect to consumption at different times and thus only affects the optimal consumption strategy. In particular, the future noncapital income should still be capitalized on the basis of the risk-free interest rate only. He then considers the case where nonzero utility, given by a utility function u, is assigned to the wealth at the time of death. An explicit solution is found only in the case where $u = v$ and there is no future noncapital income. The optimal investment strategy is then exactly as in the first case. Hakansson then studies the problem in the situation where the agent has the possibility of buying term insurance on his life. He first considers the case where the agent can enter into an insurance contract, with fixed future premium payments, at the beginning of the period. Solutions are obtained in special cases, and the optimal investment strategy is as in the case without insurance apart from the fact that the capitalized future premiums should be deducted from the wealth before the results from the case with no noncapital income and no insurance is applied. He finally considers the case where the agent can decide the amount of term insurance to buy at each decision point, but no solution is obtained in the general case.

Remark 2.5.1 Markowitz (1952) actually considers the multi-period consumption/investment problem in the general form (3.1) under the restriction that C is a pure jump process with jumps $\Delta C(0), \ldots, \Delta C(T-1)$, i.e., only consumption lumps at times $0 < \ldots < T-1$ are allowed), and argues, under the assumption of i.i.d. risky asset return distributions, that in order to get the correct solution to the problem the theoretically correct approach is dynamic programming. However, Markowitz obtains no explicit results because he does not believe that it is possible in general to determine a reasonable, and yet tractable, utility functional $\mathcal{V}$ and to solve the single-period optimization problems (at least with the available machinery at the time of his writing) and therefore makes no attempt to study the problem in detail. In particular, although he recognizes that the assumption of time additivity of utility (of consumption as well as terminal weatlh) may ease the mathematical tractability considerably, he does not pursue the problem under this assumption because he considers it to be "uneconomical".

2.6 Continuous-time portfolio selection

In a continuous-time portfolio selection problem the investor chooses a new portfolio at each time $t \in [0, T]$, and new asset prices materialize continuously over time. Unless otherwise specified, the consumption stream is restricted to being absolutely continous wrt. the Lebesgue measure, i.e., of the form $C(t) = \int_0^t c(s)\, ds,\ t \in [0, T]$, where c is the consumption rate. The continuous-time setup has become the standard setup in the modern literature on portfolio optimization. It is of course more general than the single- and multi-period setups, as the investor can still use a simple (discrete-time) strategy if he likes. However, it should also be noted that the setup relies (in its purest form, at least) on very strong assumptions, and it adds complexity to the underlying mathematics.

A large part of the continuous-time models studied in the literature are based on Brownian motions, and it is therefore convenient to introduce a *standard Brownian motion model* of a financial market. In this model the uncertainty is driven by a standard D-dimensional Brownian motion $W = (W_1(t), \ldots, W_D(t))'_{t \in [0,T]}$. There is a locally risk-free asset with price dynamics given by

$$dS_0(t)/S_0(t) = r(t)\, dt,\ 0 \leq t \leq T, \tag{6.1}$$

and n (locally) risky assets with price dynamics of the form

$$dS_i(t)/S_i(t) = \alpha_i(t)\, dt + \sum_{d=1}^{D} \sigma_{id}(t)\, dW_d(t),\ 0 \leq t \leq T, \tag{6.2}$$

$i = 1, \ldots, n$. The interest rate, drift, and volatility processes, r, α_i, $i = 1, \ldots, n$, and σ_{id}, $i = 1, \ldots, n$, $d = 1, \ldots, D$, respectively, are referred to as the *coefficients* of the model; these are in general adapted processes, which must satisfy certain integrability conditions (see Karatzas and Shreve (1998)). In this model the assumption that the market is free of arbitrage opportunities implies (Karatzas and Shreve (1998), Theorem 1.4.2) that there exists an adapted D-dimensional process $\lambda = (\lambda_1(t), \ldots, \lambda_D(t))'_{t \in [0,T']}$ such that (for Lebesgue-almost-every $t \in [0, T']$ and almost surely)

$$\alpha_i(t) = r(t) + \sum_{d=1}^{D} \lambda_d(t) \sigma_{id}(t),\ i = 1, \ldots, n, \tag{6.3}$$

in particular, we must have $D \geq n$. In general, λ is not uniquely determined from the price process dynamics (6.1)-(6.2), but for a given process λ satisfying (6.3), the components $\lambda_1, \ldots, \lambda_D$, are interpreted as the *market prices of risk* related to the D

risk sources $W_1, \ldots, W_D$ (the market prices of risk are then also included when we refer generically to the *coefficients* of the model, of course). However, if the $n \times D$ matrix $\sigma(t)$ of volatility coefficients has full rank D (for Lebesgue-almost-every $t \in [0, T']$ and almost surely), then λ is uniquely determined.

If the process λ is given, the relation (6.3) is sometimes inserted directly in (6.2) in a given model. We denote by $\Sigma(t)$ the $n \times n$ variance matrix $(\sigma(t))(\sigma(t))'$. We refer to Karatzas and Shreve (1998) for details and more properties of this market model (be ware, however, that they use the term *standard* model in a slightly more restrictive way than we have done here).

For notational convenience we shall sometimes use other subscripts than numerical indices on the price processes and Brownian motions, and we may also equip these with superscripts.

Remark 2.6.1 Some authors allow the components of the Brownian motion to be correlated. This does not represent an essential generalization of the standard Brownian motion model introduced above: If Φ denotes the (fixed) instantaneous correlation matrix (of W), then there exists a $D \times D$ matrix B such that $\Phi = BB' = BIB'$, where I is the D-dimensional unit matrix. Then $W \overset{d}{=} B\widetilde{W}$, where $\widetilde{W}$ is a standard D-dimensional Brownian motion, so an equivalent model is obtained upon replacing W by $\widetilde{W}$ and the volatility matrix $\sigma(t)$ by $\sigma(t)B$. However, it should be noted that the interpretation of each $\widetilde{W}_i$ is different from the interpretation of W_i.

Merton (1969) is a companion paper of Samuelson (1969) that deals with the continuous-time version of the problem. Together with Merton (1971) the paper initiated the study of optimal consumption and investment in continuous time, and these have been some of the most celebrated papers in the financial literature since then. The results in Merton (1969) that are of interest here are also included in Merton (1971), to which we therefore turn our attention.

Merton (1971) works within the standard Brownian motion model. In the basic setup there are n risky assets with strictly positive price processes governed by the dynamics

$$dS_i(t)/S_i(t) = \alpha_i\, dt + \sigma_i\, dW_i(t),\ i = 1, \ldots, n, \tag{6.4}$$

where $W_1, \ldots, W_n$ may be correlated with coefficients ρ_{ij}, $i, j = 1, \ldots, n$. The risky asset coefficients, i.e., the α_i's and σ_i's, are in general allowed to depend on $(t, S(t))$, but most of the explicit results are obtained when they are assumed to be constant. In this case the price processes are geometric Brownian motions. There may or may not be a risk-free asset; both cases are treated. It is assumed that the variance matrix $\Sigma = (\Sigma_{ij})_{i,j=1,\ldots,n} = (\sigma_i \rho_{ij} \sigma_j)_{i,j=1,\ldots,n}$ is non-singular, which implies that the model is arbitrage free (and complete in the case with a risk-free asset).

For a given consumption rate process c and a risky asset portfolio fraction process w the wealth develops according to the dynamics

$$dX(t) = X(t)(r + w'(t)(\alpha - r1_n))\, dt - c(t)\, dt + X(t)w'(t)D(\sigma)\, dW(t),\ t \in [0, T].$$

When there is no risk-free asset the constraint $w'(t)1_n = \sum_{i=1}^n w_i(t) = 1$, $\forall t \in [0,T]$, must be imposed, and r then drops out of the dynamics as it should. No constraints on w are imposed when there is a risk-free asset; as noted in the introduction the fraction of wealth invested in the risk-free asset at any time $t \in [0,T]$ is then given residually by $1 - w'(t)1_n$.

The objective is

$$\max_{C,w} \mathrm{E}\left(\int_0^T v_1(t, c(t)) + v_2(T, X(T))\right),$$

where, for each $t \in [0,T]$, $v_1(t,\cdot)$ is a strictly concave utility function, and $v_2(T,\cdot)$ is a concave utility function.

The first important result is a *separation* or *mutual fund* theorem: When the risky asset coefficients are constant, there exists a unique (up to a nonsingular transformation) pair of "mutual funds" constructed from linear combinations of the assets such that, independent of preferences (i.e., the form of the utility functions), wealth, and time, the investor is indifferent between choosing from a linear combination of these two funds or a linear combination of all the original assets. Once again we emphasize the importance of such results for pension funds (note in particular the independence of preferences).

When there is a risk-free asset, the funds can be chosen such that one consists of the risk-free asset only and the other consists of risky assets only. The "risky" portfolio is then determined by the vector

$$\tilde{w} = \Sigma^{-1}(\alpha - r1_n). \tag{6.5}$$

In the case with no risk-free asset two mutual funds can be obtained by inserting two arbitrary values $r_1 < r_2 < \bar{r} := (1_n'\Sigma^{-1}\alpha)/(1_n'\Sigma^{-1}1_n)$ of the "free" parameter r in (6.5) and normalizing by the sum of the obtained vector components to obtain genuine portfolios. Note that the price (per share) processes of the mutual funds are geometric Brownian motions themselves.

Apart from the fact that the proportions of risky assets in the risky funds in the continuous-time case are determined by the *instantaneous* means and standard deviations, this is in perfect accordance with the results obtained in the single-period mean-variance analysis, see (4.1) and also Remarks 2.4.1 and 2.4.2, which apply here as well. The agreement of the form of the optimal portfolios here with the form obtained in the single-period mean-variance analysis is, in loose terms, due to the fact that at any time $t \in [0,T)$, the expected change of utility (as measured by the value function) over an infinitesimally short interval $[t + dt)$ is determined uniquely by the mean and variance of the instantaneous return of the portfolio in that interval. It should be noted, though, that the general assertions of the mutual fund theorem are based on the necessary first order conditions of optimality obtained from the dynamic programming principle only, and they therefore need not hold in full generality.

Remark 2.6.2 In an equivalent *standard* model, where the Brownian motions are independent (see Remark 2.6.1), it can be shown, by use of simple linear algebra, that the *loadings* on $\widetilde{W}_1, \ldots, \widetilde{W}_n$ (note that $n = D$ here) of the mutual fund are proportional to the market price of risk vector λ (cf. (6.3)). In other words, the diffusion term of the optimally invested wealth is proportional to

$$\lambda' \, d\widetilde{W}(t)$$

for each $t \in [0, T]$. There is no similar characterization when the Brownian motions are correlated, since there is no meaningful and consistent definition of the market prices of risk of the individual Brownian motions in this case.

Assume now that there exists a risk free asset and that the risky asset coefficients are constant. Because of the mutual fund theorem it may then be assumed without loss of generality that there is just a single risky asset (with price process denoted by S_1). It is furthermore assumed that v_1 is given by $v_1(t, y) = e^{-\rho t} V(y)$, where V is a HARA utility function parameterized as

$$V(y) = \frac{1-\gamma}{\gamma} \left(\frac{\beta y}{1-\gamma} + \eta \right)^{\gamma}, \quad \frac{\beta y}{1-\gamma} + \eta > 0,$$

with $\beta > 0$, $\gamma \in \mathbb{R} \setminus \{0, 1\}$, and for simplicity that $v_2(T, X(T)) \equiv 0$. Note that for $\eta = 0$, V is logarithmic in the limiting case $\gamma \to 0$, and for $\eta = 1$ it is exponential in the limiting case $\gamma \to -\infty$.

Under certain feasibility conditions (see Remark 2.6.3 below) the HJB equation can be solved explicitly, and the value function is given by

$$J(X, t) = \frac{\delta}{\gamma} \beta^{\gamma} e^{-\rho t} \left[\frac{\delta \left(1 - e^{-(\rho - \gamma\nu)(T-t)/\delta}\right)}{\rho - \gamma\nu} \right]^{\delta} \left[\frac{X}{\delta} + \frac{\eta \left(1 - e^{-r(T-t)}\right)}{\beta r} \right]^{\gamma}, \quad (6.6)$$

where $\delta = 1 - \gamma$, $\nu = r + (\alpha - r)^2/(2\delta\sigma^2)$, and $(1 - e^{-r(T-t)})/r$ is interpreted as $T - t$ if $r = 0$. At any time $t \in [0, T]$, the optimal amount to hold in the risky asset is given by

$$w^*(t) X(t) = \frac{\alpha - r}{\delta\sigma^2} \left(X(t) + f(t) \right), \quad (6.7)$$

where

$$f(t) = \frac{\delta\eta \left(1 - e^{-r(T-t)}\right)}{\beta r}, \quad t \in [0, T].$$

The optimal amount is affine in the wealth $X(t)$, that is, it is the sum of a deterministic amount and an amount that is proportional to $X(t)$. The deterministic amount, given by $f(t)(\alpha - r)/(\delta\sigma^2)$, decreases over time (if $\eta(\alpha - r) > 0$) and tends to 0 as $t \to T$.

Remark 2.6.3 The feasibility restriction $\rho - \gamma\nu > 0$ should be imposed (see Sethi and Taksar (1988)), as is done in Karatzas et al. (1986). Furthermore, as pointed out in Sethi and Taksar (1988), the solution is only correct if $\gamma < 1$ and $\eta = 0$ (corresponding to CRRA utility). If $\gamma < 1$ and $\eta > 0$ then $V'(0) < \infty$, and both wealth and consumption may become negative under the optimal strategy for any $t \in (0, T)$ (but not at time T!), thus violating the feasibility conditions $X(t), c(t) \geq 0$, $\forall t \in [0, T]$, (note, however, that these are not explicitly stated in Merton (1971) but only in Merton (1969)). Thus, unless one allows for negative consumption, J is not the value function. It *is* the value function if the constraints $X(t) \geq 0$, $\forall t \in [0, T]$, and $c(t) \geq 0$, $\forall t \in [0, T]$, are replaced by $X(T) \geq 0$ and $c(t) \geq -\delta\eta/\beta$, $\forall t \in [0, T]$, respectively. If $\gamma < 1$ and $\eta < 0$, the solution is correct only if the initial wealth satifies

$$X(0) \geq -f(0),$$

with strict inequality if $\gamma \leq 0$. Otherwise the problem has no solution.

If $\gamma > 1$, then V is constant on $[-\delta\eta/\beta, \infty)$, implying that $-\delta\eta/\beta$ is a consumption *satiation* level, where maximum utility of consumption is obtained. The solution is in this case only correct if the wealth and the consumption rate in are both allowed to become negative and, if so, only for

$$X(t) < -f(t).$$

If $X(t) \geq -f(t)$, the optimal investment strategy is to invest in the risk-free asset only.

In the case where v_1 and v_2 *both* have the form $v_i(t, y) = e^{-\rho t} V_i(y)$ with V_i given by

$$V_i(y) = \frac{1-\gamma}{\gamma}\left(\frac{\beta_i y}{1-\gamma} + \eta_i\right)^{\gamma}, \quad \frac{\beta_i y}{1-\gamma} + \eta_i > 0,$$

where $\beta_i > 0$, $i = 1, 2$, and $\gamma \in \mathbb{R} \setminus \{0, 1\}$, the value function (which is not provided in Merton (1971) but merely claimed to be of the same form as above) is given by

$$\begin{aligned} J(X,t) \quad &= \quad \frac{\delta\beta_1^{\gamma}}{\gamma} e^{-\rho t}\left(\frac{\delta\left(1 - e^{-(\rho-\gamma\nu)(T-t)/\delta}\right)}{\rho - \gamma\nu} + \left(\frac{\beta_2}{\beta_1}\right)^{\gamma/\delta} e^{-(\rho-\gamma\nu)(T-t)/\delta}\right)^{\delta} \\ &\quad \times \left(\frac{X}{\delta} + \frac{\eta_1}{\beta_1 r}\left(1 - e^{-r(T-t)}\right) + \frac{\eta_2}{\beta_2} e^{-r(T-t)}\right)^{\gamma}. \end{aligned}$$

Furthermore, if we set $v_1(t, y) \equiv 0$, i.e., only utility of terminal wealth, the terms $\delta\left(1 - e^{-(\rho-\gamma\nu)(T-t)/\delta}\right)/(\rho - \gamma\nu)$ and $\eta_1/\beta_1 r\left(1 - e^{-r(T-t)}\right)$ drop out of the first and second (large) paranthesis, respectively, and the value function can be written as

$$J(X,t) = \frac{\delta}{\gamma} e^{-\rho T} e^{\gamma\nu(T-t)}\left(\frac{\beta_2 X}{\delta} + \eta_2 e^{-r(T-t)}\right)^{\gamma}.$$

The optimal amounts to hold in the risky asset at time $t \in [0, T]$ in the two cases become, respectively,

$$w^*(t)X(t) = \frac{\alpha - r}{\delta\sigma^2}\left(X(t) + f_1(t) + g_2(t)\right),$$

where f_1 is as f with η and β replaced by η_1 and β_1, respectively, and

$$g_2(t) = \frac{\delta\eta_2}{\beta_2} e^{-r(T-t)}, \ t \in [0, T],$$

and

$$w^*(t)X(t) = \frac{\alpha - r}{\delta\sigma^2} \left(X(t) + g_2(t) \right).$$

In both cases the optimal amount is still affine in the wealth $X(t)$. However, when the utility of terminal terminal wealth is nonzero, the deterministic amount tends to $(\alpha - r)\eta_2/(\sigma^2\beta_2)$ rather than 0.

Remark 2.6.4 The flaws pertaining to the case $\eta \neq 0$ (see Remark 2.6.3) obviously also apply to the cases with HARA utility of the terminal wealth, and to the remaining analysis of Merton (1971); we refer to Sethi and Taksar (1988) for proper reformulations and solutions in this case.

On the other hand, when the solution is correct, the sufficiency conditions of optimality are fulfilled, and the assertions of optimality are correct. In particular, this also goes for the mutual fund theorem.

We now briefly compare the optimal strategies in the special cases of CRRA, exponential, and quadratic utility, when (for simplicity) $v_1 \equiv 0$. They are given, respectively, by

$$w^*(t)X(t) = X(t)\frac{\alpha - r}{\delta\sigma^2}, \tag{6.8}$$

$$w^*(t)X(t) = \frac{\alpha - r}{\beta_2\sigma^2} e^{-r(T-t)},$$

and

$$w^*(t)X(t) = \frac{\alpha - r}{\sigma^2} \left(\frac{\eta_2}{\beta_2} e^{-r(T-t)} - X(t) \right),$$

(as long as $X(t) \leq \frac{\eta_2}{\beta_2} e^{-r(T-t)}$).

There is a considerable difference between the optimal strategies in the three cases. The CRRA utility functions yield an intuitively appealing strategy, as the optimal *proportion* of wealth to hold in the risky asset, $w^*(t)$, is constant and in particular independent of time and wealth. This in particular would make it fairly easy to implement for a portfolio manager representing many investors (with the same degree of risk aversion). Moreover, the wealth never becomes negative with this strategy. With exponential utility, the optimal *amount* to hold in the risky asset is deterministic and in particular independent of the wealth, which is rather implausible. With quadratic utility, the optimal *amount* in the risky asset is proportional to the difference between the (present value of the) satiation level and the wealth $X(t)$. In other words, the wealthier the investor is, the less should he invest in the risky assets, which is also counterintuitive. Moreover, the wealth may become negative in the cases of exponential and quadratic utility.

We emphasize two important points: First, the form of the utility function plays an extremely important role, and this is also the case in more general models than the one considered here. There are very few specific results on optimal investment that hold for all utility functions. Second, of the three forms considered above, the CRRA utility functions are by far the most reasonable ones judging by the intuitive appeal of the optimal strategies, and it is therefore no surprise that they are the most widely used ones in the modern literature on portfolio optimization. Note that with CRRA utility the proportion of wealth invested in the risky asset(s) should be constant; there is no reason (in this basic setup) that it should decrease when t approaches T, as is commonly stated advice by practitioners.

Different variations of the original problem are also studied, and we shall report the most interesting results.

If the agent has a positive, non-capital income stream (e.g., salary) with deterministic rate $y = (y(t))_{t\in[0,T]}$, the optimal amount to hold in the risky asset at any time $t \in [0,T]$ in the case with utility of both consumption and terminal wealth becomes

$$w^*(t)X(t) = \frac{\alpha - r}{\delta\sigma^2}\left(X(t) + b(t) + f_1(t) + g_2(t)\right), \tag{6.9}$$

where

$$b(t) = \int_t^T e^{-r(s-t)}y(s)\,ds,\ 0 \le t \le T.$$

In the case with zero utility of terminal wealth (resp. consumption) the term $g_2(t)$ (resp. $f_1(t)$) simply drops out. Thus, the optimal amount can be calculated as in the case with no income by simply replacing the wealth by the sum of the wealth and the present value of the future income, $b(t)$. In particular, in the CRRA utility case, where $f_1 \equiv g_2 \equiv 0$, the fraction of wealth to hold in the risky asset now decreases over time as the present value of future income decreases. It is also worth making the somewhat counterintuitive observation that for any fixed $t \in [0,T)$, the optimal fraction of wealth to invest in risky assets, $w^*(t)$, is decreasing in the wealth $X(t)$ (regardless of the choice of utility function). Thus, the wealthier one is in terms of current wealth, the smaller is the optimal fraction of wealth to invest in risky assets.

The individual may take the uncertainty regarding his lifetime into account. If the (random) time of death, denoted by τ, is taken to be independent of W and to have a probability distribution given by $\mathrm{P}(\tau > t) = \exp(-\int_0^t \mu(s)\,ds)$, $t \ge 0$, for some deterministic, positive function μ with $\int_0^\infty \mu(s)\,ds = \infty$, and if the objective has the form

$$\max_{C,w} \mathrm{E}\left(\int_0^{T\wedge\tau} e^{-\rho t}V_1(c(t)) + 1_{(\tau>T)}e^{-\rho T}V_2(X(T))\right),$$

where V_1 and V_2 are given as above, and $T \le \infty$, then the optimization problem can be solved explicitly (Merton only shows this in the special case when μ is constant and $T = \infty$, however). The objective basically just changes the preference structure with respect to time, as can be seen by recasting it in the form

$$\max_{C,w} \mathrm{E}\left(\int_0^T e^{-\rho t-\int_0^t \mu(s)\,ds} V_1(c(t)) + e^{-\rho T-\int_0^T \mu(s)\,ds} V_2(X(T))\right),$$

Thus, only the optimal consumption strategy is affected; the optimal investment strategy is left unchanged (regardless of the involved utility functions). The latter is also true when a fixed, non-capital income stream is included, that is, the uncertainty regarding the lifetime plays no role in the evaluation of the future income stream. These results correspond to those obtained by Hakansson (1969). The more general — and perhaps more interesting — case where a nonzero utility is assigned to the wealth at the time of death if death occurs before T does not appear to have a closed-form solution. However, the first order conditions and the solution to the discrete-time analogue of this problem, which exist and is given by Hakansson (1969), clearly suggest that it is optimal to invest as in the case with no uncertainty regarding the lifetime.

More general long term price behaviours, where the risky asset coefficients are not necessarily constant, are also considered. In the general case it is possible to obtain an explicit solution with logarithmic utility (only): One should then always invest as in the case with constant coefficients but with the current values of the coefficients inserted. Two special cases, one where the drift coefficient is mean reverting, and one where prices are expected to adjust to a certain long term "normal" price level, are then considered. Explicit results are obtained only in the case with just a single risky asset and CARA utility, however (note that since the mutual fund theorem requires constant coefficients it is a restriction to work with just one risky asset here), but they indicate that the optimal portfolio rules may change significantly when the price dynamics change.

Merton (1973) studies optimal consumption and portfolio strategies in a capital market model, which is an extension of the model employed in Merton (1971). Here, the *investment opportunity set* (i.e., the coefficients in the price process dynamics) are allowed to be stochastic processes with the restriction that the vector X of price and coefficient processes must be Markovian. Such a model is known as a *factor model*, and the components of X are called *factors*. The main focus is on derivation of equilibrium asset prices from the optimal asset-demand behaviour of investors, but general results on investment (and consumption) are obtained in the process. It should be noted that the entire analysis is based on the first order conditions of optimality only, however, so the results should be considered with caution.

The optimal investment strategy is such that the demand for the i'th asset for the k'th investor is given by the sum of the usual term obtained in Merton (1971) and an additional term reflecting that the investor wants to hedge against unfavourable changes in the opportunity set. An unfavourable change is one that leads to a fall in the future consumption for a given level of (future) wealth. It is shown that if X_k is a component such that a positive change in X_k is unfavourable, then a risk averse investor will demand more (less) of the i'th asset the more positively (negatively) correlated its return is with changes in X_k.

Special attention is paid to two special cases of the general model. In the first one, the opportunity set is taken to be constant. This is the model also employed in Merton (1971), where optimal portfolio strategies were derived. It can be shown under the condition that the market portfolio is efficient in equilibrium that the returns of the n risky assets will satisfy

$$\alpha_i - r = \beta_i(\alpha_M - r),\ i = 1, \ldots, n,$$

where $\beta_i = \sigma_{iM}/\sigma_M^2$, σ_{iM} is the covariance of the return on the i'th asset with the return on the market portfolio, and α_M is the expected return on the market portfolio. This is a continuous-time analogue of the Security Market Line of the classical CAPM.

The second model is a single-factor model where the factor is taken to be the interest rate, which is (of course) stochastic. It is furthermore assumed that one of the assets (the n'th one, by convention) has a return that is perfectly negatively correlated with changes in r (e.g., a default free zero coupon bond). Merton obtains a *three-fund theorem* analogue to the mutual fund theorem in Merton (1971), with the third fund consisting purely of the n'th asset. In equilibrium one then obtains the relation

$$\alpha_i - r = \frac{\sigma_i(\rho_{iM} - \rho_{in}\rho_{nM})}{\sigma_M(1-\rho_{nM}^2)}(\alpha_M - r) + \frac{\sigma_i(\rho_{in} - \rho_{iM}\rho_{nM})}{\sigma_n(1-\rho_{nM}^2)}(\alpha_M - r),\ i = 1, \ldots, n-1,$$

where the ρ's are the correlation coefficients in the underlying driving Brownian motion.

Merton then returns to the general model, where m state variables are required to describe the opportunity set. The utility functions are also allowed to be state dependent. The three-fund theorem is generalized to an $(m + 2)$-fund theorem, i.e., there exist $m + 2$ mutual funds such that all investors can invest optimally by investing in the funds only. These funds can be chosen in such a way that the $(m+1)$'th one is the classical Merton (1971) portfolio of risky assets given by (6.5) with the current coefficients inserted, the $(m+2)$'th consists of the locally risk-free asset only, and, for $i = 1, \ldots, m$, the i'th has maximum possible correlation with the i'th state variable.

Richard (1975) studies a consumption and investment problem for an individual in a setup similar to the basic one in Merton (1971). However, the agent explicitly takes the uncertainty regarding his time of death into account and is provided with the possibility of purchasing (infinitesimal) life insurance coverage with a variable premium rate p. Furthermore, the agent has a fixed future income stream with deterministic rate y (contingent upon survival). The problem is akin to the discrete-time problem studied by Hakansson (1969). Denoting by τ the time of death (which is assumed to occur in $[0, T]$), and by $Z(\tau)$ the legacy in case of death at time τ (i.e., the sum of the wealth and the insurance sum), the objective is given by

$$\max_{c,p,w} \mathrm{E}\left(\int_0^\tau v_1(s, c(s))\, ds + v_2(\tau, Z_\tau)\right),$$

where c, p, and w are the controllable processes, i.e., the consumption rate, the premium rate, and the vector of portfolio fractions, respectively. The utility functions $v_i(s,\cdot)$, $i=1,2$, are assumed to be strictly concave.

The life insurance coverage takes the form of a sum of the amount $p(t)/\mu(t)$ to be paid out upon death at time t. Thus, $\mu(t)$, which is taken to be deterministic, is the price (charged by some life insurance company) per unit of insurance coverage at time t. Obviously, p must be predictable. It is assumed that $\mu(t) \geq \lambda(t)$, $\forall t \in [0,T]$, where λ denotes the actual mortality intensity (also deterministic), i.e., the company is allowed a risk loading. There are no (a priori) restrictions on the size or sign of p (see Remark 2.6.5 below).

As for general results, a mutual fund theorem also obtains here, and the composition of risky assets is as in Merton (1971). Furthermore, it is shown that if the future income stream is relinquished at time t in favour of the *certainty equivalent value* $b(t)$ given by

$$b(t) = \int_t^T e^{-\int_t^s r+\mu(\tau)\,d\tau} y(s)\,ds, \tag{6.10}$$

which is added to the current wealth, $X(t)$, then the expected future utility and the optimal strategy remain unchanged. This is analogue to the result in the case with fixed income in Merton (1971). However, $b(t)$ is here calculated with the discount factor $e^{-\int_t^s r+\mu(\tau)\,d\tau}$ (whereas the corresponding discount factor in Merton's result is $e^{-\int_t^s r\,d\tau}$). The explanation is that adding $b(t)$ to $X(t)$ reduces the optimal sum insured by $b(t)$ and thus the optimal premium rate by $b(t)\mu(t)$, and since $b(t)$, placed at the risk free interest rate, finances a certain payment stream with rate $y(s)-b(s)\mu(s)$, $t \leq s \leq \tau$, the agent is indifferent (for comparison, Merton's result was obtained in a somewhat different setup with no possibility to buy life insurance and no utility of the legacy upon death).

Explicit solutions are obtained when, for $i=1,2$, v_i has the form $v_i(t,c) = h_i(t)u(c)$, $0 \leq t \leq T$, where $h_i(t) > 0$, $\forall t \in [0,T]$, and u is a CRRA utility function (independent of i). Not surprisingly, the optimal investment strategy corresponds to (6.9) with b of course given by (6.10) (and $\delta = 1-\gamma$, $f_1 \equiv g_2 \equiv 0$). The optimal premium rate p^* is given by the equation

$$Z^*(t) := X(t) + \frac{p^*(t)}{\mu(t)} = \left(\frac{h_2(t)\lambda(t)}{a^\delta(t)\mu(t)}\right)^{1/\delta} (X(t)+b(t)),\ 0 \leq t \leq T,$$

where $a:[0,T] \to \mathbb{R}$ is given by

$$a(t) = \int_t^T \left((\lambda(s)/\mu(s))^{\gamma/\delta} (\lambda(s)h_2(s))^{1/\delta} + h_1^{1/\delta}(s) \right) e^{-\frac{1}{\delta}\int_t^s \lambda(\tau)-\gamma(\nu+\mu(\tau))\,d\tau} ds,$$

$0 \leq t \leq T$, with $\nu = r + (\alpha-r)^2/(2\delta\sigma^2)$.

A noteworthy feature of the optimal premium rate is that, for any $t \in [0,T]$, the corresponding (optimal) legacy $Z^*(t)$ is linear in $X(t)+b(t)$. We also see that if $h_2(t)\lambda(t)/(a^\delta(t)\mu(t)) \geq 1$ then $p^*(t) \geq 0$. Otherwise the optimal premium rate is strictly negative when the wealth is sufficiently large; in this situation the individual actually *sells* insurance to the company, i.e., he receives strictly positive "premiums" from the company, which in turn is entitled to receive a strictly positive lump sum amount upon his death. However, the optimal legacy will always be nonnegative, that is, his wealth will always be large enough to cover any sum that must be paid to the company.

Remark 2.6.5 The premium rate is not only allowed to become strictly negative; as we have seen it will in fact do so in some circumstances under the optimal strategy. This may seem unrealistic since selling insurance in this way is not possible in practice. However, the setup is not at all unrealistic; all that is required is that the individual has his wealth managed as a reserve by a company under a highly flexible insurance scheme that at all times allows him to make deposits and withdrawals and to determine the size of the *sum insured*, i.e., the total sum to be paid out should he die within the following small time period, subject only to the constraint that it must be nonnegative. The difference between the sum insured and the reserve would then simply constitute the actuarial sum at risk, from which the actuarial risk premium could be calculated and subtracted from the reserve. Positive and negative premium rates in Richard's setup would then simply correspond to positive and negative risk premiums, and both cases indeed occur in standard insurance schemes.

Aase (1984) sets up an optimal portfolio (and, secondarily, consumption) problem in a generalized market, where the price processes have drift, diffusion and jump parts (with deterministic relative jump sizes), and where the coefficients are in general allowed to be adapted processes so that the model is not necessarily Markovian. However, the proposed general solution method is erroneous, and in a special case with Markovian dynamics, which is also considered, the stated HJB equation is incorrect and thus leads to incorrect solutions. Nevertheless, the general problem is (partly) solved in the cases with logarithmic and linear utility (in the latter case neither borrowing nor short-selling is allowed), where the solutions are in fact correct in spite of the erroneous approach due to certain nice properties of these functions. Only a semi-explicit solution is provided in the logarithmic utility case, though. In the linear utility case the optimal strategy is, of course, to invest the entire wealth at any time $t \in [0, T]$ in the asset with the largest expected instantaneous return.

Karatzas et al. (1986) study a general consumption and investment problem in a setup similar to the basic one studied in Merton (1971), but with an infinite time horizon and with the explicit constraint that consumption must be nonnegative. The objective is to maximize

$$\mathrm{E}\left(\int_0^\tau e^{-\rho t} v(c(t))\, dt + e^{-\rho\tau} P\right),$$

where $\rho > 0$, τ is the time of bankruptcy, i.e., with X denoting the wealth process, $\tau = \inf\{t \geq 0 : X(t) = 0\}$ with $\inf \emptyset = \infty$, and $P \in [-\infty, \infty]$ is an arbitrary value assigned to the state of bankruptcy. There are several possibilities for specification and interpretation of P, a natural one is to set P equal to the "present utility" of zero consumption from the time of bankruptcy to infinity with the obvious interpretation that no consumption can take place after bankruptcy.

The utility function $v : [0, \infty) \to [-\infty, \infty)$ is assumed to belong to $C^3((0, \infty))$ and to satisfy $v(0) = \lim_{c \searrow 0} v(c)$. It is furthermore assumed that v is strictly increasing and strictly concave with $\lim_{c \to \infty} v'(c) = 0$.

The model is interesting only if $P < \frac{1}{\rho} \lim_{c \to \infty} v'(c)$; otherwise there is no optimal strategy (one should consume so as to get to bankruptcy as quickly as possible). On the other hand, if $P < \frac{1}{\rho} v(0)$ then the optimal investment and consumption strategies are exactly as in the case $P = \frac{1}{\rho} v(0)$, and they never lead to bankruptcy. It is therefore assumed that $\frac{1}{\rho} v(0) \leq P < \frac{1}{\rho} \lim_{c \to \infty} v'(c)$.

Explicit solutions for the optimal strategies and the value function V are obtained under a certain technical feasibility condition, which ensures that $V(x) \in \mathbb{R}$, $\forall x \in (0, \infty)$. It is shown in particular that the value function is twice continuously differentiable. A mutual fund theorem similar to the one obtained in Merton (1971) holds true, and the optimal proportion to invest in the risky fund is

$$\pi = -\frac{(\alpha - r)V'(x)}{\sigma^2 x V''(x)},$$

which in general depends on the value function V. In the case where v is a CRRA utility function with CRRA coefficient $1 - \gamma$, the abovementioned technical feasibility condition is equivalent to the condition $\rho - \gamma\nu > 0$, see Remark 2.6.3.

Bajeux-Besnainou and Portait (1998) analyze dynamic mean-variance efficient portfolios, i.e., portfolios in markets with continuous trading whose return (or terminal value) has minimum variance for given expectations. They show that in a market that offers a zero-coupon bond with maturity T but otherwise can be quite general in the sense that the price processes are merely semimartingales, the efficient frontier is generated by the static combinations of the zero coupon bond and some arbitrarily chosen efficient portfolio, which, if the market is complete, can be chosen to be a short position in the portfolio yielding a payoff equal to the state price deflator. These results substantially generalize those obtained by Merton (1971) in the quadratic utility case. Explicit results are obtained in some complete, Brownian motion-driven models, where the payoff of the state price deflator can be obtained using constant relative portfolios.

Khanna and Kulldorff (1998) generalize the mutual fund theorem of Merton (1971) to the case where the utility function of the investors are merely assumed to satisfy the very mild assumption that they should be nondecreasing (in particular, they are allowed to be linear or even convex). There exists a mutual fund such that if an optimal investment strategy exists, then there also exists an optimal strategy that invests in the mutual fund and the risk-free asset only, and if no optimal strategy exists, then for any strategy there is a strategy involving only the mutual fund and the risk-free asset, which is as good or better.

2.7 The martingale method

A major breakthrough in the field of financial mathematics was achieved with the seminal papers Harrison and Kreps (1979) and Harrison and Pliska (1981, 1983), which were mentioned in Paragraph 2.1.D. They have had tremendous impact on both foundations and applications of the modern theory in this area and have led to a powerful approach to portfolio optimization, sometimes called the *martingale method*, which can be used in general, particularly non-Markovian, models. In this section we explore some of the literature on this method and thereby make a slight digression from the main line of this survey in the sense that our focus will move somewhat from concrete, explicit results on optimal investment strategies to the methods employed to derive them (although (semi-)explicit results will also appear).

Pliska (1986) exploits the results of Harrison and Pliska (1981, 1983) to formulate a new approach to the optimal investment problem within the semimartingale model of Harrison and Pliska (1981, 1983) and under the assumption that an equivalent martingale measure (in short, an EMM) exists (the idea behind this approach was in fact mentioned already in Harrison and Pliska (1981, 1983), though). It begins with the seemingly simple observation that when viewed as random variables, the terminal wealths that can be obtained at some time horizon T from a given initial wealth x_0 via admissible and x_0-financed portfolio strategies employed over $[0,T]$ correspond one-to-one to the attainable contingent claims with initial price x_0. The set of *discounted* attainable contingent claims with initial price x_0 is then identified as a certain subset A of $L^1(\Omega, \mathcal{F}, \mathrm{Q})$, where Q is an arbitrary EMM called the *reference measure*. The investment problem is then formulated in terms of two subproblems: The first of these is to determine an optimal (discounted) terminal wealth (or, equivalently, an optimal (discounted) contingent claim with initial price x_0), according to some preference structure. This amounts to choosing an optimal element of A, and to this end the theory of convex analysis (see Appendix B) is utilized. The second is to determine an x_0-financed portfolio process generating this terminal wealth; this is then an optimal portfolio process.

More specifically, Pliska works with the objective of maximizing

$$\mathrm{E}\left(u(X(\omega), \omega)\right)$$

over all $X \in A$, where $u : \mathbb{R} \times \Omega \to \mathbb{R}$ is a *state-dependent* utility function, which is concave and strictly increasing in $x \in \mathbb{R}$ for each fixed $\omega \in \Omega$. The first result, based on results from the theory of convex optimization, gives a sufficient condition for optimality, which in the general case is somewhat abstract and difficult to apply but nonetheless deserves to be mentioned.

Suppose that $X^* \in A$ and that Y^* is a bounded random variable such that

$$\mathrm{E}^Q\left(XY^*\right) = x_0 \mathrm{E}^Q\left(Y^*\right), \ \forall X \in A. \tag{7.1}$$

If

$$u(X^*(\omega), \omega)\Lambda(\omega) - X^*(\omega)Y^*(\omega) = \sup_{x \in \mathbb{R}} \left\{u(x, \omega)\Lambda(\omega) - xY^*(\omega)\right\}, \text{ a.s.}, \tag{7.2}$$

where $\Lambda = d\mathrm{P}/d\mathrm{Q}$, then X^* is optimal, i.e.,

$$\mathrm{E}\left(u(X^*(\omega), \omega)\right) \geq \mathrm{E}\left(u(X(\omega), \omega)\right), \ \forall X \in A. \tag{7.3}$$

Furthermore, we must have $Y^* \geq 0$, a.s.

Conversely, it is shown that if the functional

$$X \mapsto \mathrm{E}\left(u(X(\omega), \omega)\right), \ X \in L^1(\Omega, \mathcal{F}, \mathrm{Q}),$$

is finite and continuous on $L^1(\Omega, \mathcal{F}, \mathrm{Q})$, then the above condition is also necessary in the sense that if $X^* \in A$ satisfies (7.3), then there exists a bounded random variable Y^* satisfying (7.1) such that (7.2) holds.

As mentioned, this condition for optimality is not easy to apply because it is in general difficult to characterize the set A in a way that allows for a method to determine a candidate for X^*. Furthermore, these results do not settle the question of whether an optimal terminal wealth exists at all. However, based on a result from Harrison and Pliska (1981, 1983) Pliska shows that when the market is complete one has

$$A = \{X \in L^1(\Omega, \mathcal{F}, \mathrm{Q}) : \mathrm{E}^Q\left(X\right) = x_0\},$$

and if $X^* \in A$ and Y^* is a bounded random variable such that (7.1) and (7.2) are satisfied, then Y^* must be a strictly positive constant, i.e., $Y^* = y$, a.s., for some $y > 0$. It can then be seen that according to the optimality condition (7.2), any $y > 0$ leads to an optimal attainable (discounted) wealth candidate $X^*(\omega, y)$ in closed form, obtained as the maximizer of the function

$$x \mapsto u(x, \omega)\Lambda(\omega) - xy, \ x \in \mathbb{R}, \tag{7.4}$$

(almost surely), provided that this function can be maximized (a.s.). Now, (7.1) must be satisfied (or, equivalently, we must have $X^* \in A$ and therefore $\mathrm{E}^Q\left(X\right) = x_0$). Thus, if $y^* > 0$ is a number such that (7.1) is satisfied with $Y^* = y^*$, then $X^*(\omega, y^*)$ is an optimal (discounted) terminal wealth.

It is crucial that the function (7.4) can be maximized (a.s.), and, as is easily seen, this depends heavily on the utility function u, or more specifically, on the function $x \mapsto u_\omega(x) := u(x, \omega)$, $x \in \mathbb{R}$, defined for each $\omega \in \Omega$. Pliska shows in particular that if $u_\omega \in C^1(\mathbb{R})$ and satisfies the conditions

$$\begin{aligned} \lim_{x \to \infty} u'_\omega(x) &= 0, \\ \lim_{x \to -\infty} u'_\omega(x) &= \infty, \end{aligned}$$

almost surely, then the function (7.4) can be maximized (a.s.), and an optimal attainable (discounted) wealth candidate $X^*(\omega, y)$ can thus be obtained. Furthermore, if there exist finite scalars ϵ and β such that

$$0 < \epsilon \leq u'_\omega(x_0)\Lambda(\omega) < \beta, \text{ a.s.,}$$

then there exists $y^* > 0$ such that $X^*(\omega, y^*) \in A$. These conditions are therefore sufficient to ensure that a solution to the problem exists, but they are not necessary.

What remains is to solve the second subproblem, i.e., to determine the optimal investment strategy. General results in this direction are not provided, but apparently the optimal strategy typically depends heavily on u.

Pliska finally solves the first subproblem in the case of exponential utility, and the corresponding optimal investment strategy is obtained explicitly in the special case of a Black-Scholes market. This is a special case of the results of Merton (1971), to which we refer for concrete expressions.

The approach taken by Pliska (1986) was further developed — and extended so as to incorporate optimization of consumption — in independent studies by Karatzas et al. (1987) and Cox and Huang (1989, 1991) in the context of complete standard Brownian motion models. The objectives considered in these papers can be written in the form

$$\max_{(c,\eta)\in\mathcal{A}(x)} \mathrm{E}\left(k_1 \int_0^T H(t)u(t,c(t))\,dt + k_2 H(T)u(X(T))\right),$$

where $\mathcal{A}(x)$ denotes the set of admissible pairs of consumption and investment strategies c and η corresponding to the investor's initial wealth x; $k_1, k_2 \in \{0,1\}$ are constants, and H is a given, subjective *discount factor process* of the agent, which is continuous and strictly positive. As for k_1 and k_2, all three cases of interest ($k_1 = 1 - k_2 = 0$, $k_1 = 1 - k_2 = 1$, and $k_1 = k_2 = 1$) are considered by Karatzas et al. (1987); Cox and Huang (1989, 1991) only consider the last of these. The factor involving H only appears in Karatzas et al. (1987), where u in turn is not allowed to depend on t.

The papers differ to some extent in terms of aims and scope as well as technical assumptions in the setups. Cox and Huang (1989) work with Markovian price processes (satisfying certain regularity conditions) and focus on constructing explicit optimal consumption and portfolio strategies. Karatzas et al. (1987) work with price processes whose coefficient processes are merely adapted but also uniformly bounded, and deal with existence issues as well as construction of (semi-)explicit expressions. Cox and Huang (1991) work with slightly more general price processes and focus on existence issues. In all cases, existence of an EMM is either explicitly or implicitly assumed, and it is furthermore unique due to the assumption of market completeness. A common requirement in the setups is that the wealth process of the agent must be nonnegative throughout $[0, T]$.

Although results on optimal consumption are in general taken to be outside the aims and scope of this survey, certain aspects of the general method developed in Karatzas et al. (1987) and Cox and Huang (1989, 1991) pertaining to the consumption issue as such cannot be ignored. The approach to optimal consumption is based on an idea quite analogous to the idea behind the approach to optimal investment of Pliska (1986). The observation that any affordable consumption rate process can be viewed as an attainable (stochastic) payment stream leads to the idea that determining an optimal consumption strategy is a matter of determining an optimal attainable payment stream. If the set of attainable payment streams can be characterized in a suitably simple way, one then needs to find the optimal element of this set, according to some optimization objective. If this can be achieved, then what remains is to find an investment strategy such that the resulting wealth process exactly finances the chosen consumption process.

We emphasize two important points of this approach: First, the method does not depend on the preference structure for consumption (although the actual optimization problem involved does, of course), so in this sense it supersedes the main argument for ignoring consumption rules in this survey. Second, it shows that in problems involving preferences for consumption, optimal strategies for consumption and investment are in general not only related, they are in a certain sense two sides of the same story: An optimal investment strategy is, by definition, one that exactly finances an optimal consumption stream. As a curiosity we mention here the observations made by Samuelson (1969) and Merton (1969) that in the case of CRRA utility, the optimal consumption and investment decisions are *independent* of one another. This, however, is a consequence of the particular properties of the CRRA utility functions and only goes for the decisions made at any *given* point in time; it does not impose any contradiction.

Having commented on the general results on optimal consumption we once again turn our attention to the issue of optimal investment strategies under objectives involving the terminal wealth only. The corresponding optimization problems of Karatzas et al. (1987) and Cox and Huang (1989, 1991) can then be viewed as special cases of the problem studied by Pliska (1986), and the proposed approaches are similar. Both Karatzas et al. (1987) and Cox and Huang (1989, 1991) provide sufficient conditions on u (and on the market coefficients) for existence of an optimal terminal wealth $X^{\hat{\eta}}(T)$. If it exists, it has the form

$$X^{\hat{\eta}}(T) = I\left(\mathcal{Y}(x)\frac{Z(T)}{H(T)S_0(T)}\right), \tag{7.5}$$

where $I := (u')^{-1}$ is the (generalized) inverse of the derivative of u, $Z(T)$ is the Radon-Nikodym derivative of the (unique) EMM with respect to the real probability measure P, and $\mathcal{Y}(x)$ is a strictly positive constant ensuring that $X^{\hat{\eta}}(T)$ is attainable with the initial wealth x, i.e.,

$$\mathrm{E}\left(X^{\hat{\eta}}(T)Z(T)/S_0(T)\right) = x.$$

In particular, if $H(T)$ is constant (a fairly plausible assumption), the optimal terminal wealth can be characterized by the property that its marginal utility is proportional to the *state price deflator* $Z(T)/S_0(T)$. In other words, the marginal utility of the optimal wealth divided by the state price is the same in (almost) all states, which is very appealing intuitively: If this quantity, which can be viewed as a measure of the (local) additional utility per unit of initial cost in each state, were larger in some states than in others, then it would be possible to increase the expected utility with the same initial wealth by, loosely speaking, adding more to the payoffs in the former states and reducing the payoffs in the latter.

Still, an optimal portfolio cannot easily be obtained explicitly in an applicable form in general. The assumption of market completeness implies that the optimal terminal wealth is attainable, but in general the hedging portfolio can only be stated in a form that involves the (abstract) integrand in the martingale representation of the Q-martingale

$$\left(\mathrm{E}^{\mathrm{Q}}\left(X^{\hat{\eta}}(T)/S_0(T)\,\middle|\,\mathcal{F}_t\right)\right)_{t\in[0,T]}$$

with respect to the D-dimensional Q-Brownian motion $W(\cdot)+\int_0^{\cdot}\lambda(s)\,ds$, where λ is the market price of risk process. However, under further regularity conditions (slightly different in the two papers) on the utility functions and the coefficient processes, implying in particular that the price processes are Markovian, Cox and Huang (1989) and Karatzas et al. (1987) show that two certain functions, from which explicit expressions of the optimal portfolios as well as the value function of the problem can be obtained, may be characterized as the unique solutions of two certain linear SDE's. In particular, explicit results that are consistent with the results of Merton (1971, 1973), are derived.

The explicit nonnegativity constraint on the wealth process generally plays an important role for the problem. If u is defined on all of $\mathbb{R}$ (which need not be the case for the constrained problem, of course), one may consider the corresponding unconstrained problem (without the nonnegativity constraint). Now, if $u'(0) := \lim_{x\searrow 0} u'(x) = \infty$ then the constraint is not binding, that is, the solutions of the two problems coincide. Otherwise it is binding in general, and Cox and Huang (1989) show that the optimal *constrained* terminal wealth corresponding to the initial wealth x_0, $\hat{X}^{x_0,c}(T)$, has the form

$$\hat{X}^{x_0,c}(T)=\left(\hat{X}^{y_0,u}(T)\right)^{+}=\hat{X}^{y_0,u}(T)+\left(\hat{X}^{y_0,u}(T)\right)^{-},$$

where $\hat{X}^{y_0,u}(T)$ is the optimal *unconstrained* terminal wealth corresponding to an initial wealth of y_0, which is determined in such a way that $\hat{X}^{x_0,c}(T)$ has the initial value x_0. Thus, the optimal *constrained* wealth is equivalent to the positive part of $\hat{X}^{y_0,u}(T)$ or, alternatively, to the sum of $\hat{X}^{y_0,u}(T)$ and its negative part. As for the latter characterization, $\hat{X}^{y_0,u}(T)$ obviously has the intial price y_0, and $\left(\hat{X}^{y_0,u}(T)\right)^{-}$ therefore must have the price x_0-y_0, in particular we must have $y_0\le x_0$. In other words, the optimal *constrained* wealth can be obtained by investing a part, y_0, of the initial wealth in a portfolio generating the optimal *unconstrained* terminal wealth $\hat{X}^{y_0,u}(T)$ and buying a put option on $\hat{X}^{y_0,u}(T)$ with exercise price 0.

The martingale method was extended to the case of incomplete markets by He and Pearson (1991a,b) and independently by Karatzas et al. (1991). He and Pearson (1991a) consider a discrete-time model with a finite underlying probability space but otherwise very general price processes, while He and Pearson (1991b) and Karatzas et al. (1991) consider continuous-time standard Brownian motion models. For brevity we focus on the latter two papers; the methodology of the former is similar.

When the market is incomplete there are infinitely many EMM's and therefore infinitely many state price deflators. Consequently, it becomes increasingly difficult to identify an optimal portfolio, but He and Pearson (1991b) and Karatzas et al. (1991) provide sufficient conditions for optimality and obtain existence results. The basic idea is as follows: If the optimal terminal wealth that is attainable in some (fictitious) *complete* market containing the assets in the actual (incomplete) market is such that it is attainable also in the actual market, then it must be optimal there as well, and the maximum expected utility in the fictitious complete market must in particular equal the maximum expected utility in the actual market. The optimal terminal wealth in any (fictitious) complete market can be obtained using the results from Cox and Huang (1989, 1991) and Karatzas et al. (1987), and the class of fictitious complete markets, which can be characterized and parameterized by the market price of risk processes θ, thus yields a class of candidate optimal terminal wealths $\widehat{X}^{\theta}(T)$. Karatzas et al. (1991) show, under certain mild regularity conditions on the utility function, that the abovementioned idea holds true. Another equivalent sufficient condition for optimality of a candidate $\widehat{X}^{\theta}(T)$ obtained from a fictitious completion is that

$$\mathrm{E}\left(Z(T)\widehat{X}^{\theta}(T)\right) \leq x_0$$

for any state price deflator $Z(T)$ (see Remark 2.7.1 below). This is intuitively reasonable because of the well-known fact that the attainable claims in an incomplete market are those whose price is the same under all consistent pricing measures.

He and Pearson (1991b) focus on another optimality condition. In order to be optimal, a candidate terminal wealth $\widehat{X}^{\theta}(T)$ obtained from a fictitious completion must satisfy

$$\mathrm{E}\left(u(\widehat{X}^{\theta}(T))\right) \leq \mathrm{E}\left(u(\widehat{X}^{\nu}(T))\right) \tag{7.6}$$

for all possible market price of risk processes ν. In other words, θ must represent the "worst possible completion" of the market in terms of (maximum) expected utility for the agent. The intuitive explanation since is that since the expected utility of the optimal terminal wealth must be smaller than the expected utility of the optimal terminal wealth in any fictitious completion, it must in particular be smaller than in the worst one. Although this condition is necessary (as is easily seen and also shown by Karatzas et al. (1991)), it is not sufficient in general. However, it naturally leads to the so-called dual problem of minimizing the maximum expected utility over all fictitious completions. He and Pearson (1991b) show that a solution to the original problem can be obtained from a solution to the dual problem.

Both He and Pearson (1991b) and Karatzas et al. (1991) obtain results on existence of a solution to the original problem using results from convex duality theory. We shall not go into details but simply state the most important results. Assume that $u \in C^2(0, \infty)$. He and Pearson (1991b) show that a solution exists if u is bounded from above, and

$$\frac{-xu''(x)}{u'(x)} \geq 1, \ \forall x > 0.$$

On the other hand, Karatzas et al. (1991) show, under certain regularity conditions, that a solution exists if $u(0) > -\infty$ and

$$\frac{-xu''(x)}{u'(x)} \leq 1, \ \forall x > 0.$$

In particular, a solution exists for any CRRA utility function (assuming sufficient regularity).

Remark 2.7.1 The term "state price deflator" is not necessarily fully appropriate here, because in general we only have

$$\mathrm{E}\left(S_0(T)Z(T)\right) \leq 1,$$

and the inequality may be strict. This does not imply existence of arbitrage opportunities for the investor, however, due to the explicit restriction that the agent's wealth process must remain nonnegative.

As for explicit results, Karatzas et al. (1991) show that in the case of logarithmic utility, the optimal relative portfolio strategy is always given by the vector (6.5), with the now random coefficients inserted. This is a substantial generalization of the corresponding result of Merton (1971). Thus, a logarithmic investor simply adjusts his relative portfolio according to the current coefficient values and in particular does not care about their (anticipated) future development. In the case of power utility they show that the relative portfolio obtained by Merton (1971) (with the random coefficients inserted, of course) is optimal in the special case where all the market coefficients are *totally unhedgeable*, i.e., adapted to a filtration generated by a (multi-dimensional) Brownian motion that is independent of the (multi-dimensional) Brownian motion involved in the diffusion term of the risky asset price dynamics. Thus, a power utility investor simply ignores the undiversifiable risks. The Merton relative portfolio is also optimal if the market is complete and the interest rate and the market price of risk (but not necessarily the volatility matrix) are deterministic.

He and Pearson (1991b) show how a solution may be obtained in the special case of a Markov factor model by solving a quasi-linear PDE. In particular they obtain results that are consistent with those obtained by Merton (1973) and generalize those obtained by Cox and Huang (1989).

Kramkov and Schachermayer (1999) extend the martingale approach further to the general case of semimartingale financial markets, treating the incomplete-market case as well as the complete-market case. They work with the discounted price processes and consider the problem of maximizing the expected utility of the discounted terminal wealth. It is assumed that the utility function $u : (0, \infty) \to \mathbb{R}$ is in $C^1(0, \infty)$ and satisfies the conditions

$$\begin{aligned} u'(0) &:= \lim_{x \searrow 0} u'(x) = \infty, \\ u'(\infty) &:= \lim_{x \to \infty} u'(x) = 0. \end{aligned}$$

In particular, this means that the derivative, u', has a well-defined inverse $I := (u')^{-1} : (0, \infty) \to (0, \infty)$. The main result is that if the *asymptotic elasticity* of the utility function u is strictly less than 1, i.e., if

$$\limsup_{x\to\infty} \frac{xu'(x)}{u(x)} < 1, \tag{7.7}$$

then an optimal (discounted) terminal wealth $\hat{X}(T)$ exists and is unique. The given characterization of $\hat{X}(T)$ is somewhat abstract, and we provide here a characterization adapted to the present context. The optimal (discounted) terminal wealth $\hat{X}(T)$ is given by

$$\hat{X}(T) = I\left(\mathcal{Y}(x)\hat{Y}(T)\right), \tag{7.8}$$

where $\mathcal{Y}(x) > 0$ is a certain constant, and $\hat{Y} = (\hat{Y}(t))_{t\in[0,T]}$ is the solution (which is also shown to exist and be unique) to the dual problem of minimizing the expected value

$$\mathrm{E}\left(u\left(I\left[\mathcal{Y}(x)Y(T)\right]\right) - \mathcal{Y}(x)Y(T)I\left[\mathcal{Y}(x)Y(T)\right]\right)$$

over the set of nonnegative supermartingales $Y = (Y(t)_t)_{t\in[0,T]}$ with initial value $Y(0) = 1$ and with the property that for any admissible portfolio process with corresponding wealth process X, the process YX is a supermartingale. The constant $\mathcal{Y}(x)$ is determined by

$$x = \mathrm{E}\left(\hat{X}(T)\hat{Y}(T)\right) = \mathrm{E}\left(I\left(\mathcal{Y}(x)\hat{Y}(T)\right)\hat{Y}(T)\right),$$

and $\hat{X}\hat{Y}$ is a uniformly integrable martingale.

This characterization is of course somewhat abstract. However, the expression (7.8) is similar to (7.5) apart from the discount factor $H(T)$ in (7.5), which only pertains to the setup in Karatzas et al. (1991), and the factor $S(T)$, which has dropped out because Kramkov and Schachermayer (1999), as mentioned, work with the expected utility of the *discounted* terminal wealth. The random variable $\hat{Y}(T)$ corresponds to $Z(T)$: It is easy to note that the abovementioned set of nonnegative supermartingales contains the density processes of the equivalent local martingale measures of the market. In the general semimartingale model it is necessary to consider this larger set of supermartingales in order to get an existence result, but apart from that we have a similar characterization as above ((7.6) and the remarks following it): The supermartingale $\hat{Y}$ represents, in a certain sense, the worst possible completion of the market.

It is worth noting that this existence result relies on the assumptions about the utility function only, and that the characterization of the optimal terminal wealth depends on the utility function and the optimal supermartingale $\hat{Y}$ of the dual problem. It does not depend on the characteristics of the semimartingale modelling the market as such. Again, this emphasizes the extremely important role of the utility function. Kramkov and Schachermayer (1999) also show that the requirement (7.7) is essentially also necessary for existence of an optimal terminal wealth in general. In the survey article Schachermayer (2002), utility functions satisfying (7.7) are therefore said to have *reasonable asymptotic elasticity*.

Cvitanić et al. (2001) extend the approach to the general case of an incomplete semimartingale market where the investor has an unhedgeable, but uniformly bounded endowment process (not necessarily nondecreasing). They obtain a result on existence of an optimal strategy under some fairly mild assumptions including the assumption of reasonable asymptotic elasticity of the investor's utility function, but we shall not go into details.

Gerber and Shiu (2000) consider an optimal investment problem in a complete market with a fixed interest rate. The optimization problem is solved for HARA utilities using an approach that is based on the so-called Esscher transform but essentially is similar to the martingale technique. In the case of CRRA utilities with a positive CRRA coefficient, they allow for a minimum guarantee. Their setup and results are essentially a special case of earlier studies, but the paper provides a nice, self-contained exposition of the problem, the solution method, and the results.

2.8 Problems with stochastic interest rates

Although the general models considered in the various papers on the martingale methodology (Pliska (1986), Karatzas et al. (1987), Cox and Huang (1989, 1991) etc.) allow in particular for stochastic interest rates, these papers contain only a few explicit results that can be readily applied in concrete portfolio optimization problems. However, investment problems with specific stochastic interest rate models, which can lead to explicit strategies, have received a lot of attention in recent years.

Sørensen (1999), Brennan and Xia (2000), Bouiler et al. (2001), Korn and Kraft (2001), Jensen and Sørensen (2001), and Bajeux-Besnainou et al. (2003) study pure investment problems in complete, *standard* models (as introduced in the beginning of Section 2.6) with (nondegenerate) Gaussian term structures of interest, i.e., with all (instantaneous) forward rates, as well as the short rate, normally distributed.

In particular, Sørensen (1999), Bouiler et al. (2001), Korn and Kraft (2001), and Bajeux-Besnainou et al. (2003) consider single-factor models, where the short rate of interest has dynamics of the form

$$dr(t) = (\theta(t) - \kappa r(t))\, dt - \sigma_r\, dW_r(t), \tag{8.1}$$

where $\theta : [0, T] \to \mathbb{R}$ is a deterministic function, and $\kappa, \sigma_r \in \mathbb{R}$ are constants (to prevent these models from being extremely unrealistic, the so-called *mean reversion* parameter κ should be nonnegative, although this is only assumed explicitly in Korn and Kraft (2001) and Bajeux-Besnainou et al. (2003)). It is assumed that there is a locally risk-free asset, with price dynamics given by

$$dS_0(t) = S_0(t) r(t)\, dt,$$

of course. Sørensen (1999), Bouiler et al. (2001), and Bajeux-Besnainou et al. (2003) adopt the Vasicek (1977) term structure model, where $\theta(\cdot) \equiv \theta$ is constant (for $\kappa > 0$, θ/κ is then the long-run mean of the interest rate, so θ should also be nonnegative, but this is not assumed in these papers). Korn and Kraft (2001) consider the extended Vasicek (1977) model due to Hull and White (1990) as well as the Ho and Lee (1986) model where $\kappa = 0$; the function θ is in both cases determined so as to make the model fit the initial term structure. Thus, in each case there is a T'-bond (i.e., a zero-coupon bond with maturity T') for some $T' \in [T, \infty)$, which has price dynamics of the form

$$dS_1^{T'}(t)/S_1^{T'}(t) = \big(r(t) + \lambda_r \sigma_r B(t, T')\big)\, dt + \sigma_r B(t, T')\, dW_r(t), \tag{8.2}$$

where $B(t, T') = (1 - e^{-\kappa(T'-t)})/\kappa$ when $\kappa > 0$ and $B(t, T') = T' - t$ when $\kappa = 0$. Here we have equipped S_1 with the topscript T' to emphasize the maturity of the bond. The market price of (interest rate) risk parameter λ_r is assumed to be constant (Korn and Kraft (2001) allow λ_r to vary deterministically over time, but essentially this is not a generalization).

Furthermore, there is a stock with price dynamics given by

$$dS_2(t)/S_2(t) = (r(t) + \lambda_r \sigma_{2r} + \lambda_s \sigma_{2s})\, dt + \sigma_{2r}\, dW_r(t) + \sigma_{2s} dW_s(t),$$

where $\sigma_s, \lambda_s \in \mathbb{R}$ are constants (or deterministic functions), and W_s is independent of W_r. Thus, the stock may be correlated with the interest rate (and thus the bond), but there is also independent stock risk with market price λ_s.

In terms of the optimal portfolio proportions, the solution for a CRRA investor with CRRA coefficient $\delta > 0$ is given by

$$\begin{aligned}
\begin{bmatrix} w_1(t) \\ w_2(t) \end{bmatrix} &= \frac{1}{\delta} \begin{bmatrix} \sigma_r^2 B^2(t, T') & \sigma_{2,r}\sigma_r B(t, T') \\ \sigma_{2,r}\sigma_r B(t, T') & \sigma_{2,r}^2 + \sigma_{2,s}^2 \end{bmatrix}^{-1} \begin{bmatrix} \lambda_r \sigma_r B(t, T') \\ \lambda_r \sigma_{2,r} + \lambda_s \sigma_{2,s} \end{bmatrix} \\
&\quad - \frac{1}{\delta} \begin{bmatrix} (1-\delta) B(t, T)/B(t, T') \\ 0 \end{bmatrix} \\
&= \frac{1}{\delta} \begin{bmatrix} (\lambda_r \sigma_{2,s} - \lambda_s \sigma_{2,r})/(\sigma_{2,s}\sigma_r B(t, T')) - (1-\delta) B(t, T)/B(t, T') \\ \lambda_s/\sigma_{2,s} \end{bmatrix},
\end{aligned}$$

and $w_0(t) = 1 - w_1(t) - w_2(t)$, where $B(T,T)/B(T,T)$ is interpreted as 1 if $T' = T$ (note, though, that Korn and Kraft (2001) assume $\delta < 1$). The proportion invested in the stock, $\lambda_s/\delta\sigma_{2,s}$, is the same as in the case with constant interest rate, cf. (6.7). From the expression in the first line it is seen that the optimal relative portfolio is the sum of the (generalized) classical Merton (1971) portfolio (cf. (6.5) and (6.8)), and a hedge term used to hedge the interest rate risk; this term involves the T'-bond only. In particular, a high risk averter (with a large δ) will always have a large proportion in the T'-bond.

If $T' = T$ (as in Sørensen (1999), Bouiler et al. (2001), and Bajeux-Besnainou et al. (2003)), then the optimal strategy has the somewhat undesirable feature that the bond position explodes at the time horizon, as $B(T) = 0$ (this does not impose admissibility problems, though). The reason is that the volatility term of the bond thereby also tends to 0, so that in order to make the interest rate risk factor stay involved the portfolio, which is necessary for effective diversification and to benefit from the interest rate risk premium, the position in the bond must explode. This is of course mostly a theoretical problem. As a sort of remedy, Bouiler et al. (2001) and Bajeux-Besnainou et al. (2003)) show that the optimal portfolio can be written in terms of constant fractions by introducing another interest rate derivative, interpreted as a "rolling horizon bond" with constant time to maturity, with dynamics given by

$$dS_3(t)/S_3(t) = (r(t) + \lambda_r\sigma_{3,r})\,dt + \sigma_{3,r}\,dW_r(t),$$

for some constant $\sigma_{3,r}$ (this derivative is redundant in the sense that it can be mimicked perfectly in the model by a dynamically managed portfolio in the locally risk-free asset and the T'-bond). Then $w_1(t)$ becomes $-(1-\delta)/\delta$, $w_2(t)$ is as before, $w_3(t)$ is given by $(\lambda_r\sigma_{2,s} - \lambda_s\sigma_{2,r})/(\delta\sigma_{2,s}\sigma_{3,r})$, and $w_0(t) = 1 - \sum_{i=1}^{3} w_i(t)$. In particular, a mutual fund consisting only of shares of the bank account, the stock, and the "rolling horizon bond" with stock proportion $\lambda_s/\sigma_{2,s}$, "rolling horizon bond" proportion $(\lambda_r\sigma_{2,s} - \lambda_s\sigma_{2,r})/(\sigma_{2,s}\sigma_{3,r})$, and the rest in the bank, could be shared by all CRRA investors; an investor with coefficient δ and time horizon T would then invest a proportion $1/\delta$ in this mutual fund and $1 - 1/\delta$ in a T-bond.

Korn and Kraft (2001), on the other hand, assume that $T' > T$, which prevents the position in the zero-coupon bond from exploding.

Brennan and Xia (2000) consider a more general two-factor model, in which the interest rate dynamics are given by

$$dr(t) = (\theta(t) + Y(t) - \kappa r(t))\;dt - \sigma_r\,dW_r(t),$$

where Y is a random factor with dynamics

$$dY(t) = (\lambda_Y\sigma_Y - \rho Y(t))\;dt - \sigma_Y\,dW_Y(t).$$

Here, $\lambda_Y, \rho, \sigma_Y$ are constants (ρ should again be strictly positive), and W_r and W_Y are allowed to be correlated. There are two zero-coupon bonds with different maturities $T_1 < T_2$; their price dynamics are given by

$$\begin{aligned} dS_i^{T_i}(t)/S_i^{T_i}(t) &= (r(t) + \lambda_r\sigma_r B(t, T_i) + \lambda_Y \sigma_Y C(t, T_i))\, dt \\ &\quad + \sigma_r B(t, > T_i)\, dW_r(t) + \sigma_Y C(t, T_i)\, dW_Y(t), \end{aligned}$$

$i = 1, 2$, where $C(t, T_i)(t) = e^{-\kappa(T_i - t)}/(\kappa(\kappa - \rho)) - e^{-\rho(T_i - t)}/(\rho(\kappa - \rho)) + 1/(\kappa\rho)$. It seems necessary to assume that $T_1 \geq T$, although this is not done explicitly in the paper. Moreover, as in the abovementionded papers there is a stock with price dynamics

$$dS_3(t)/S_3(t) = (r(t) + \lambda_s\sigma_s)\, dt + \sigma_s\, dW_s(t),$$

where W_s is in general correlated with W_r and W_Y. Again, the optimal portfolio is the sum of the usual (generalized) Merton (1971) portfolio and two portfolios constructed so as to hedge the two state variables. However, it is shown that this model effectively reduces to a single-factor model in the sense that a T-bond is sufficient to hedge the risk inherent in the two state variables. In other words, if $T_1 = T$, then the hedging portfolio is just the T-bond, and the optimal portfolio proportions in the (locally) risky assets are given by

$$\begin{bmatrix} w_1 \\ w_2 \\ w_3 \end{bmatrix} = \frac{1}{\delta}\Sigma^{-1} \begin{bmatrix} \lambda_r\sigma_r B(t, T_1) + \lambda_Y\sigma_Y C(t, T_1) \\ \lambda_r\sigma_r B(t, T_2) + \lambda_Y\sigma_Y C(t, T_2) \\ \lambda_s\sigma_s \end{bmatrix} + \left(1 - \frac{1}{\delta}\right) \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix},$$

where Σ is the variance matrix of the instantaneous returns. If $T < T_1$, then two bonds and the bank account can be used to generate a T-bond.

Jensen and Sørensen (2001) extend the problem studied in Sørensen (1999) to the general case of Gaussian term structure models without specific assumptions about the short rate. Furthermore, they allow for n risky assets with volatility coefficients that are stochastic processes. The market price of risk vector is deterministic, however. They obtain the general result that for a CRRA investor the optimal asset allocation is obtained by investing a part of the wealth in the risky assets according to the proportions given by $\tilde{w}/\delta$, where $\tilde{w}$ is calculated as in (6.5), and dividing the remaining proportion of wealth, $1 - 1_n'\tilde{w}/\delta$, between the locally risk-free asset and a zero-coupon bond maturing at the time horizon T in a way that depends on the interest rate model.

We shall briefly comment on other aspects of these papers. Sørensen (1999) focuses on implementation issues that would occur in practice, where continuous trading in particular is impossible. In particular it is shown that one can approximate the solution by solving a series of short-term mean-variance optimization problems in terms of the short-term drift and variance of the *forward price* of the optimal wealth, i.e., $X(t)/S_1^T(t)$, where X is the wealth process of the optimal investment strategy.

In Bouiler et al. (2001) the optimization problem described above is actually only a part of a bigger, initial problem, namely that of a (defined contribution) pension fund that seeks to optimally manage a policy, which specifies a deterministic continuous premium rate and a deterministic *minimum* annuity rate beginning at the time of retirement of the policy holder (who is implicitly assumed to live until the annuity runs out, i.e., mortality risk is ignored). This initial problem is formulated with the objective of maximizing the expected utility of the surplus (total assets minus liabilities) pertaining to the policy at the time of retirement, subject to the constraint that this surplus must be nonnegative. It is shown that this problem can be solved by short-selling and buying bonds at time 0 in such a way that the stream of premiums and minimum annuity benefits are perfectly matched, and then using the proceeds (which must be nonnegative in order for the problem to have a solution) as the initial value of a portfolio, which can be managed independently so as to maximize the expected utility. This latter problem is the one described above.

Korn and Kraft (2001) start out with a more general interest rate model given by the dynamics

$$dr(t) = a(t, r(t))\, dt - \sigma_r\, dW_r(t),$$

where $a : [0, T] \times \mathbb{R} \to \mathbb{R}$ is an arbitrary function (satisfying certain regularity conditions). Obviously, (eq:GaussianSingleFactorShortRateDynamics) appears as a special case. They use the dynamic programming approach and prove a non-standard verification theorem applicable in this general their model. In this regard it is worth pointing out that most standard verification theorems (e.g. Fleming and Rishel (1975), Fleming and Soner (1993), Yong and Zhou (1999)) rely on a certain Lipschitz assumption, which is not satisfied in optimal investment problems with an unbounded interest rate.

Jensen and Sørensen (2001) also consider the situation where the investor is constrained by the requirement that the terminal wealth must exceed a certain minimum amount, say G. This is also known as *portfolio insurance*. As in Cox and Huang (1989), the solution in this case is to take a fixed, long position in zero-coupon bonds that yield the payoff G and use the remaining initial wealth to buy a European call option, with strike price G, written on an asset that is formed by investing an initial amount according to the optimal relative portfolio in the unconstrained case; the initial amount of this portfolio is determined in such a way that the price of this option exactly equals the remaining initial wealth. Of course, if this option is not readily available, one can (in principle, at least) mimick it perfectly via a self-financing investment strategy. The authors focus on measuring the loss of expected utility for the investor implied by the constraint and find that it may be quite severe if G is relatively large and the investor is relatively risk tolerant. They also measure the implied loss of expected utility imposed on the investor if the wealth is invested according to a different CRRA coefficient, which may be the case if many investors with different risk attitudes enter into a common investment pool, e.g. a pension fund. Their results indicate that such losses may also be quite severe, in particular for low risk averters. However, in the case of a minimum guarantee, a high risk averter is somewhat compensated by the guarantee.

Brennan and Xia (2000) and Bajeux-Besnainou et al. (2003) discuss the optimal strategies and compare them with "popular advice" from certain financial advisors. In particular, both papers address an apparent puzzle pointed out by Canner et al. (1997), namely that according to many financial advisors the ratio of the bond and stock proportions in the portfolio of (locally) risky assets should increase with risk aversion, which seems to contradict the mutual fund theorems. Brennan and Xia (2000) explain that this advice *is* consistent with their results when the zero-coupon bond(s) in the optimal portfolio are taken into account as part of the portfolio of (locally) risky assets. Bajeux-Besnainou et al. (2003) also demonstrate consistency, but from a different perspective. Note first that if one considers the T-bond as the risk-free asset (for an investor with horizon T) and considers the mutual fund with constant proportions in the bank account, the stock, and the "rolling horizon bond" as the risky-asset fund, then there is indeed an inconsistency when CRRA investors are considered. However, the authors show that the optimal strategy for an investor with HARA utility with $c \neq 0$ is to buy and hold a T-bond with payoff equal to $-d/c$ and invest the remaining wealth in the optimal portfolio of a CRRA investor with CRRA coefficient $1/c$ (note that a short position in this portfolio may be appropriate, as the remaining wealth after the initial purchase of the T-bond may be negative). This means in particular that the ratio of bond and stock proportions in the risky asset-fund is no longer constant but does indeed increase with the risk aversion.

Canestrelli and Pontini (2000) also study a pure investment problem where the short interest rate follows the Vasicek (1977) dynamics, but in contrast to the abovementioned papers they assume that there are n risky assets modelled as in the basic setup in Merton (1971), i.e., with price dynamics given by (6.4) with constant coefficients. In particular, the market is incomplete. Although the Brownian motion $W = (W_1, \ldots, W_n)'$ (where the components are in general correlated) is allowed to be correlated with W_r, interest rate derivatives are unattainable in general, and there is in particularly no zero-coupon bonds available.

An explicit solution is obtained under certain regularity conditions on the market coefficients and the CRRA coefficient δ. In consistency with the results of Merton (1973), the optimal investment strategy can be expressed in terms of three mutual funds: The first of these is the usual well-diversified Merton (1971) portfolio given by (6.5) and generalized, as the current interest rate is inserted at any time; the second is determined by the vector $\Sigma^{-1}\Gamma'$, where Γ is the (row) vector of covariances between the (instantaneuos) movements in the interest rate and the n risky asset prices, i.e.,

$$\Gamma = (g\sigma_1\eta_{r1}, \ldots, g\sigma_n\eta_{rn}),$$

with $\eta_{ri}\, dt = dW_r(t)\, dW_i(t)$, $i = 1, \ldots, n$; the third consists of the locally risk-free asset only. The proportion of wealth that should be invested in the Merton (1971) portfolio at time t is given by $1_n'\Sigma^{-1}(\alpha - r(t)1_n)/\delta$. This is shown to be decreasing as a function of the interest rate, which is quite natural. The purpose of the second fund is to hedge against changes in the interest rate. The investor's demand for the second and third funds depend on the market coefficients in a somewhat complex manner, but it should be noted that the optimal proportions depend randomly on $r(t)$ and t only. Moreover, the optimal proportion in the second fund tends to zero as t tends to T, which is also quite natural.

A major drawback of the Gaussian interest rate models considered in Sørensen (1999), Canestrelli and Pontini (2000), Bouiler et al. (2001), Korn and Kraft (2001), and Jensen and Sørensen (2001), is that they allow for negative interest rates. Deelstra et al. (2000) study the pure investment problem in a setup with a Cox-Ingersoll-Ross interest rate model (Cox et al. (1985)), where r is governed by the dynamics

$$dr(t) = (\theta - \kappa r(t))\, dt - \sigma_r \sqrt{r(t)}\, dW_r(t),$$

where $\sigma_r, \theta, \kappa > 0$ are constants, and $2\theta \geq \sigma_r^2$ so that $r(t) > 0,\ \forall t \in [0, T]$. The market consists of a bank account, a T-bond, and a stock, with price dynamics given by

$$\begin{aligned} dS_0(t)/S_0(t) &= r(t)\, dt, \\ dS_1(t)/S_1(t) &= (r(t) + \lambda_r \sigma_r h(t) r(t))\ dt + \sigma_r h(t) \sqrt{r(t)}\, dW_r(t), \\ dS_2(t)/S_2(t) &= (r(t) + \lambda_r \sigma_{2,r} r(t) + \lambda_s \sigma_{2,s})\ dt + \sigma_{2,r} \sqrt{r(t)}\, dW_r(t) + \sigma_{2,s}\, dW_s(t), \end{aligned}$$

respectively. Here, $h : [0, T] \to [0, \infty)$ is a certain differentiable function depending on the model coefficients with $h(T) = 0$. The market price of interest rate risk is $\lambda_r \sqrt{r(t)}$, whereas the market price of stock risk is λ_s; both λ_r and λ_s are constants. An explicit solution is obtained in the CRRA utility case (still with $\delta > 0$ denoting the CRRA coefficient), under certain conditions on the risk aversion parameter. The optimal portfolio is given by the bond and stock proportions

$$\begin{aligned} w_1(t) &= \frac{\lambda_r \sigma_{2,s} - \lambda_s \sigma_{2,r}}{\delta \sigma_{2,s} \sigma_r h(t)} + \frac{k_\delta(t)}{h(t)}, \\ w_2(t) &= \frac{\lambda_s}{\delta \sigma_{2,s}}, \end{aligned}$$

where $k_\delta : [0, T] \to \mathbb{R}$ is a certain differentiable function, depending on δ (and the model coefficients), with $k_\delta(T) = 0$. It is seen that the proportion held in the stock is again the same as in the case with constant interest rate, in particular it is time independent. The first term in the bond proportion is similar to the first term in Sørensen (1999), but the hedge term $k_\delta(t)/h(t)$ is no longer constant. Although $\lim_{t \nearrow T} (k_\delta(t)/h(t))$ exists and is finite, the bond position once again explodes at T. However, as in Bouiler et al. (2001), the optimal portfolio can be expressed in terms of portfolio proportions in the bank account, the stock, and a "rolling bond" with price dynamics given by

$$dS_3(t)/S_3(t) = (r(t) + \lambda_r \sigma_{3,r} r(t))\, dt + \sigma_{3,r} \sqrt{r(t)}\, dW_r(t),$$

for some $K > 0$. Then $w_1(t) = 0$, $w_2 = \lambda_{2,s}/\sigma_{2,s}$ (as above), $w_3(t) = k_\delta(t)\sigma_r/\sigma_{3,r} + (\lambda_r \sigma_{2,s} - \lambda_s \sigma_{2,r})/(\delta \sigma_{2,s} \sigma_{3,r})$, and $w_0(t) = 1 - \sum_{i=1}^{3} w_i(t)$ (or, equivalently, $w_1(t) =$

$k_\delta/h(t)$, $w_2 = \lambda_{2,s}/\sigma_{2,s}$, $w_3(t) = (\lambda_r\sigma_{2,s} - \lambda_s\sigma_{2,r})/(\delta\sigma_{2,s}\sigma_{3,r})$, and $w_0(t) = 1 - \sum_{i=1}^{3} w_i(t))$. It is shown that $w_1(t) + w_3(t)$ is increasing in t if and only if $\delta < 1$. Thus, low risk averters will gradually take out money from the bank and increase their position in bonds, whereas high risk averters ($\delta > 1$) will do the opposite. Note, however, that investing in the "rolling bond" close to maturity means going very short in cash and very long in zero coupon bonds, as indicated earlier.

Let us mention a few other papers adressing investment problems in models with more or less specific stochastic interest rate models. Deelstra et al. (2003) study the problem in a generalized complete term structure model that includes the Vasicek and Cox-Ingersoll-Ross models as special cases. Munk and Sørensen (2004) study the problem in a more general, but still complete, term structure model. Finally, Korn and Kraft (2004) point out some problems arising in investment problems with stochastic interest rates if the technical conditions for optimality are not satisfied.

2.9 Effects of inflation

In the investment problems considered so far the objectives have been maximization of expected utility of terminal wealth. At the end of the day, however, what really matters is not so much the level of wealth, but rather the purchasing power of wealth. It is therefore very reasonable to take inflation into account.

Brennan and Xia (2002) study an investment problem under inflation. There is a money market account, a stock, and two nominal zero-coupon bonds with different maturities (both exceeding the investor's time horizon. The price level is modelled as a diffusion process Π with dynamics given by

$$d\Pi(t)/\Pi(t) = \pi(t)\,dt + \sigma_{\Pi S}\,dW_S(t) + \sigma_{\Pi r}\,dW_r(t) + \sigma_{\Pi\pi}\,dW_\pi(t) + \sigma_{\Pi u}\,dW_u(t),$$

where $\pi(t)$ is the instantaneous *expected* rate of inflation, W_S, W_r, W_π, and W_u are Brownian motions, and $\sigma_{\Pi S}, \sigma_{\Pi r}, \sigma_\pi$, and $\sigma_{\Pi u}$ are constant correlation coefficients. The processes W_S, W_r, and W_π are in general correlated, and they determine the random development of the stock, the real interest rate, denoted by r, and the process π, respectively. The process W_u is an independent risk source. The processes π and r are both modelled as Ornstein-Uhlenbeck processes, with W_π and W_r as the respective "driving" Brownian motions. The stock price process is a geometric Brownian motion "driven" by W_S.

The model is incomplete whenever $\sigma_{\Pi u} \neq 0$, as there is no (locally) risk-free asset in real terms. In real terms, the model is thus a special case of the model studied by Merton (1971) without a risk-free asset. The investor has the objective of maximizing the expected utility, according to a CRRA utility function, of his terminal *real* wealth. Thus, the results of Merton (1971) in the case of CRRA utility can be applied to obtain an explicit solution.

Corresponding to the results obtained in the models with stochastic interest rates of the previous section, the optimal relative portfolio in the stock and the two bonds can be expressed as a sum of the Merton (1971) portfolio and two hedging portfolios aimed at hedging the risks pertaining to real interest rate and the rate of inflation. Akin to the result obtained by Karatzas et al. (1991) in the special case with totally unhedgeable coefficients, the optimal portfolio does not depend on $\sigma_{\Pi u}$, the volatility corresponding to the totally unhedgeable risk, i.e., this risk is ignored. The parts of the inflation risk that can be partially hedged should be treated similarly as other types of risks in a standard Brownian motion model with constant coefficients.

2.10 Some concluding remarks

We end this survey by offering a few concluding remarks. However, most of the abovementioned results are discussed at length by the authors who have obtained them, and we shall not go into a detailed and exhaustive discussion.

For long-term investors, such as pension savers, the classical single-period approach is insufficient; multi-period (with many periods) or continuous-time models are called for. In particular, continuous-time models with Brownian motions as basic building blocks, which are the most widely used ones in the modern literature, can often yield explicit results on optimal investment; Merton's results constitute the primary example.

Certain specific aspects concerning pension saving should be taken into account. For example, minimum guarantees are often involved, and the long-term nature of pension schemes calls for models with a stochastic interest rate in order to be sufficiently realistic. However, models with a constant interest rate can be considered very reasonable, even with long time horizons, if the interest rate is interpreted as the *real* interest rate, rather than the nominal one, and the wealth and the asset prices are taken to be denominated in *real* terms, as they arguably should be (see also Section 2.9). With this interpretation the assumption is that the real interest rate, rather than the nominal one, is constant, which is more plausible.

As pointed out, the choice of utility function plays an extremely important role for the optimal strategy. In particular, what is optimal cannot be stated without reference to the underlying utility function employed. This also goes for the many interesting problems that have been left out of this survey. One particularly nice and simple class of utility functions, which (as we have seen above) has very desirable features, is the class of CRRA utility functions, possibly extented to allow for minimum guarantees. It is therefore no coincidence that these utility functions are the most widely used ones in the literature.

Chapter 3

Optimal Bonus Strategies in Life Insurance: The Markov Chain Interest Rate Case

Sections 3.1-3.5 of this chapter constitute an adapted version of Nielsen (2005), with only minor corrections and modifications. Section 3.6 has been added for this thesis and thus does not appear in Nielsen (2005).

We study the problem of optimal redistribution of surplus in life and pension insurance when the interest rate is modelled as a continuous-time Markov chain with a finite state space. We work with traditional participating life insurance policies with payments consisting of a specified contractual payment stream and an unspecified additional bonus payment stream. Our model allows for interest rates below the technical interest rate. We apply stochastic control techniques in our search for optimal strategies, and we prove the dynamic programming principle for our particular type of problem. Furthermore, we state and prove a verification theorem and obtain an explicit solution that leads to a characterization of optimal strategies, indicating that some widely used redistribution schemes are suboptimal.

3.1 Introduction

A participating life or pension insurance contract specifies, at the time of issue, a premium plan and a fixed stream of (so-called) guaranteed benefits, both depending randomly only on the future course of life of the insured. On top of this the insured is entitled to additional *bonus* benefits, which, on the contrary, are determined by the company currently during the policy term as the systematic (portfolio) surplus emerges, as a consequence of the fact that all policies are charged (implicitly) with a safety loading.

The bonus payments typically contribute significantly to the total benefits and thus constitute one of the most important decision problems in traditional life and pension insurance, because it is up to the company to decide on the dividend/bonus plan (the distinction between *dividends* and *bonuses* is unimportant at this point and will be explained later, although we make clear that dividends go to the insured, not to owners or shareholders). Two concerns are involved: On the one hand, the company is urged to hand out dividends to the insured on a regular basis, partly to satisfy current customers (the insured) and perhaps even try to gain shares in the competitive market of the insurance business, and partly by legislative demands. On the other hand, the companies must see to it that they do not hand out more than they can afford, that is, they should always be able to meet all future obligations. It should be noted that dividends and bonuses, which have been credited/paid out to a policy at one stage, cannot be reclaimed later on.

The purpose of this chapter is to formulate and solve the problem of designing an optimal dividend plan as a control problem, taking into account the abovementioned concerns, in a stochastic interest rate environment. The paper is to some extent related to the comprehensive existing literature on optimal investment and consumption in continuous-time models, initiated by the classical papers Merton (1969) and Merton (1971). Virtually countless different variations and generalizations of investment/consumption problems have been studied since then, see, e.g. Duffie (1996), Ch. 9, or Karatzas and Shreve (1998), Sec. 3.11, for extensive surveys.

The problem studied in this paper differs from the vast majority of this literature in several respects. The interest rate (intensity) is modelled as a continuous-time Markov chain with a finite state space, and we consider the corresponding (locally) risk-free asset as the only investment opportunity and thus deal with the consumption/dividend aspect only. This means that the process under control is not a diffusion, but a PDP (piecewise deterministic process). Furthermore, our "agent" faces a certain non-hedgeable stream of positive and negative payments, but has no initial endowment.

Our objective (formulated in Section 3.3) constitutes another distinction, as it resembles the type of objectives employed in *singular control problems*. This, combined with the fact that our purpose is to find an optimal dividend strategy, leads to a connection to the class of problems known as dividend optimization problems, see e.g. Radner and Shepp (1996), or Taksar (2000) for a survey focusing on insurance companies. Such problems, aiming at maximizing (in some sense) the value of future dividends to shareholders for a firm with some initial reserve and a certain stochastic revenue process (typically a diffusion) through control of dividends and possibly risk exposure and/or investments, are often formulated as singular control problems. However, the connection is not as apparent as it may seem. Firstly, dividends play an entirely different role in this paper (as mentioned above). Secondly, our "revenue process" is not a diffusion, and, as we shall see, the optimal control process therefore need not have a singular component. Thirdly, our problem has a known, finite horizon, whereas the time horizon in most dividend optimization problems is either infinity or the (random) time of ruin.

Lakner and Slud (1991) study a problem akin to ours: They consider an agent, who only invests in a locally risk-free asset with an interest rate intensity that is adapted to a point process, and who seeks to optimize his consumption stream. However, the objective has a distinctly different form from ours, and there are no exogenous payment streams involved.

There exists a literature on optimization problems of various kinds within life insurance and/or pension funding. Optimal contribution (premium) rates for pension funds with *defined benefit* (DB) schemes are studied in Haberman and Sung (1994) and Chang et al. (2003). This problem is extended to include the problem of optimal asset allocation in Cairns (2000), Josa-Fombellida and Rincon-Zapatero (2001) and Taylor (2002). As for the contribution rates, the objectives in those papers are to minimize a certain loss function measuring (in various ways) the total sum of squared distances to certain targets of the controllable contribution rate and the aggregate fund level. A different approach, also within DB models, is taken in Haberman (1994) and Haberman and Wong (1997), where the aim is to find optimal time periods for distributing surpluses or deficits in the fund level through adjustment of the contribution rates.

For *defined contribution* schemes optimal investment strategies constitute the main issue, and this has been studied in Bouiler et al. (2001), Deelstra et al. (2003), Devolder et al. (2003), Vigna and Haberman (2001) and Haberman and Vigna (2002).

The present study is distinct from all the abovementioned: Firstly, we consider the dividend problem only, as we take the returns from investments as exogenously given (although stochastic). Secondly, our modelling approach is different, as we consider an individual (generic) policy with certain contractual payments and consider the associated surplus process throughout the policy term, whereas the abovementioned papers largely take more of a bird's eye view and consider the dynamics of the aggregate fund level without bringing the life history stochastics of the individual policies to the surface. Thirdly, our objective is (subjective) value maximization, not risk minimization.

The rest of the chapter is organized as follows: We briefly introduce the basic setup in Section 3.2 and then go on to state the optimal stochastic control problem in Section 3.3. We present theoretical results in Section 3.4 and obtain an explicit characterization of optimal dividend strategies. Section 3.5 contains main discussions and general conclusions. Finally, Section 3.6 contains generalizations of the results.

A quick word on notation: Integrals of the form $\int_{[t,t']} f_s \, dA_s$ for functions or processes (f predictable, A FV and right-continuous) are interpreted as $f_t \Delta A_t + \int_{(t,t']} f_s \, dA_s$, with $\Delta A_t = A_t$ if A_s is not defined for $s < t$. Correspondingly, $\int_{(t,t')} f_s \, dA_s$ means $\lim_{\tau \nearrow t'} \int_{(t,\tau]} f_s \, dA_s$ and so on.

3.2 Basics

A. The insurance policy and the contractual payments.
We work throughout on the basis of the general model in Norberg (1999), which is a basic reference for notions such as surplus, bonus and dividends, although we make specializations for the purpose of our study. In this section we briefly recapitulate the notions and results used in the following, referring to Norberg (1999) for details and discussions.

Consider a generic multi-state life insurance policy issued at time 0 and terminating at time $T > 0$. The state of the policy at time $t \in [0, T]$ is denoted by Z_t, and $Z = (Z_t)_{t\in[0,T]}$ is taken to be a right-continuous Markov process defined on some probability space $(\Omega, \mathcal{F}, \mathrm{P})$, with a finite state space $\mathcal{Z} = \{1, \ldots, k\}$ and with a finite number of jumps in $[0, T]$. It is assumed that the initial state is 1, $Z_0 = 1$.

The contractual payments associated with the policy, i.e. premiums and *guaranteed* benefits, are represented by a stochastic *payment process* $B = (B_t)_{t\in[0,T]}$ specifying the accumulated payments from the company to the insured, that is, B_t denotes benefits less premiums payable in $[0, t]$, $t \in [0, T]$. It is assumed that B has a form that allows of continuous payments during sojourns in policy states as well as lump sum payments at a finite number of fixed times and in connection with jumps between policy states. Thus, B has an absolutely continuous component and a pure jump component, and all randomness associated with the payments dB_t in a small time interval $(t, t + dt]$ is due to the (random) behaviour of Z in that interval.

Even though this formalization of the payment process opens up for a wide range of insurance types, the problem that we deal with in this paper is more relevant for some types of insurance than others. Typically we have in mind a pension scheme with the following characteristics: The policyholder pays a level premium continuously for a certain number of years from the time of issue, possibly until the time of retirement, and receives a continuous life annuity, possibly temporary, commencing at a specific time (e.g. the time of retirement); with all payments being contingent on the policyholder being alive.

The payment process B must satisfy the actuarial *equivalence principle* under a technical *first order model* that lays down a (technical) interest rate $r^* \in \mathbb{R}$ and a (technical) distribution of Z. We assume for simplicity that the distribution of Z under the first order model coincides with its distribution under P, which, in the terminology of Norberg (1999), corresponds to the *second order model.* This assumption can be justified by the fact that the interest rate risk is considerably larger than the risk connected to systematic changes in, say, mortality, for pension schemes as the one outlined above. Thus, upon introducing the first order state-wise (prospective) reserves V^{*j}, $j \in \mathcal{Z}$, in this setting given by

$$V_t^{*j} = \mathrm{E}\left(\int_{(t,T]} e^{-(\tau-t)r^*}\, dB_\tau \,\middle|\, Z_t = j \right),$$

for $t \in [0, T]$, $j \in \mathcal{Z}$, the equivalence principle can be expressed by the equation

$$V_0^{*1} = -B_0.$$

It is tacitly assumed here, and throughout, that $\mathrm{E}\,(\mathrm{Var}(B)_T) < \infty$, where $\mathrm{Var}(B)$ is the variation process of B. This ensures that the (state-wise) reserves are well-defined and bounded. Moreover, it is assumed that $V_t^{*j} \geq 0$, $\forall t \in [0, T]$, $j \in \mathcal{Z}$, as should always be the case in practice.

B. The mean portfolio surplus.
The first order model represents a prudent, conservative estimate of the economic development. The real development, in contrast, is represented by the probability measure P introduced above.

It is assumed that the company's investment portfolio bears interest with intensity $r = (r_t)_{t \in [0,T]}$, which is a stochastic process (defined on $(\Omega, \mathcal{F})$), independent of Z under P. The information about the development of r is formalized by the filtration $\mathbb{G} = \{\mathcal{G}_t\}_{t \in [0,T]}$, i.e. $\mathcal{G}_t = \sigma\{r_\tau, \tau \in [0, t]\}$.

The *mean portfolio surplus* at time $t \in [0, T]$ is defined by

$$S_t = \mathrm{E}\left(S_t^{\mathrm{ind}} \,\middle|\, \mathcal{G}_t \right), \tag{2.1}$$

where S_t^{ind} is the individual surplus (based on full information about the policy) defined naturally by

$$S_t^{\mathrm{ind}} = -\int_{[0,t]} e^{\int_\tau^t r}\, dB_\tau - V_t^{*Z_t}, \tag{2.2}$$

i.e. as the difference between the value of the net income from the policy in $[0, t]$, compounded with the experienced interest rate, and the first order reserve at time t. Here and throughout, $e^{\int_s^t r}$ is used as short-hand notation for $e^{\int_s^t r_\tau \, d\tau}$.

It can now easily be shown (see Norberg (1999)) that for each $t \in [0, T]$ we have

$$S_t = \int_0^t e^{\int_\tau^t r}\, c_\tau \, d\tau, \tag{2.3}$$

where, for $t \in [0, T]$,

$$c_t = \sum_{j \in \mathcal{Z}} \mathrm{P}\left(Z_t = j\right)\, \left(r_t - r^*\right) V_t^{*j} = \left(r_t - r^*\right) V_t^*, \tag{2.4}$$

with

$$V_t^* = \sum_{j \in \mathcal{Z}} \mathrm{P}\left(Z_t = j\right)\, V_t^{*j} \tag{2.5}$$

denoting the mean first order reserve at time t. This form shows how the surplus emerges, with a straightforward interpretation of the *contribution rate* c: It is the difference between the rates of actually earned interest and technically accrued interest on the mean first order reserve.

Equivalently, from (2.1) we easily obtain

$$S_t = -\int_{[0,t]} e^{\int_\tau^t r}\, db_\tau - \int_{(t,T]} e^{-(\tau-t)r^*}\, db_\tau,$$

where b is defined by $b_t = \mathrm{E}\,(B_t)\,,\ t \in [0,T]$. Thus, the mean portfolio surplus is nothing but the difference between the compounded value of past net income and the present value of future net outgoes (evaluated on the first order basis) corresponding to the (artificial) deterministic payment stream b.

C. The dividend reserve.
The surplus (or a part of it) is repaid to the insured by crediting *dividends* to his or her account. The stream of dividends is represented by the stochastic process $D = (D_t)_{t\in[0,T]}$, and as for the payment process B, D_t denotes the accumulated dividends during $[0,t]$, $t \in [0,T]$. Dividends can only be credited on the basis of the information that is currently known, and since we work with the mean portfolio surplus we must therefore require that D be $\mathbb{G}$-adapted. This means that D represents the mean dividends *per policyholder* in the entire portfolio (including those who have died). Thus, the policy that we consider should be thought of as a truly generic "mean" policy representing the portfolio, not as a specific policy. By taking this approach we avoid having to deal with policy state variables in the control problem addressed below.

Furthermore, dividends must be non-negative and cannot be reclaimed once they have been credited, so D must be non-negative and non-decreasing. Finally, D is taken to be right-continuous.

The *dividend reserve* U_t at time $t \in [0,T]$ is defined as the value at time t of past contributions less dividends, compounded with interest, i.e.

$$U_t = \int_{[0,t]} e^{\int_\tau^t r}\,(c_\tau\, d\tau - dD_\tau). \tag{2.6}$$

Credited dividends may be used in various ways and need not be paid out immediately. The actual payouts, in turn, are called *bonuses.* In some simple cases, e.g. if dividends are paid out currently as they are credited, then the distinction between dividends and bonus payments is not important. In most cases it is, however, and the distribution and repayment of surplus can then be viewed as a two-step procedure: First, dividends are determined and distributed among the policies according to some *dividend scheme.* Then, according to a *bonus scheme*, it is decided how the distributed dividends are transformed into bonus payments. We refer to Norberg (1999, 2001) for more detailed discussions on this quite flexible system.

From (2.6) it is easily seen that U develops as

$$U_0 = -D_0, \tag{2.7}$$

$$dU_t = r_t U_t\, dt + c_t\, dt - dD_t, \tag{2.8}$$

for $t \in (0, T]$. These equations have a straightforward interpretation. It is seen that for any $0 \leq s \leq t \leq T$ we have

$$U_t = U_s \, e^{\int_s^t r} + \int_{(s,t]} e^{\int_\tau^t r} \, (c_\tau \, d\tau - dD_\tau). \tag{2.9}$$

It is a widely accepted actuarial principle, stated in Norberg (1999), that in order to reestablish equivalence at time T, the equivalence principle should be exercised conditionally, given the experienced development. In our setup this amounts to requiring that

$$U_T = 0. \tag{2.10}$$

However, this is not possible unless $U_{T-} \geq 0$, and in this paper we shall work with an interest rate model under which the event that $U_{T-} < 0$ occurs with strictly positive probability, as is arguably also the case in the real world. We shall thus ignore the requirement (2.10), since it plays no role in the problem formulation anyway. For further discussions of this we refer to Section 3.5.

3.3 Statement of the problem

A. Interest rate model.

We adopt the Markov chain interest rate model studied in Norberg (1995) and Norberg (2003). Thus, only a finite number of interest rate levels are possible. More specifically, the interest rate (intensity) process is defined by

$$r_t = r^{Y_t} = \sum_{e \in \mathcal{Y}} I_{(Y_t = e)} \, r^e, \; t \in [0, T], \tag{3.1}$$

where $Y = (Y_t)_{t \in [0,T]}$ is a right-continuous pure jump Markov process with a finite state space $\mathcal{Y} = \{1, \ldots, q\}$, defined on $(\Omega, \mathcal{F}, P)$, which we now assume is complete.

All jump intensities λ^{ef}, $e, f \in \mathcal{Y}, e \neq f$, exist, and for simplicity they are taken to be constant. Put $\mathcal{Y}^e = \mathcal{Y} \setminus \{e\}$ and $\lambda^{e\cdot} = \sum_{f \in \mathcal{Y}^e} \lambda^{ef}$, $e \in \mathcal{Y}$, and denote by N^{ef}, $e, f \in \mathcal{Y}, e \neq f$, the counting processes counting the jumps from state e to f. We assume that $r^e \geq 0$, $e \in \mathcal{Y}$, and we put $\overline{r} = \max_{e \in \mathcal{Y}} r^e$ and $\underline{r} = \min_{e \in \mathcal{Y}} r^e$.

Let $\mathbb{G} = (\mathcal{G}_t)_{0 \leq t \leq T}$ denote the natural filtration of Y augmented by the null sets of $\mathcal{F}$. Then $\mathbb{G}$ is right-continuous (see e.g. Davis (1993)).

B. The Markov process under control.

The uncontrolled process $(Y, U) = \{(Y_t, U_t)\}_{t \in [0,T]}$ assumes its values in the state space $\mathcal{Y} \times \mathbb{R}$. It is a *piecewise deterministic process* (PDP) because Y is a pure jump process and between jumps of Y, U is governed by the ordinary (deterministic) differential equation

$$dU_t = U_t \, r^{Y_t} \, dt + c_t^{Y_t} \, dt,$$

where

$$c_t^e = (r^e - r_t^*) V_t^*, \; t \in [0, T], \; e \in \mathcal{Y}.$$

Control of the dividend reserve process is carried out through the dividend process D, yielding the (stochastic) dynamics (2.8).

A key reference in the theory of control of PDP's is Davis (1993), which, for this chapter, however, has served primarily as a source of inspiration, because the problem at hand does not quite fit into its framework (e.g. due to the general form of D).

We denote by E the state space for the triple process $\{(t, Y_t, U_t)\}_{t\in[0,T]}$, that is, $E = [0,T] \times \mathcal{Y} \times \mathbb{R}$. For notational convenience we also define $E^0 = [0,T) \times \mathcal{Y} \times \mathbb{R}$.

Note that the policy process Z no longer plays any role, because we work with the mean portfolio surplus.

C. Objective.
We want to design a control strategy that is optimal in some sense, and to this end we wish to maximize

$$\mathrm{E}\left(\int_{[0,T)} e^{-\int_0^s r}\, dD_s + e^{-\int_0^T r}\psi(U_{T-})\right), \tag{3.2}$$

the sum of the expected discounted dividends and the expected discounted "utility" of the dividend reserve immediately before T as measured by some concave and strictly increasing function $\psi : \mathbb{R} \to \mathbb{R}$. Note that we only integrate over $[0,T)$ because any terminal dividend lump sum ΔD_T is taken directly from U_{T-} and thus (indirectly) measured by ψ. In other words, all control action takes place in $[0,T)$. We refer to Section 3.5 for a discussion of this objective as well as principles regarding U_{T-}.

3.4 Control theory

A. Dynamic programming framework.
We apply the method of dynamic programming, and to this end we need to consider each point $(t,e,u) \in E^0$ as an (imaginary) initial point for the controlled process. For technical reasons (needed in the proof of the dynamic programming principle, Theorem 3.4.5 below) we set up the dynamic programming framework rigorously, as in Fleming and Rishel (1975) or Yong and Zhou (1999).

For any $(t,e) \in [0,T)\times\mathcal{Y}$ we define $\mathcal{D}(t,e)$ as the set of all 5-tuples $(\Omega, \mathcal{F}, \mathrm{P}, Y, D)$ such that $(\Omega, \mathcal{F}, \mathrm{P})$ is a complete probability space, $Y = (Y_s)_{s\in[t,T]}$ is a continuous-time Markov chain defined on $(\Omega, \mathcal{F}, \mathrm{P})$ with state space $\mathcal{Y}$ and transition intensities λ^{ef}, $e, f \in \mathcal{Y}, e \neq f$, and with $Y_t = e$ (P-a.s.), and $D = (D_s)_{s\in[t,T]} : \Omega \times [t,T] \to [0,\infty)$ is a measurable, $\mathbb{G}$-adapted, right-continuous and increasing process with $\mathrm{E}_{(t,e)}(D_T) < \infty$. Here, $\mathbb{G} = (\mathcal{G}_s)_{s\in[t,T]}$ denotes the filtration generated by Y (augmented by all the P-null sets of $\mathcal{F}$), $r = (r_s)_{s\in[t,T]}$ is defined as in (3.1) on $[t,T]$, and $\mathrm{E}_{(t,e)}(\cdot)$ is the expectation operator corresponding to P, equipped with the subscript (t,e) to emphasize the initial point.

For notational simplicity we shall usually write $D \in \mathcal{D}(t,e)$ as shorthand for $(\Omega, \mathcal{F}, \mathrm{P}, Y, D) \in \mathcal{D}(t,e)$ when no confusion can arise. Now, for any $(t,e,u) \in E^0$

and any $D \in \mathcal{D}(t, e)$ we define

$$U_{t'}^{(t,e,u,D)} = ue^{\int_t^{t'} r} + \int_{[t,t']} e^{\int_\tau^{t'} r}(c_\tau \, d\tau - dD_\tau), \ t' \in [t, T], \tag{4.1}$$

i.e. the dividend reserve at time t' corresponding to the starting point (t, e, u) and the dividend strategy D (see (2.9)), and

$$\Phi^D(t, e, u) = \mathrm{E}_{(t,e)} \left(\int_{[t,T)} e^{-\int_t^s r} \, dD_s + e^{-\int_t^T r} \, \psi \left(U_{T-}^{(t,e,u,D)} \right) \right). \tag{4.2}$$

In addition we put $\Phi^D(T, e, u) = \psi(u)$, $(e, u) \in \mathcal{Y} \times \mathbb{R}$, and we have thereby defined the *performance function* $\Phi^D : E \to \mathbb{R}$. Note that in (4.1) u acts only as a parameter, and in (4.2) it therefore appears only in $U_{T-}^{(t,e,u,D)}$, not in the subscript of the expectation operator.

We define correspondingly the *value function* $\Phi : E \to \mathbb{R}$ by

$$\Phi(t, e, u) = \sup_{D \in \mathcal{D}(t,e)} \Phi^D(t, e, u), \ (t, e, u) \in E^0,$$

which is well defined by Assumption 3.4.1 below, and $\Phi(T, e, u) = \psi(u)$, $(e, u) \in \mathcal{Y} \times \mathbb{R}$. For a given starting point $(t, e, u) \in E^0$ the objective is to find an optimal control, i.e. a control $D \in \mathcal{D}(t, e)$ satisfying

$$\Phi(t, e, u) = \Phi^D(t, e, u),$$

and for our original problem the objective is thus expressed by letting $t = 0, e = Y_0$ and $u = 0$.

Assumption 3.4.1 *We assume throughout that*

(i) $\psi \in C^1(\mathbb{R})$,

(ii) there exists a $u_0 \in \mathbb{R}$ *such that* $\psi'(u_0) > 1$,

(iii) $L := \lim_{u \to -\infty} \psi'(u) < \infty$.

Condition (i) can be weakened, but it is enforced because it simplifies the results and proofs in the following. Condition (ii) ensures that the control problem is well posed: Under this condition it is fairly easy to show that the value function is finite. Conversely, if it is not fulfilled, then any control $D' \in \mathcal{D}(t, e)$ can be dominated by a "larger" one, that is, if we define $D'' := D' + K$ for some constant $K > 0$ then $\Phi^{D''}(t, e, u) \geq \Phi^{D'}(t, e, u)$, leading to a highly counter-intuitive dividend strategy. The interpretation of this condition is that ultimate losses must not be punished too lightly. Condition (iii) is purely technical and is needed in certain proofs in the following. This condition ensures that ψ is Lipschitz continuous (with Lipschitz constant L). However, as we shall see, it has no effect on the optimal strategies (if properly applied).

We shall need the following technical result. For $t \in [0,T]$, let $A[t,T]$ denote the space of right-continuous functions $d : [t,T] \to \mathcal{Y}$, equipped with the σ-algebra generated by the coordinate projections $\{(d \mapsto d_s)_{d\in A[t,T]}\}_{s\in[t,T]}$.

Lemma 3.4.2 *Let $D \in \mathcal{D}(t,e)$ for fixed $(t,e) \in [0,T) \times \mathcal{Y}$. There exists a measurable mapping $\eta : [t,T] \times A[t,T] \to [0,\infty)$ such that*

$$D_s(\omega) = \eta\left(s, Y_{\cdot\wedge s}(\omega)\right),\ \forall s \in [t,T],\ \mathrm{P}-a.s.$$

Proof. Since D is $\mathbb{G}$-adapted, there exists for each $s \in [t,T]$ a measurable mapping $\eta_s : A[t,T] \to [0,\infty)$ such that $D_s = \eta_s(Y_{\cdot\wedge s})$ $\mathrm{P}-a.s.$ Now, for $n \geq 1$, put $t_k = t + (T-t)k/2^n$, $k = 0,\dots,2^n$, and define $\eta^n : [t,T] \times A[t,T] \to [0,\infty)$ by

$$\eta^n(s,d) = 1_{\{t\}}(s)\eta_t(d) + \sum_{k=1}^{2^n} 1_{(t_{k-1},t_k]}(s)\eta_{t_k}(d),\ (s,d) \in [t,T] \times A[t,T].$$

Each η^n is clearly measurable, and we have $D_{t_k} = \eta^n(t_k, Y_{\cdot\wedge t_k}),\ \forall k = 0,\dots,2^n,\ \mathrm{P}-a.s.$ Since D is also right-continuous, the mapping $\eta : [t,T] \times A[t,T] \to [0,\infty)$ given by

$$\eta(s,d) = \lim_{n\to\infty} \eta^n(s,d),\ (s,d) \in [t,T] \times A[t,T],$$

is well defined and fulfills the assertion. $\square$

B. Basic properties of the value function.
The following results are hardly surprising.

Proposition 3.4.3 *For any $(t,e) \in [0,T] \times \mathcal{Y}$, $\Phi(t,e,\cdot)$ is increasing and concave as a function of $u \in \mathbb{R}$.*

Proof. Let $(t,e) \in [0,T) \times \mathcal{Y}$ (for $t = T$ the assertions are obvious). It is obvious that $\Phi(t,e,\cdot)$ is increasing. To prove concavity, let $u' < u'' \in \mathbb{R}$, and let $D', D'' \in \mathcal{D}(t,e)$. Let $0 \leq \gamma \leq 1$, and put

$$u^\gamma = \gamma u' + (1-\gamma)u''.$$

We now claim that there exists a $D^\gamma \in \mathcal{D}(t,e)$ such that

$$\Phi^{D^\gamma}(t,e,u^\gamma) \geq \gamma\Phi^{D'}(t,e,u') + (1-\gamma)\Phi^{D''}(t,e,u''). \tag{4.3}$$

Then, for any $\varepsilon > 0$, by choosing D' and D'' such that $\Phi^{D'}(t,e,u') \geq \Phi(t,e,u') - \varepsilon$ and $\Phi^{D''}(t,e,u'') \geq \Phi(t,e,u'') - \varepsilon$, the inequality (4.3) yields

$$\Phi(t,e,u^\gamma) \geq \gamma\Phi(t,e,u') + (1-\gamma)\Phi(t,e,u'') - \varepsilon,$$

from which concavity follows because ε is arbitrary.

To prove the existence of $D^\gamma \in \mathcal{D}(t,e)$ fulfilling (4.3) we may, by suitable use of Lemma 3.4.2, assume that the underlying probability spaces and Y-processes corresponding to D' and D'' are identical. Thus, we can define D^γ on this probability space by

$$D^\gamma = \gamma D' + (1-\gamma)D''.$$

It is now easily seen that

$$U_{T-}^{(t,e,u^\gamma,D^\gamma)} = \gamma U_{T-}^{(t,e,u',D')} + (1-\gamma)U_{T-}^{(t,e,u'',D'')},$$

and by concavity of ψ we therefore get (4.3). □

Proposition 3.4.4 *For any $(t,e) \in [0,T] \times \mathcal{Y}$ and any $D \in \mathcal{D}(t,e)$, $\Phi(t,e,\cdot)$ and $\Phi^D(t,e,\cdot)$ are Lipschitz continuous as functions of $u \in \mathbb{R}$.*

Proof. Let $(t,e) \in [0,T) \times \mathcal{Y}$ (for $t = T$ the assertions are obvious, cf. Assumption 3.4.1). Let $u' < u'' \in \mathbb{R}$, and let $D \in \mathcal{D}(t,e)$. We have

$$\begin{aligned}
\left|\Phi^D(t,e,u'') - \Phi^D(t,e,u')\right| &\le \left|\mathrm{E}_{(t,e)}\left(e^{-\int_t^T r}\left[\psi(U_{T-}^{(t,e,u'',D)}) - \psi(U_{T-}^{(t,e,u',D)})\right]\right)\right| \\
&\le \mathrm{E}_{(t,e)}\left|\psi(U_{T-}^{(t,e,u'',D)}) - \psi(U_{T-}^{(t,e,u',D)})\right| \\
&\le Le^{T\overline{r}}|u'' - u'|,
\end{aligned}$$

which shows that $\Phi^D(t,e,\cdot)$ is Lipschitz continuous with Lipschitz constant $Le^{T\bar{r}}$. Since $D \in \mathcal{D}(t,e)$ was arbitrary, we conclude that the same goes for $\Phi(t,e,\cdot)$. □

Measurability of Φ follows from the fact that also $\Phi(\cdot,e,u)$ is continuous on $[0,T)$ for any $(e,u) \in \mathcal{Y} \times \mathbb{R}$. We choose to omit the proof of this, however, since it is tedious and technically rather complicated, and since continuity in t on $[0,T)$ is intuitively obvious.

C. The dynamic programming principle.

The following theorem establishes the dynamic programming principle (DPP) for the control problem considered in this chapter. A rigorous proof covering this type of problem has not been found in the literature. The proof, which is rather technical, is highly inspired by the proof of Yong and Zhou (1999), Theorem 4.3.3.

Theorem 3.4.5 *For any $(t,e,u) \in E^0$ and any $t' \in [t,T]$,*

$$\Phi(t,e,u) = \sup_{D\in\mathcal{D}(t,e)} \mathrm{E}_{(t,e)}\left(\int_{[t,t')} e^{-\int_t^s r} dD_s + e^{-\int_t^{t'} r}\,\Phi\left(t', Y_{t'}, U_{t'-}^{(t,e,u,D)}\right)\right).$$

Proof. Let $(t,e,u) \in E^0$ and $t' \in [t,T]$. To ease notation we write E instead of $\mathrm{E}_{(t,e)}$ throughout this proof. For any $D \in \mathcal{D}(t,e)$,

$$\begin{aligned}
\Phi^D(t,e,u) &= \mathrm{E}\left(\int_{[t,t')} e^{-\int_t^s r} dD_s\right) \\
&\quad + \mathrm{E}\left(e^{-\int_t^{t'} r}\,\mathrm{E}\left(\int_{[t',T)} e^{-\int_{t'}^s r} dD_s + e^{-\int_{t'}^T r}\psi\left(U_{T-}^{(D,t,e,u)}\right)\middle|\, \mathcal{G}_{t'}^t\right)\right).
\end{aligned}$$

We now claim that we almost surely have

$$\mathrm{E}\left(\int_{[t',T)} e^{-\int_{t'}^{s} r} dD_s + e^{-\int_{t'}^{T} r}\psi\left(U_{T-}^{(D,t,e,u)}\right)\middle|\,\mathcal{G}_{t'}^{t}\right) \le \Phi\left(t', Y_{t'}, U_{t'-}^{(D,t,e,u)}\right), \quad (4.4)$$

leading to

$$\Phi^D(t,e,u) \le \mathrm{E}\left(\int_{[t,t')} e^{-\int_{t}^{s} r} dD_s + e^{-\int_{t}^{t'} r}\Phi\left(t', Y_{t'}, U_{t'-}^{(D,t,e,u)}\right)\right),$$

which in turn implies the inequality "$\le$". Although (4.4) may seem intuitively very reasonable, it takes a few lines of technicalities to prove it with some degree of rigor: Since Y and D are right-continuous, there exists a set $\Omega_0 \in \mathcal{F}$ with $\mathrm{P}(\Omega_0) = 1$ such that the following holds true: For any fixed $\omega_0 \in \Omega_0$, Y and D are (almost surely) deterministic on $[t,t']$ under $\mathrm{P}\left(\cdot\,|\,\mathcal{G}_{t'}^{t}\right)(\omega_0)$, where $\mathrm{P}\left(\cdot\,|\,\mathcal{G}_{t'}^{t}\right)$ denotes a regular conditional probability (of P) given $\mathcal{G}_{t'}^{t}$. And furthermore, by Lemma 3.4.2 we can deduce that under $\mathrm{P}\left(\cdot\,|\,\mathcal{G}_{t'}^{t}\right)(\omega_0)$, the process D restricted to $[t',T]$ is adapted to the process Y restricted to $[t',T]$, and the latter is a continuous-time Markov process starting in the state $Y_{t'}(\omega_0)$ and otherwise inheriting the distributional properties of Y. Thus, $(\Omega, \mathrm{P}\left(\cdot\,|\,\mathcal{G}_{t'}^{t}\right)(\omega_0), \mathcal{F}, Y|_{[t',T]}, D) \in \mathcal{D}(t', Y_{t'}(\omega_0))$. Furthermore, we have

$$U_{T-}^{(D,t,e,u)} = U_{T-}^{(D,t',Y_{t'},U_{t'-}^{(D,t,e,u)})}, \quad (4.5)$$

with $Y_{t'}$ and $U_{t'-}^{(D,t,e,u)}$ being deterministic under $(\Omega, \mathrm{P}\left(\cdot\,|\,\mathcal{G}_{t'}^{t}\right)(\omega_0), \mathcal{F})$. Substituting (4.5) into the left hand side of (4.4) and using the rule $\mathrm{E}\left(\cdot\,|\,\mathcal{G}_{t'}^{t}\right)(\omega_0) = \mathrm{E}^{\mathrm{P}\left(\cdot\,|\mathcal{G}_{t'}^{t}\right)(\omega_0)}(\cdot)$ yields the desired inequality.

Conversely, for any $\varepsilon > 0$, by Proposition 3.4.4 there exists $\delta > 0$ such that for any $e \in \mathcal{Y}$ and $u^1, u^2 \in \mathbb{R}$ with $|u^1 - u^2| < \delta$ we have

$$\left|\Phi^D(t',e,u^1) - \Phi^D(t',e,u^2)\right| + \left|\Phi(t',e,u^1) - \Phi(t',e,u^2)\right| \le \varepsilon,\ \forall D \in \mathcal{D}(t',e).$$

Let $(I^n)_{n\ge1}$ be a partition of $\mathbb{R}$ with each I^n being an interval of length $m(I^n) \le \delta$. Now choose, for each $n \ge 1$, a $u^n \in I^n$. For any $e \in \mathcal{Y}$ and $n \ge 1$ there exists $D^{e,n} \in \mathcal{D}(t',e)$ such that

$$\Phi^{D^{e,n}}(t',e,u^n) \ge \Phi(t',e,u^n) - \varepsilon.$$

Thus, for any $u \in I^n$ we have

$$\Phi^{D^{e,n}}(t',e,u) \ge \Phi^{D^{e,n}}(t',e,u^n) - \varepsilon \ge \Phi(t',e,u^n) - 2\varepsilon \ge \Phi(t',e,u) - 3\varepsilon.$$

For each $e \in \mathcal{Y}$ and $n \ge 1$ there exists by Lemma 3.4.2 a measurable mapping $\eta^{e,n} : [t'T] \times A[t',T] \to [0,\infty)$ such that $D_s^{e,n} = \eta^{e,n}(s, Y_{\cdot\wedge s})$, $\mathrm{P}-a.s.$ Now, take any control $D \in \mathcal{D}(t,e)$, and define a new control $\widetilde{D} \in \mathcal{D}(t,e)$ by $\widetilde{D}_s = D_s$ for $s < t'$ and

$$\widetilde{D}_s = D_{t'-} + \sum_{f\in\mathcal{Y}}\sum_{n\ge1} 1_{\{(e,I^n)\}}\left((Y_{t'}, U_{t'}^{(D,t,e,u)})\right)\eta^{e,n}\left(s, (Y_{\tau\wedge s})_{t'\le\tau\le T}\right),\ t' \le s \le T.$$

This new control is seen to be admissible, and

$$
\begin{aligned}
\Phi(t,e,u) &\geq \Phi^{\widetilde{D}}(t,e,u) \\
&= \mathrm{E}\left(\int_{[t,T)} e^{-\int_t^s r} d\widetilde{D}_s + e^{-\int_t^T r} \psi\left(U_{T-}^{(\widetilde{D},t,e,u)}\right)\right) \\
&= \mathrm{E}\left(\int_{[t,t')} e^{-\int_t^s r} dD_s\right) \\
&\quad + \mathrm{E}\left(e^{-\int_t^{t'} r}\, \mathrm{E}\left(\int_{[t',T)} e^{-\int_{t'}^s r} d\widetilde{D}_s + e^{-\int_{t'}^T r} \psi\left(U_{T-}^{(\widetilde{D},t,e,u)}\right) \middle| \mathcal{G}_{t'}^t\right)\right).
\end{aligned}
$$

Now, with arguments similar to the ones used in the proof of (4.4), we almost surely have

$$
\mathrm{E}\left(\int_{[t',T)} e^{-\int_{t'}^s r} d\widetilde{D}_s + e^{-\int_{t'}^T r} \psi\left(U_{T-}^{(\widetilde{D},t,e,u)}\right) \middle| \mathcal{G}_{t'}^t\right) \geq \Phi(t', Y_{t'}, U_{t'-}^{(D,t,e,u)}) - 3\varepsilon,
$$

and thus, by arbitraryness of $\varepsilon > 0$,

$$
\Phi(t,e,u) \geq \mathrm{E}\left(\int_{[t,t')} e^{-\int_t^s r} dD_s + e^{-\int_t^{t'} r} \Phi\left(t', Y_{t'}, U_{t'-}^{(D,t,e,u)}\right)\right),
$$

This proves the inequality "$\geq$", because $D \in \mathcal{D}(t,e)$ was arbitrary. □

D. Variational inequalities and heuristically optimal controls.
The value function can be viewed as a function with q components, namely the state-wise value functions $\Phi^e(\cdot,\cdot) = \Phi(\cdot, e, \cdot) : [0,T] \times \mathbb{R} \to \mathbb{R},\ e \in \mathcal{Y}$, corresponding to the different states in $\mathcal{Y}$. Dropping the generic argument (t,u) to ease notation, the HJB equation of dynamic programming becomes a system of variational inequalities for these functions, which can be expressed (using the notation ϕ for the generic function argument) by

$$
0 = \max\left(\frac{\partial}{\partial t}\phi^e - \phi^e\, r^e + \frac{\partial}{\partial u}\phi^e\,(r^e u + c_t^e) + \sum_{f \in \mathcal{Y}^e} \lambda^{ef}(\phi^f - \phi^e), 1 - \frac{\partial}{\partial u}\phi^e\right),\ e \in \mathcal{Y}, \tag{4.6}
$$

with the boundary condition

$$
\phi^e(T,u) = \psi(u),\ (e,u) \in \mathcal{Y} \times \mathbb{R}. \tag{4.7}
$$

A heuristic argument for this system of equations goes as follows: Assume for a moment that only absolutely continuous dividend processes are allowed, such that any control is of the form $dD_t = \delta_t\, dt$ for some *dividend rate* process δ. Then, assuming furthermore that $\Phi \in C^1(E^0)$, the dynamic programming principle leads to the HJB equation

$$\begin{aligned} 0 &= \frac{\partial}{\partial t}\Phi^e(t,u) - \Phi^e(t,u)r^e + \frac{\partial}{\partial u}\Phi^e(t,u)\left(r^e u + c_t^e\right) \\ &\quad + \sum_{f \in \mathcal{Y}^e} \lambda^{ef}\left(\Phi^f(t,u) - \Phi^e(t,u)\right) + \sup_{\delta \geq 0}\left(\delta - \frac{\partial}{\partial u}\Phi^e(t,u)\delta\right). \end{aligned} \tag{4.8}$$

It is seen that if $\frac{\partial}{\partial u}\Phi^{Y_t}(t,U_t) > 1$, the optimum is obtained by putting $\delta_t = 0$. However, if $\frac{\partial}{\partial u}\Phi^{Y_t}(t,U_t) < 1$ then no optimum exists, but obviously δ should be chosen as large as possible.

Now, allowing dividend strategies that are not absolutely continuous, we can in fact achieve an intuitively optimal strategy by including lump sum dividends when appropriate, i.e. when $\frac{\partial}{\partial u}\Phi^{Y_t}(t,U_t) < 1$. Since such a lump sum has an effect on the value equal to the jump size, it is seen that we actually cannot have $\frac{\partial}{\partial u}\Phi^e(t,u) < 1$, but rather must have $\frac{\partial}{\partial u}\Phi^e(t,u) = 1$ for u sufficiently large.

Thus, for any $(t,e,u) \in E^0$ we have

$$\begin{aligned} 0 &\geq 1 - \frac{\partial}{\partial u}\Phi^e, \\ 0 &\geq \frac{\partial}{\partial t}\Phi^e - \Phi^e r^e + \frac{\partial}{\partial u}\Phi^e\left(r^e u + c_t^e\right) + \sum_{f \in \mathcal{Y}^e}\lambda^{ef}\left(\Phi^f - \Phi^e\right), \end{aligned}$$

with at least one of the inequalities being an equality, and these variational inequalities are combined in (4.6).

E. A verification theorem.

We now state a verification theorem that identifies (4.6) as the appropriate equation to study and yields sufficient conditions for optimality, and we apply it in Paragraph F below to find an optimal control strategy. However, the classical notion of a solution turns out to be insufficient and therefore needs to be generalized slightly. Thus, by *a generalized solution to (4.6)-(4.7)* we shall mean a function $\phi = (\phi^e)_{e\in\mathcal{Y}} : E \to \mathbb{R}$ satisfying the following requirements:

(i) There exists a finite partition $0 = t_0 < t_1 < \ldots < t_n = T$ such that each ϕ^e is continuously differentiable and fulfills (4.6) on $(t_{i-1}, t_i) \times \mathbb{R}$ for $i = 1, \ldots, n$.

(ii) Each ϕ^e fulfills (4.7) and is continuous on $[0,T] \times \mathbb{R}$ except possibly on the subset $\{(T,u) : \psi'(u) < 1\}$.

A solution is necessarily unique.

Theorem 3.4.6 *If ϕ is a generalized solution to (4.6)-(4.7), which is concave in u, then $\phi \geq \Phi$. Furthermore, for any $(t,e,u) \in E^0$, if $D \in \mathcal{D}(t,e)$ and the corresponding dividend reserve process $U^D \equiv U^{(t,e,u,D)}$ satisfy (a.s.)*

$$\lim_{t \nearrow T} \phi^{Y_t}(t, U_t^D) = \phi^{Y_T}(T, U_{T-}^D) = \psi(U_{T-}^D) \tag{4.9}$$

and

$$\int_{[t,T)} e^{-\int_t^s r} \sum_{e\in\mathcal{Y}} I_{(Y_s=e)}(dA_s^e + dJ_s^e) = 0, \tag{4.10}$$

where

$$\begin{aligned} dA_s^e &= \left(\frac{\partial}{\partial t}\phi^e(s,U_s^D) - r^e\phi^e(s,U_s^D) + \frac{\partial}{\partial u}\phi^e(s,U_s^D)(r^eU_s^D + c_s^e)\right) ds \\ &\quad + \sum_{f\in\mathcal{Y}^e}\left(\phi^f(s,U_{s-}^D) - \phi^e(s,U_{s-}^D)\right)\lambda^{ef}\, ds + \left(1 - \frac{\partial}{\partial u}\phi^e(s,U_{s-}^D)\right) dD_s, \\ dJ_s^e &= \phi^e(s,U_s^D) - \phi^e(s,U_{s-}^D) - \frac{\partial}{\partial u}\phi^e(s,U_{s-}^D)\Delta U_s^D \\ &= \phi^e(s,U_s^D) - \phi^e(s,U_s^D + \Delta D_s) + \frac{\partial}{\partial u}\phi^e(s,U_s^D + \Delta D_s)\Delta D_s, \end{aligned}$$

then $\Phi^D(t,e,u) = \phi(t,e,u) = \Phi(t,e,u)$*, i.e. D is optimal.*

Proof. Let ϕ be as stated in the theorem. Let $(t,e,u) \in E^0$, and let $D \in \mathcal{D}(t,e)$ be arbitrary. Define the process $X = (X_s)_{s\in[t,T]}$ by

$$X_s = e^{-\int_t^s r}\,\phi^{Y_s}(s,U_s^D),\ s \in [t,T].$$

Note that $X_t = \phi^e(t,u+D_t)$. By Itô's formula (applied on each interval $(t_{i-1},t_i]$),

$$X_{T-} - \phi^e(t,u) = \int_{[t,T)} e^{-\int_t^s r} \sum_{e\in\mathcal{Y}} \left[I_{(Y_s=e)}(dA_s^e + dJ_s^e - dD_s) + dM_s^e\right],$$

where

$$dM_s^e = \sum_{f\in\mathcal{Y}^e}\left(\phi^f(s,U_{s-}^D) - \phi^e(s,U_{s-}^D)\right)\left(dN_s^{ef} - I_{(Y_{s-}=e)}\lambda^{ef} ds\right).$$

Now, ϕ satisfies (4.6), so we have $dA_s^e \le 0$, and $dJ_s^e \le 0$ because ϕ^e is concave in u. Furthermore, since ϕ has continuous partial derivatives on E^0 and $\frac{\partial}{\partial u}\phi^e(t,u) \ge 1$ on E^0,

$$X_{T-} = e^{-\int_t^T r}\lim_{t\nearrow T}\phi^{Y_t}(t,U_t^D) \ge e^{-\int_t^T r}\psi(U_{T-}^D).$$

Thus, by taking expectations we obtain

$$\phi^e(t,u) \ge \mathrm{E}_{(t,e)}\left(\int_{[t,T)} e^{-\int_t^s r} dD_s + e^{-\int_t^T r}\psi(U_{T-}^D)\right) = \Phi^D(t,e,u). \tag{4.11}$$

Since $D \in \mathcal{D}$ was arbitrary we conclude that $\Phi \le \phi$.

Now, assume that $D \in \mathcal{D}$ satisfies (4.9) and (4.10). Then, by similar calculations the inequality (4.11) becomes an equality, i.e. $\phi^e(t,u) = \Phi^D(t,e,u)$ and thereby $\phi^e(t,u) = \Phi(t,e,u)$. □

Now, suppose that we have found a solution ϕ as stated in Theorem 3.4.6. An analysis of the expressions for dA_s^e and dJ_s^e then shows that the optimal strategy can be characterized (loosely) as follows: We decompose the state space into disjoint regions as

$$E = K^\circ \cup \partial K \cup (E \setminus K),$$

where $K = \{(t,e,u) \in E : \frac{\partial}{\partial u}\phi^e(t,u) \leq 1\}$. At any time the dividend allocation should be based on the state of the controlled process. We first note that one should always put $dD_s = 0$ if $(s, Y_s, U_{s-}^{(t,e,u,D)}) \in E \setminus K$. Furthermore, to ensure that $dJ_s^e \equiv 0$, lump sum dividends should not be credited unless $(s, Y_s, U_{s-}^{(t,e,u,D)}) \in K^\circ$, and then only to the extent that the resulting state $(s, Y_s, U_{s-}^{(t,e,u,D)} - \Delta D_s)$ is kept out of $E \setminus K$ (note that for any $(t,e,u) \in K$ we actually have $\frac{\partial}{\partial u}\phi^e(t,u) = 1$).

The condition (4.9) means that as we reach the terminal time T, the dividend reserve must not exceed the boundary. Thus, although we have not yet specified precisely the optimal behaviour in $K^\circ \cup \partial K$, we interpret K° as the "jump region", $E \setminus K$ as the "no action region", and ∂K as the "optimal boundary": If $(s, Y_s, U_{s-}^{(t,e,u,D)}) \in K^\circ$ we immediately make a jump to ∂K, if $(s, Y_s, U_{s-}^{(t,e,u,D)}) \in E \setminus K$ we put $dD_s = 0$, and if $(s, Y_s, U_{s-}^{(t,e,u,D)}) \in \partial K$ we choose dD_s in such a way that the process stays in ∂K (if possible).

For $(t,e) \in [0,T) \times \mathcal{Y}$ we can consider the (t,e)-sections $(K^\circ)_{(t,e)}$, $(\partial K)_{(t,e)}$ and $(E \setminus K)_{(t,e)}$ consisting, respectively, of those $u \in \mathbb{R}$ such that (t,e,u) is in K°, ∂K and $E \setminus K$. Since ϕ is concave in u, they are of the form $(\tilde{u}(t,e), \infty)$, $\{\tilde{u}(t,e)\}$ and $(-\infty, \tilde{u}(t,e))$, respectively, for some $\tilde{u}(t,e) \in \mathbb{R} \cup \{\infty\}$. Thus, for each $e \in \mathcal{Y}$, the e-section $(K^\circ)_e$ (if non-empty) "sits on top of" $(E \setminus K)_e$, with $(\partial K)_e = \{\tilde{u}(t,e), t \in [0,T]\}$ constituting the boundary.

F. A closed-form solution.
Let

$$\tilde{u} = \inf_{u \in \mathbb{R}} \{u \in \mathbb{R} : \psi'(u) \leq 1\}$$

with $\inf \emptyset = \infty$, and define $\tilde{\psi} : \mathbb{R} \to \mathbb{R}$ by

$$\tilde{\psi}(u) = \begin{cases} \psi(u), & \text{if } u \leq \tilde{u}, \\ \psi(\tilde{u}) + u - \tilde{u}, & \text{otherwise.} \end{cases}$$

It is easily seen that $\tilde{\psi}'(u) \geq 1$ for all $u \in \mathbb{R}$.

Proposition 3.4.7 *Let $\phi = (\phi^e)_{e \in \mathcal{Y}} : E \to \mathbb{R}$ be given by*

$$\phi^e(t,u) = \mathrm{E}_{(t,e)} \left(e^{-\int_t^T r} \tilde{\psi} \left(u e^{\int_t^T r} + \int_t^T e^{\int_\tau^T r} c_\tau \, d\tau \right) \right), \ (t,e,u) \in E^0, \tag{4.12}$$

and $\phi^e(T,u) = \psi(u)$, $(e,u) \in \mathcal{Y} \times \mathbb{R}$. Then ϕ is a generalized solution to (4.6)-(4.7), which is concave in u. Furthermore, for $t \in [0,T]$, let

$$\tilde{u}(t) = \tilde{u} e^{-\underline{r}(T-t)} - \int_t^T e^{-\underline{r}(\tau - t)} \, (\underline{r} - r^*) V_\tau^* \, d\tau, \tag{4.13}$$

and let $D \in \mathcal{D}_{(t,e)}$ *be such that for all* $s \in [t,T]$ *(with* $U^D \equiv U^{(t,e,u,D)}$*),*

$$dD_s = \left(r^{Y_s} - \underline{r}\right)\left(\tilde{u}(s) + V_s^*\right) I_{(U_s^D = \tilde{u}(s))}\, ds + \left(U_{s-}^D - \tilde{u}(s)\right) I_{(U_{s-}^D > \tilde{u}(s))}. \tag{4.14}$$

Then D *is optimal.*

Remark 3.4.8 The expression on the right-hand side of (4.12) does not depend on the underlying probability space in the 5-tuple $(\Omega, \mathcal{F}, \mathrm{P}, Y, D) \in \mathcal{D}(t,e)$, and ϕ is therefore well defined.

Proof. It is easily seen that ϕ is concave in u. Since $\tilde{\psi}'$ is bounded by Assumption 3.4.1 (iii) we have for $(t,e,u) \in E^0$ that

$$\frac{\partial}{\partial u}\phi^e(t,u) = \mathrm{E}_{(t,e)}\left(\tilde{\psi}'\left(ue^{\int_t^T r} + \int_t^T e^{\int_\tau^T r}\, c_\tau\, d\tau\right)\right).$$

Thus, each ϕ^e has continuous partial derivatives in u with $\frac{\partial}{\partial u}\phi^e(t,u) \geq 1$, $(t,u) \in (0,T) \times \mathbb{R}$.

To examine the smoothness properties of the ϕ^e wrt. t we adopt the ideas of Norberg (2005). By conditioning on whether Y leaves the initial state e in (t,T) and, in case it does, on the transition, we get the integral equations

$$\begin{aligned}\phi^e(t,u) \quad &= \quad e^{-(\lambda^{e\cdot}+r^e)(T-t)}\, \tilde{\psi}\left(\widetilde{U}_T^{(t,e,u,0)}\right) \\ &\quad + \int_t^T e^{-(\lambda^{e\cdot}+r^e)(\tau-t)} \sum_{f\in\mathcal{Y}^e} \lambda^{ef}\phi^f\left(\tau, \widetilde{U}_\tau^{(t,e,u,0)}\right) d\tau,\end{aligned} \tag{4.15}$$

where $\widetilde{U}_\tau^{(t,e,u,0)}$ is the (deterministic) dividend reserve at time $\tau \in (t,T]$ when the initial state is (t,e,u) and when no jumps occur and no dividends are credited in $[t,\tau]$, i.e.

$$\widetilde{U}_\tau^{(t,e,u,0)} = ue^{r^e(\tau-t)} + \int_t^\tau e^{r^e(\tau-s)} c_s^e\, ds.$$

Now, introduce the transformed state-wise functions $w^e : [0,T) \times \mathbb{R} \to \mathbb{R}$, $e \in \mathcal{Y}$, given by

$$w^e(t,u) = e^{-(\lambda^{e\cdot}+r^e)t}\phi^e\left(t, ue^{r^e t} + \int_0^t e^{r^e(t-\tau)}\, c_\tau^e\, d\tau\right) = e^{-(\lambda^{e\cdot}+r^e)t}\phi^e\left(t, \widetilde{U}_t^{(0,e,u,0)}\right),$$

$(t,e,u) \in E^0$. The integral equations for the ϕ^e can be expressed in terms of the w^e as

$$w^e(t,u) = e^{-(\lambda^{e\cdot}+r^e)T}\, \tilde{\psi}\left(\widetilde{U}_T^{(0,e,u,0)}\right) + \sum_{f\in\mathcal{Y}^e}\int_t^T e^{\kappa^{ef}\tau}\, \lambda^{ef} w^f(\tau, W_\tau^{ef}(u))d\tau, \tag{4.16}$$

where $\kappa^{ef} = \lambda^{f\cdot} + r^f - \lambda^{e\cdot} - r^e$, and

$$W_\tau^{ef}(u) = ue^{(r^e-r^f)\tau} - \int_0^\tau e^{-r^f s}\, c_s^f\, ds + e^{(r^e-r^f)\tau}\int_0^\tau e^{-r^e s}\, c_s^e\, ds,\ \tau \in (0,T).$$

The advantage of working with the transformed functions is that the integral equations (4.16) are much simpler than (4.15) because the variable t only appears as the lower limit of the integral on the right hand side of (4.16). It is thus easy to see that each w^e has continuous partial derivatives wrt. t given by

$$\frac{\partial}{\partial t} w^e(t,u) = -\sum_{f \in \mathcal{Y}^e} e^{\kappa^{ef} t} \lambda^{ef} w^f(t, W_t^{ef}(u)), \ (t,u) \in (0,T) \times \mathbb{R}.$$

Now, from the expression of ϕ^e in terms of w^e,

$$\phi^e(t,u) = e^{(\lambda^{e\cdot} + r^e)t} w^e \left(t, u e^{-r^e t} - \int_0^t e^{-r^e \tau} c_\tau^e \, d\tau \right),$$

it is seen that the partial derivative of ϕ^e wrt. t exists for all $(t,u) \in (0,T) \times \mathbb{R}$, and after some calculations one arrives at

$$\frac{\partial}{\partial t} \phi^e(t,u) = r^e \phi^e(t,u) - \frac{\partial}{\partial u} \phi^e(t,u)(r^e u + c_t^e) - \sum_{f \in \mathcal{Y}^e} \lambda^{ef} \left(\phi^f(t,u) - \phi^e(t,u) \right). \tag{4.17}$$

The partial derivative is discontinuous exactly when c^e is discontinuous, which is the case at most at a finite number of time points where fixed contractual lump sum payments are due.

We conclude that ϕ is a generalized solution to (4.6)-(4.7).

To show that D given by (4.14) is optimal we verify that (4.9) and (4.10) are fulfilled. Since $\tilde{\psi}'(u) \geq 1$, $u \in \mathbb{R}$, we see that

$$\frac{\partial}{\partial u} \phi^e(t,u) = \mathrm{E}_{(t,e)} \left(\tilde{\psi}' \left(u e^{\int_t^T r} + \int_t^T e^{\int_\tau^T r} c_\tau \, d\tau \right) \right) = 1$$

if and only if

$$u e^{\int_t^T r} + \int_t^T e^{\int_\tau^T r} c_\tau \, d\tau \geq \tilde{u}, \quad a.s.,$$

which holds only if $u \geq \tilde{u}(t)$. By inspection of D, and by (4.17), we see that (4.10) holds. Now, for $s \in (t,T]$ we have

$$d\tilde{u}(s) = (\underline{r}\tilde{u}(s) - (r^* - \underline{r}) V_s^*) \, ds.$$

By using (2.8) it is straightforward to verify that if the dividend reserve is at the boundary at some time, say $U_t^{(t,e,u,D)} = \tilde{u}(t)$, it will remain at the boundary during $(t,T]$, implying that (4.9) holds. We skip the details. □

The outlined optimal control process consists of an absolutely continuous part and a jump part. With the terminology from Paragraph E the optimal boundary is given by $\tilde{u}(\cdot)$ — it is in particular independent of $e \in \mathcal{Y}$. It is seen that $\tilde{u}(T) = \tilde{u}$, which explains the choice of notation. Furthermore we have $\tilde{u}(t) < \infty$, $\forall t \in [0,T]$, if and only if $\tilde{u} < \infty$.

From (4.13) we see that $\tilde{u}(t)$ can be interpreted as the present value of $\tilde{u}$ minus the future surplus contributions in the worst case scenario (where $r_s = \underline{r}$ for $t < s \le T$). Thus, dividends should be credited only if the total wealth $V^* + U^{(t,e,u,D)}$ is sufficient to ensure that the contractual payments can be met and that $\tilde{u}$ can be reached at time T in the worst case scenario. From (4.14) we then see that in this case all interest earned on the wealth in excess of what would be earned in the worst case scenario should be credited as dividends. This will keep the dividend reserve at the boundary for the remaining part of the policy term (as shown in the proof of Proposition 3.4.7).

A lump sum payment should be made if $U^D_{s-} > \tilde{u}(s)$, but this will actually never happen under the outlined optimal strategy D (unless $\tilde{u}(0) < 0$), which is therefore absolutely continuous. However, we admit that our choice of objective function actually allows for (infinitely) many optimal strategies if the boundary is reached (at time t, say) because the optimal value is retained if dividends are withheld and credited (with compounded interest) at a later point $s \in (t, T)$.

In most realistic situations the surplus contributions will be negative in the worst case scenario, and the optimal boundary will therefore be decreasing in t (except possibly in the very beginning of the policy period). A considerate amount of time will normally pass before the boundary is reached. In return, dividends can be credited quite aggressively towards the end of the policy term. This is contrary to what has traditionally been common practice, namely that dividends are credited throughout the entire policy terms.

Finally, as mentioned in the introduction, the objective (3.2) makes our problem look like a singular control problem, and the optimal strategy is indeed similar to what is usually obtained in singular control problems, since in both cases it is characterized by the different regions of the state space and the optimal control to perform in those regions. However, since our controlled process U is of finite variation, the optimal control process D need not have a singular component, as it is possible to absorb U at the boundary. In contrast, in problems where the controlled wealth process is a diffusion it is only possible to reflect it at the boundary, and the optimal dividend process typically has a non-zero singular component.

3.5 Discussions and conclusion

A. The objective function.
The control problem considered in this paper has been formulated on the basis of two objectives. Firstly, dividends should be credited concurrently, more or less, with the realization of surplus during the term of the policy. Secondly, the company should always make sure that all future obligations can be met, even in scenarios with very poor investment performance.

The second objective is quite obvious, whereas the first one may not be, depending on what type of insurance one has in mind and also on how the dividends are converted into bonus payments: For an insured with, say, a life annuity serving as his or her retirement benefits, a large terminal bonus lump sum upon termination of the contract (at the time of death) would be somewhat useless, as opposed to bonus payments handed out along with the retirement benefits. On the other hand, for a life insurance or an endowment insurance, where the benefits consist of a single lump sum payment upon termination of the contract (due to death of the insured or expiration of the policy), a terminal bonus payment would be quite appropriate.

However, the problem addressed in this paper is much more relevant for those types of insurance for which large reserves are built up, and they typically have a large element of retirement benefits of some kind, e.g. life annuities. Therefore, the first objective is indeed reasonable. There are other reasons as well, see the introduction, but we shall not discuss this any further.

One could work with a fixed or deterministic interest rate instead of r in the discounting factor. However, for several reasons we have chosen not to do so. Although there are numerous possible bonus schemes, the prevailing ones share the property that no bonuses are paid out prior to the benefit period, since there is no particular need for this. Hence, all dividends credited before the benefit period gain interest by the rate r until they are paid out. So, even from the insured's point of view, r is a relevant interest rate under such bonus schemes, at least until the beginning of the benefit period.

The case with a fixed interest rate $\beta > 0$, say, might be considered more relevant for specific insurance types, or if all dividends were paid out immediately as cash bonus upon allotment. A major drawback is that (in reasonable specifications of the interest rate model) it would lead to optimal dividend schemes, which, at least from an actuarial point of view, would be counter-intuitive: One would again obtain the variational inequalities (4.6)-(4.7) with the term $-\phi^e r^e$ replaced by $-\phi^e\beta$, thus yielding the same types of optimal strategies. However, it can then be shown that if $e, f \in \mathcal{Y}$ with $r^e > r^f$, then

$$\frac{\partial}{\partial u}\phi^e > \frac{\partial}{\partial u}\phi^f,$$

from which it is seen that the "optimal boundary" corresponding to r^e would be above the one corresponding to r^f. In particular, it might be optimal, at a given point (t, U_t), to put $dD_t = 0$ if $r_t = r^e$ and $dD_t > 0$ if $r_t = r^f$!

The majority of papers on optimal investment and consumption only allow for absolutely continuous consumption, and the utility of the consumption stream is then measured by the expected value of an integral of the (discounted) utility of the consumption rate for some utility function. By this approach it is implicitly assumed that this utility function is *time additive*.

In some cases the use of a strictly concave utility function on consumption rates makes it possible to provide an analytical solution to the problem. However, measuring dividend or consumption rates with a strictly concave utility function must be done with care. As pointed out above, a large part of the dividends credited to a life insurance policy are typically accumulated and not paid out immediately.

Therefore, the abovementioned assumption about time additivity is hard to justify, and we have thus preferred to work with linear utility of dividends, which yields a *local substitution* effect.

Another issue is that if the marginal utility of consumption is large when the rate is close to zero (for power utility laws it tends to infinity), then the resulting optimal consumption rate may be strictly positive even when the wealth is close to 0. That is certainly not what an actuary would consider optimal.

We also point out that our choice of objective function yields optimal controls that are easy to characterize and understand.

On the other hand, it is quite reasonable to let ψ be strictly concave. Firstly, U_{T-} is a lump sum payment or loss, so the abovementioned time additivity problem is not an issue here. Secondly, since our problem is most relevant for pension schemes such as life annuities as opposed to life insurances, the insured's utility of a lump sum (extra) bonus payment at time T will typically be small compared to his or her utility of (extra) dividends (if, of course, they are converted into bonus payments prior to T), and this interest is being considered by choosing a strictly concave ψ.

B. Principles governing the dividends.
As mentioned, it is a widely accepted actuarial principle that one should reestablish equivalence at time T when the economic-demographic development during the policy term is known. This can be expressed by the requirement $U_T = 0$, which in the setup in Norberg (1999, 2001) can be fulfilled, since it is assumed there that S_{T-} is always non-negative.

In contrast, the model considered in this paper can lead to a negative terminal dividend reserve, and thus a loss for the company, in particular if dividends have been credited too aggressively in the past. Also in practice, this risk is present.

Thus, one may have to allow for the possibility that $U_T > 0$ so that the company makes a positive expected profit (as is typically the case in practice anyway) to establish fairness, e.g. by paying out $\max(0, U_{T-} - K)$ for some suitable $K > 0$. In this paper we have more or less ignored this issue and simply taken ψ as a given function (which of course should depend on the terminal payout function) measuring the utility of the dividend reserve just before T. Alternatively, as proposed in Norberg (2001), profits could be covered by expenses.

As a different theoretical approach to this delicate issue one could employ the principle of no arbitrage, which is fundamental in mathematical finance, see e.g. Duffie (1996). A general framework for this approach (within life and pension insurance) is provided in Steffensen (2001), where the redistribution of surplus for a traditional life insurance policy is taken to be a part of the contract terms so that the policy as a whole becomes a contingent claim in a financial market. More precisely, a policy does not only lay down a first order payment stream at issue, but also a dividend stream, both with payments that are functions of time and the current (random) state of an underlying *index* S, which is an observable, uncontrollable Markov vector process of a certain form. In short, the idea is that some of the component processes of S are traded on a financial market Z, whereas others may be non-marketed processes, representing e.g. the life history of the policyholder. It is then assumed that this market is free of arbitrage, and the model therefore provides a way to specify contracts that are fair for the insurance company as well as the policyholders in the sense that they are consistent with market prices of different kinds of risk in the index S. The model also explicitly yields reserves at market value.

This is of course very appealing from a theoretical point of view. However, we have chosen another approach, which reflects perhaps a more pragmatic view on life insurance: Typically, the durations of life and pension insurance contracts exceed those of even the longest term bonds available in the markets (which, by the way, partly explains why we have considered the first order payment stream as non-hedgeable). For this reason alone there are no unique arbitrage free values of such contracts. Furthermore, for traditional life and pension insurance contracts the dividends are not explicitly linked to some underlying index, but rather determined throughout the policy term by the company. In addition, assumptions of idealized financial markets such as unlimited and frictionless trading in both long and short positions are not fully met in reality.

Nevertheless, we shall briefly outline how the control problem addressed in this paper could be formulated within the framework in Steffensen (2001). One would have to introduce an equivalent martingale measure Q, which would amount to specifying an equivalent set of transition intensities of Y (see e.g. Norberg (2003)). Then, all dividend strategies would have to fulfill the requirement

$$\mathrm{E}^{\mathrm{Q}}\left(e^{-\int_0^T r}\, U_T\right) = 0. \tag{5.1}$$

In particular, strictly positive values of the terminal *surplus* U_T would have to be allowed, in contrast to the situation in Norberg (1999), where (5.1) would imply that $U_T = 0$ always. Thus, one could try to find optimal dividend streams among those fulfilling (5.1). However, we shall not pursue this further.

C. Concluding remarks.

The overall aim of this paper has been to investigate the problem of fixing dividends in life and pension insurance, which is a problem with two obviously conflicting interests: The company wants to hand out as much as possible, but it does not want to lose money systematically. The main motivation for this problem has been the fact that over the past few years many large companies have faced an increasing risk of losing money in this respect. This situations has been triggered by decreasing interest rates, but one could argue that it is also caused by too aggressive dividend strategies. The obtained results indicate that the prevailing strategies, by which dividends are credited to the insured throughout the entire term, are suboptimal. This hardly comes as a surprise, as dividends handed out early on in the policy term certainly increase the risk of a terminal loss, while the benefit for the insured is relatively small.

3.6 Generalizations

A. Synopsis.
In this section we explore some possibilities for generalizations of the results of this chapter to other and more general interest rate models. We continue to work with Markovian models throughout this section in order to be able to apply results from the theory of dynamic programming. We leave all components of the problem (i.e., setup, assumptions, notation, etc.) as they stand, unless they are explicitly modified below. In particular, we stick to the objective (3.2).

In Paragraph B we consider a fairly general Brownian motion-driven (diffusion) interest rate model and show that the results concerning optimal dividend strategies obtained in the finite-state Markov chain case carry over with straightforward modifications.

In Paragraph C we discuss further generalizations of the interest rate model and argue that similar results hold for general (suitably well-behaved) Markovian interest rate models as long as the interest rate is bounded from below, which is crucial. In this paragraph we leave out rigorous proofs and technical details and keep the discussion fairly loose in order to ease the readability.

B. A diffusion interest rate model.
In this paragraph we leave our finite-state Markov chain model and turn to a quite different class of interest rate models by assuming that $r = (r_t)_{t\in[0,T]}$ is a nonnegative continuous-time Markov diffusion process with dynamics of the form

$$dr_t = \alpha(t, r_t)\, dt + \sigma(t, r_t)\, dW_t, \tag{6.1}$$

where $\alpha, \sigma : [0,T]\times\mathbb{R} \to \mathbb{R}$ are measurable functions and W is a standard Brownian motion. We need to impose a few technical conditions. First, to ensure existence and uniqueness of a solution to the SDE (6.1), we assume that, for some $K > 0$,

$$\begin{aligned} |\alpha(t,e) - \alpha(t,f)| + |\sigma(t,e) - \sigma(t,f)| &\leq K\,|e-f|\,, \\ \alpha^2(t,e) + \sigma^2(t,e) &\leq K^2(1+e^2), \end{aligned}$$

for every $t \in [0,T]$, $e, f \in [0,\infty)$. These requirements also ensure that r is square integrable, and thus that the process

$$\left(\int_0^t \sigma(s, r_s)\, dW_s\right)_{t\in[0,T]}$$

is a (square-integrable) martingale (see, e.g. Karatzas and Shreve (1991)). Obviously, this model includes various popular short rate models as special cases.

The interest rate process is now continuous, and its state space, denoted by I, is therefore an interval in $[0,\infty)$. As in the finite-state Markov chain case, the greatest lower bound for the interest rate process, given here by

$$\underline{r} = \sup\{e \geq 0 : r(t) \geq e,\ \forall t \in [0,T],\ \text{a.s.}\},$$

turns out to play an important role for the optimal dividend strategy. For simplicity we assume that $\underline{r}$ is the greatest lower bound for any initial condition, that is, conditionally, given $(r_t = e)$ for some $(t, e) \in [0, T) \times I$, r_s may get arbitrarily close to $\underline{r}$ with strictly positive probability for $s \in (t, T]$.

The state space for the triple process $(t, r_t, U_t)_{t \in [0,T]}$, which (with a slight abuse of notation) is still denoted by E, is now given by $E = [0, T] \times I \times \mathbb{R}$. Correspondingly, we maintain the notation (t, e, u) for a generic argument for the value function $\Phi : [0, T] \times I \times \mathbb{R} \to \mathbb{R}$. The definitions of E^0 and $\mathcal{D}(t, e)$ in Section 3.4 are modified in an obvious way.

Remark 3.6.1 The performance and value functions are well defined. To see this, note first that for any $(t, e, u) \in E^0$ and $D \in \mathcal{D}(t, e)$,

$$\begin{aligned} \mathrm{E}_{(t,e)} \left| e^{-\int_t^T r} U_{T-}^{(t,e,u,D)} \right| &= \mathrm{E}_{(t,e)} \left| e^{-\int_t^T r} u + \int_{(t,T)} e^{-\int_t^s r} (c_s\, ds - dD_s) \right| \\ &\leq |u| + \mathrm{E}_{(t,e)} \left(\int_{(t,T)} |r_s - r^*| |V_s^*|\, ds \right) + \mathrm{E}_{t,e} \left(D_T \right) \\ &< \infty. \end{aligned}$$

The assertion now follows from Assumption 3.4.1 (iii).

The HJB equation becomes (skipping the generic argument (t, e, u) for notational convenience)

$$0 = \max \left(\mathcal{L}\phi - \phi\, e, 1 - \frac{\partial}{\partial u} \phi \right), \tag{6.2}$$

where

$$\mathcal{L}\phi = \frac{\partial}{\partial t}\phi + \frac{\partial}{\partial u}\phi\, (eu + c_t) + \frac{\partial}{\partial e}\phi\, \alpha(t, e) + \frac{1}{2} \frac{\partial^2}{\partial e^2} \phi\, \sigma^2(t, e),$$

with the boundary condition

$$\phi(T, e, u) = \psi(u),\ (e, u) \in I \times \mathbb{R}. \tag{6.3}$$

As in Section 3.4 we define correspondingly *a generalized solution to (6.2)-(6.3)* to be a function $\phi : [0, T] \times I \times \mathbb{R} \to \mathbb{R}$ satisfying the following requirements:

(i) There exists a finite partition $0 = t_0 < t_1 < \ldots < t_n = T$ such that ϕ is of class $C^{1,2,1}$ and fulfills (6.2) on $(t_{i-1}, t_i) \times I^\circ \times \mathbb{R}$ for $i = 1, \ldots, n$.

(ii) The condition (6.3) is fulfilled, and ϕ is continuous on $[0, T] \times I \times \mathbb{R}$ except possibly on the subset $\{(T, e, u) : \psi'(u) < 1\}$.

A solution is necessarily unique.

The following verification theorem, along with its proof, corresponds to Theorem 3.4.6.

Theorem 3.6.2 *Let ϕ be a generalized solution to (6.2)-(6.3), which is concave in u. Assume that, for any $(t,e,u) \in E^0$ and $D \in \mathcal{D}(t,e)$, we have*

$$\mathrm{E}_{(t,e)}\left(\int_t^T \left(\frac{\partial}{\partial e}\phi(s,r_s,U_s^D)\,\sigma(s,r_s)\right)^2 ds\right) < \infty, \tag{6.4}$$

where $U^D \equiv U^{(t,e,u,D)}$ is the the corresponding dividend reserve process. Then $\phi \geq \Phi$. Furthermore, for any $(t,e,u) \in E^0$, if $D \in \mathcal{D}(t,e)$ and U^D satisfy (a.s.)

$$\lim_{t \nearrow T} \phi(t,r_t,U_t^D) = \phi(T,r_T,U_{T-}^D) = \psi(U_{T-}^D) \tag{6.5}$$

and

$$\int_{[t,T)} e^{-\int_t^s r}(dA_s + dJ_s) = 0, \tag{6.6}$$

where

$$\begin{aligned} dA_s &= \mathcal{L}\phi(s,r_s,U_s^D)\,ds - r_s\phi(s,r_s,U_s^D)\,ds + \left(1 - \frac{\partial}{\partial u}\phi(s,r_s,U_{s-}^D)\right) dD_s, \\ dJ_s &= \phi(s,r_s,U_s^D) - \phi(s,r_s,U_{s-}^D) - \frac{\partial}{\partial u}\phi(s,r_s,U_{s-}^D)\Delta U_s^D \\ &= \phi(s,r_s,U_s^D) - \phi(s,r_s,U_s^D + \Delta D_s) + \frac{\partial}{\partial u}\phi(s,r_s,U_s^D + \Delta D_s)\Delta D_s, \end{aligned}$$

then $\Phi^D(t,e,u) = \phi(t,e,u) = \Phi(t,e,u)$, i.e. D is optimal.

Proof. Let ϕ be as stated in the theorem. Let $(t,e,u) \in E^0$, and let $D \in \mathcal{D}(t,e)$ be arbitrary. Define the process $X = (X_s)_{s\in[t,T]}$ by

$$X_s = e^{-\int_t^s r}\,\phi(s,r_s,U_s^D),\ s \in [t,T].$$

Note that $X_t = \phi(t,e,u+D_t)$. By Itô's formula (applied on each interval $(t_{i-1},t_i]$),

$$X_{T-} - \phi(t,e,u) = \int_{[t,T)} e^{-\int_t^s r}(dA_s + dJ_s + dM_s - dD_s),$$

where

$$dM_s = \frac{\partial}{\partial e}\phi(s,r_s,U_s^D)\,\sigma(s,r_s)\,dW_s.$$

Now, ϕ satisfies (6.2), so we have $dA_s \leq 0$, and $dJ_s \leq 0$ because ϕ is concave in u. Moreover, $M = (M_s)_{s\in[t,T]}$ is a zero-mean martingale due to (6.4). Since ϕ has continuous partial derivatives on E^0 and $\frac{\partial}{\partial u}\phi(t,e,u) \geq 1$ on E^0,

$$X_{T-} = e^{-\int_t^T r}\lim_{t\nearrow T}\phi(t,r_t,U_t^D) \geq e^{-\int_t^T r}\psi(U_{T-}^D).$$

Thus, by taking expectations we obtain

$$\phi(t,e,u) \geq \mathrm{E}_{(t,e)}\left(\int_{[t,T)} e^{-\int_t^s r}dD_s + e^{-\int_t^T r}\psi(U_{T-}^D)\right) = \Phi^D(t,e,u). \tag{6.7}$$

Since $D \in \mathcal{D}$ was arbitrary we conclude that $\Phi \le \phi$.

Now, assume that $D \in \mathcal{D}$ satisfies (6.5) and (6.6). Then, by similar calculations the inequality (6.7) becomes an equality, i.e. $\phi(t,e,u) = \Phi^D(t,e,u)$ and thereby $\phi(t,e,u) = \Phi(t,e,u)$. □

As already stated, Theorem 3.6.2 corresponds to Theorem 3.4.6, and the (abstract) characterization of the optimal dividend strategy following the latter in Section 3.4 applies here as well, with straightforward modifications.

Similarly, the following result corresponds to Proposition 3.4.7. To avoid blurring the picture with technicalities we state the proposition in a somewhat imprecise fashion. As a consequence, the proof becomes a bit sketchy. We refer to Remark 3.6.5 below for some technical details on how rigorousness can be obtained.

Proposition 3.6.3 *Let $\phi : E \to \mathbb{R}$ be given by*

$$\phi(t,e,u) = \mathrm{E}_{(t,e)}\left(e^{-\int_t^T r}\,\tilde{\psi}\left(ue^{\int_t^T r} + \int_t^T e^{\int_\tau^T r}\,c_\tau\,d\tau\right)\right),\ (t,e,u) \in E^0, \tag{6.8}$$

and $\phi(T,e,u) = \psi(u)$, $(e,u) \in I \times \mathbb{R}$, where $\tilde{\psi}$ is defined as in Paragraph 3.4.F. Then, under certain regularity conditions, ϕ is a generalized solution to (6.2)-(6.3), which is concave in u.

Furthermore, assume that for any $(t,e,u) \in E^0$ and $D \in \mathcal{D}(t,e)$, condition (6.4) is satisfied. For $t \in [0,T]$, let

$$\tilde{u}(t) = \tilde{u}e^{-\underline{r}(T-t)} - \int_t^T e^{-\underline{r}(\tau-t)}\,(\underline{r} - r^*)V_\tau^*\,d\tau, \tag{6.9}$$

and let $D \in \mathcal{D}_{(t,e)}$ be such that for all $s \in [t,T]$ (with $U^D \equiv U^{(t,e,u,D)}$),

$$dD_s = (r_s - \underline{r})\,(\tilde{u}(s) + V_s^*)\,I_{(U_s^D = \tilde{u}(s))}\,ds + \left(U_{s-}^D - \tilde{u}(s)\right)I_{(U_{s-}^D > \tilde{u}(s))}. \tag{6.10}$$

Then D is optimal.

Remark 3.6.4 The expression on the right-hand side of (6.8) does not depend on the underlying probability space in the 5-tuple $(\Omega, \mathcal{F}, \mathrm{P}, Y, D) \in \mathcal{D}(t,e)$. Moreover, the expectation exists (cf. Remark 3.6.1), and ϕ is therefore well defined.

Proof. It is easily seen that ϕ is concave in u. Since $\tilde{\psi}'$ is bounded by Assumption 3.4.1 (iii) we have for $(t,e,u) \in (0,T) \times I^\circ \times \mathbb{R}$ that

$$\frac{\partial}{\partial u}\phi(t,e,u) = \mathrm{E}_{(t,e)}\left(\tilde{\psi}'\left(ue^{\int_t^T r} + \int_t^T e^{\int_\tau^T r}\,c_\tau\,d\tau\right)\right).$$

Thus, $\frac{\partial}{\partial u}\phi(t,e,u) \ge 1$, $(t,e,u) \in (0,T) \times I^\circ \times \mathbb{R}$.

The first assertion of the proposition will follow if we can show that

$$\mathcal{L}\phi - \phi\, e = 0 \tag{6.11}$$

on $(t_{i-1}, t_i) \times I^\circ \times \mathbb{R}$ for $i = 1, \ldots, n$. But (6.11) follows from the Feynman-Kac formula, which is valid under certain regularity conditions (see Remark 3.6.5 below), so we conclude that ϕ is a generalized solution to (6.2)-(6.3) under such conditions.

The proof that D given by (6.10) is optimal goes exactly as in the proof of Proposition 3.6.3. □

Remark 3.6.5 The proof rests on the validity of the formula (6.11), which, as stated, follows from the Feynman-Kac formula (see, e.g. Karatzas and Shreve (1991)). However, the Feynman-Kac formula is a stochastic representation formula akin to (6.8), which is valid for a *given* (sufficiently regular) solution to some PDE. In the context of the problem at hand, the Feynman-Kac result says that *if* some function $\tilde{\phi}$ were known to be a (sufficiently regular) solution to the PDE (6.11) with the boundary condition $\tilde{\phi}(T, e, u) = \tilde{\psi}(u)$, then (6.8) would hold (with $\tilde{\phi}$ in the place of ϕ, of course). One cannot deduce that the converse is true in general, i.e., we cannot be sure that ϕ, which is *defined* by (6.8), satisfies (6.11). Although there also exist various results in this direction in the literature, we have not been able to find one that covers the situation at hand. We shall therefore content ourselves with a proof that (6.11) is valid under a smoothness assumption on ϕ and a technical integrability condition ((6.12) below).

Thus, assume that there exists a finite partition $0 = t_0 < t_1 < \ldots < t_n = T$ such that ϕ is of class $C^{1,2,1}$ on $(t_{i-1}, t_i) \times I^\circ \times \mathbb{R}$ for $i = 1, \ldots, n$. This assumption can be checked in special cases of the general interest model (6.1). For any $(t, e, u) \in E^0$, the process $Y^{(t,e,u)} = (Y_s^{(t,e,u)})_{s \in [t,T]}$ defined by

$$Y_s^{(t,e,u)} = e^{-\int_t^s r} \phi(s, r_s, U_s), \; t \le s \le T,$$

where we have used the short-hand notation U_s for $U_s^{(t,e,u,0)}$, is a martingale, because

$$Y_s^{(t,e,u)} = \mathrm{E}_{(t,e)} \left(e^{-\int_t^T r} \tilde{\psi}(U_T) \,\middle|\, \mathcal{G}_s \right), \; t \le s \le T.$$

By Itô's formula we have

$$dY_s^{(t,e,u)} = e^{-\int_t^s r} \left(\mathcal{L}\phi(s, r_s, U_s) + \frac{\partial}{\partial e}\phi(s, r_s, U_s)\sigma(s, r_s)\, dW_s - \phi(s, r_s, U_s)\, r_s \right).$$

We now assume that, for any $(t, e, u) \in E^0$ and any stopping time τ taking values in $[t, T]$, we have

$$\mathrm{E}_{(t,e)} \left(\int_t^\tau \frac{\partial}{\partial e}\phi(s, r_s, U_s)\sigma(s, r_s)\, dW_s \right) = 0, \; \forall (t, e, u) \in E^0. \tag{6.12}$$

This condition is satisfied if

$$\mathrm{E}_{(t,e)} \left(\int_t^T \left(\frac{\partial}{\partial e}\phi(s, r_s, U_s)\sigma(s, r_s) \right)^2 ds \right) < \infty, \; \forall (t, e, u) \in E^0.$$

For any $(t, e, u) \in E^0$ and any stopping time τ as above we then have, by the optional sampling theorem,

$$\mathrm{E}_{(t,e)}\left(\int_t^\tau e^{-\int_t^s r}\left(\mathcal{L}\phi(s, r_s, U_s) - \phi(s, r_s, U_s)\, r_s\right) ds\right) = 0. \tag{6.13}$$

Now, let $(t, e, u) \in E^0$ be a point such that $t \in (t_{i-1}, t_i)$ for some $i = 1, \ldots, n$, and assume that

$$\mathcal{L}\phi(t, e, u) - \phi(t, e, u)\, e > 0.$$

Put

$$\tau = \inf\left\{s \in [t, T] : \mathcal{L}\phi(s, r_s, U_s) - \phi(s, r_s, U_s)\, r_s \leq 0\right\} \wedge T.$$

Since $\phi \in C^{1,2,1}((t_{i-1}, t_i) \times I^\circ \times \mathbb{R}$, we have $\tau > t$, a.s., so (6.13) cannot be true, and we have a contradiction. A similar argument shows that we cannot have

$$\mathcal{L}\phi(t, e, u) - \phi(t, e, u)\, e < 0,$$

so we conclude that (6.11) holds.

Once again, the (precise) characterization of the optimal dividend strategy following Proposition 3.4.7 applies here as well, with straightforward modifications.

C. Further generalizations of the interest rate model.
We have seen in Paragraph B that the results concerning the optimal dividend strategy in the finite-state Markov chain case carry over with obvious modifications to the present case. In this paragraph we discuss further generalizations.

We begin by arguing that it is very reasonable to believe that our results hold in general for any nonnegative Markov interest rate model as long as it is reasonably well behaved. To see this, assume that r is modelled simply as some continuous-time nonnegative Markov process, for which $\underline{r}$ is the greatest lower bound for any initial condition, as was assumed above in the diffusion case. Assume that

$$\mathrm{E}_{(t,e)}\left(\int_t^T r(t)\, dt\right) < \infty,\ \forall (t, e) \in [0, T] \times I,$$

where I again denotes the state space of the interest rate. Now, the explicit results obtained in the two cases we have studied are based on the simple observation that, due to the form of the objective (3.2), one can obtain an almost-optimal dividend strategy by holding back all dividends in the dividend reserve until t is very close to T and then simply hand out $(U_t - \tilde{u})^+$ once and for all. More precisely, if $(t_n)_{n\geq 1}$ is an increasing sequence of time points in $[0, T)$ with $\lim_{n\to\infty} t_n = T$, and $(D^{(n)})_{n\geq 1}$ is a sequence of dividend strategies given by

$$D^{(n)}(s) = \begin{cases} 0, & s < t_n, \\ (U_{t_n} - \tilde{u})^+, & s \geq t_n, \end{cases}$$

then it is intuitively very reasonable that, for any $(t, e, u) \in E^0$,

$$\lim_{n\to\infty} \Phi^{D^{(n)}}(t, e, u) = \Phi(t, e, u). \tag{6.14}$$

There is no reason why this should not be the case in general. Moreover, by definition we have

$$\Phi^{D^{(n)}}(t, e, u) = \mathrm{E}_{(t,e)}\left(\int_{[t,T)} e^{-\int_t^s r}\, dD_s^{(n)} + e^{-\int_t^T r}\, \psi\left(U_{T-}^{(t,e,u,D^{(n)})}\right)\right),$$

and it is fairly easy to verify, from the form of the $D^{(n)}$, that for $n \to \infty$ we almost surely have

$$\int_{[t,T)} e^{-\int_t^s r}\, dD_s^{(n)} + e^{-\int_t^T r}\, \psi\left(U_{T-}^{(t,e,u,D^{(n)})}\right) \to e^{-\int_t^T r}\, \tilde{\psi}\left(U_T^{(t,e,u,0)}\right).$$

This convergence is in fact dominated, so we have

$$\begin{aligned}\lim_{n\to\infty} \Phi^{D^{(n)}}(t, e, u) &= \mathrm{E}_{(t,e)}\left(e^{-\int_t^T r}\, \tilde{\psi}\left(U_T^{(t,e,u,0)}\right)\right) \\ &= \mathrm{E}_{(t,e)}\left(e^{-\int_t^T r}\, \tilde{\psi}\left(ue^{\int_t^T r} + \int_t^T e^{\int_\tau^T r}\, c_\tau\, d\tau\right)\right).\end{aligned}$$

Thus, we are led to conclude (or at least believe) that

$$\Phi(t, e, u) = \mathrm{E}_{(t,e)}\left(e^{-\int_t^T r}\, \tilde{\psi}\left(ue^{\int_t^T r} + \int_t^T e^{\int_\tau^T r}\, c_\tau\, d\tau\right)\right), \ (t, e, u) \in E^0. \tag{6.15}$$

Indeed, this identity is exactly what has been verified in the finite-state and diffusion cases (recall (4.12) and (6.8)). It is therefore natural to believe that our characterization of the optimal dividend strategy is valid in general.

It is crucial for our results that the interest rate is bounded from below (we have even assumed that it is nonnegative, but this is not important). As we have seen, the lower bound $\underline{r}$ actually plays an important role for the optimal dividend strategy. Economically, this condition is not only meaningful, it is even desirable. Mathematically, however, it may be interesting to drop this condition and consider the case where the interest rate is unbounded and may assume any negative value, as is the case for Gaussian models such as the well-known Vasicek (1977) model. The calculations above are still valid (at least if all technicalities are in place), so it is still reasonable to believe (6.15) to hold. However, the situation is entirely different: There is no optimal dividend strategy. To see this, recall from the discussion following Theorem 3.4.6 that dividends should never be credited if the current state of the controlled process is not in the set $K = \{(t, e, u) \in E : \frac{\partial}{\partial u}\phi(t, e, u) \leq 1\}$. However, with an unbounded interest rate we have $K = \emptyset$ (as can be seen by arguments similar to the ones used in the last part of the proof of Theorem 3.4.7), so dividends should never be credited!

The intuitive explanation as to why there is no optimal strategy is, of course, that with an interest rate that is unbounded from below there is always a chance (however small it may be) that a severe loss is suffered, and if dividends have been handed out before T, then a loss may be punished harder than if this were not the case. On the other hand, since (6.14) still holds, almost-optimal strategies can be obtained from the sequence $(D^{(n)})_{n\geq 1}$.

Chapter 4

Utility Maximization and Risk Minimization in Life and Pension Insurance

Sections 4.1-4.7 of this chapter constitute an adapted and modified version of Nielsen (2004). An abridged and modified version of Nielsen (2004) without Section 4.6 appears as Nielsen (2006). Sections 4.8-4.9 have been added for this thesis and thus do not appear in Nielsen (2004) or Nielsen (2006).

We consider a life insurance company that seeks to optimize the pension benefits on behalf of an insured. We take the uncertain course of life of the insured explicitly into account and thus have a non-standard financial optimization problem for which we propose a two-step approach: First, according to a certain preference structure and under a certain fairness constraint, an optimal pension payment process is obtained. This leaves the company with a non-hedgeable liability, for which we then discuss two quadratic hedging approaches. We obtain general results on dividend optimization, indicating that some widely used strategies are suboptimal, and semi-explicit expressions for the optimal bonus and investment strategies.

4.1 Introduction

The theory of optimal dynamic investment and consumption strategies in continuous time deals with the problem of an agent equipped with an initial endowment (and/or an income stream), who can invest in a financial market modelled by some multi-dimensional stochastic process, and who wants to maximize the expected utility, according to some preference structure, of consumption and/or terminal wealth. Beginning with the classical papers Merton (1969, 1971), where Markovian dynamics were assumed and dynamic programming methods were employed, the theory has developed extensively over the past few decades, and virtually countless variations and generalizations of Merton's problems have been analyzed and studied. A more modern approach, which allows of more general (in particular non-Markovian) structures, is the martingale method developed by Pliska (1986), Karatzas et al. (1987), and Cox and Huang (1989, 1991). Initially this method rested on the assumption of a complete market, but it has been generalized by the use of duality theory to incomplete markets by Karatzas et al. (1991) and further by Kramkov and Schachermayer (1999) and Cvitanić et al. (2001). We shall not give an account of the literature here; the reader is referred to, e.g. Karatzas and Shreve (1998), Sect. 3.11, or Duffie (1996), Ch. 9, for extensive surveys.

Merton's theory and its generalizations are clearly relevant for life and pension insurance companies acting on behalf of their policyholders who, apart from buying insurance coverage, are saving for retirement and thus may be thought of as economic agents. In the major part of the literature the agent is (implicitly) assumed to outlive the time horizon in consideration, and the agent's wealth and income are unambiguously defined and depend on the evolution of the financial market only. This assumption is certainly justifiable in many situations, but the payments — and thus the wealth and income — pertaining to a life insurance policy depend in general not only on the evolution of the financial market but also on the (uncertain) course of life of the insured. Ignoring this uncertainty corresponds to the implicit assumption that the non-financial part of the risk associated with a policy can be "diversified away" so that one can work with the financial part only. However, it is not clear in general exactly how the "financial part" is to be identified, and, furthermore, the implicit diversification assumption need not be well founded.

In this chapter we consider an optimization problem where the individual policy risk is taken explicitly into account. More specifically, we consider a life insurance company or a pension fund (henceforth referred to as the *company*) aiming at maximizing the expected utility of the pension benefits for a policyholder through dividend allocation and investment. We work with a general multi-state policy, for which the case with an uncertain life time is just one of several possible specializations. The company can invest in a complete financial market driven by a d-dimensional Brownian motion and with random coefficients. We do not assume that the market provides hedging opportunities regarding the policy risk, however, so the financial market is only complete when looked at in isolation. In other words, the policy payment process is not assumed to be put on the market.

For the company it is therefore not just a matter of investing for the insured as if he or she were an economic agent. Rather, the genuine risk associated with the policy must be taken into account, and it is not clear in general how this risk should be handled and how the company should invest. The company faces an overall problem involving two concerns: optimization on behalf of the insured and optimization for the company itself, according to suitable criteria. We propose a two-step solution approach that, to our knowledge, has not been proposed in the existing literature. The first step consists in utility maximization of the benefit stream for the insured (as explained in more detail below), taking his or her objective as given. The second step consists in treating the optimal benefit stream obtained in the first step as a given (random) liability for the company and determine an investment strategy that takes this liability into account and is optimal for the company in a suitable sense. We consider risk minimization as the chosen criterion.

Although our life insurance policy model is quite general and allows for various specifications (in particular purely financial ones), it is set-up as a general so-called *participating* or *with-profit* policy, characteristic of which are its contractual payment streams with built-in safety margins. These lead to the emergence of a systematic surplus, which forms the source of bonus benefits paid to the policyholder in addition to the contractual benefits, see, e.g. Norberg (1999). This chapter deals with optimization of strategies for bonus distribution and investment within the class of strategies satisfying a certain fairness constraint based on a technical assessment of a fair price of the policy risk.

A widely used method in practice is to allocate parts of the emerged surplus to the policyholder during the policy term through *dividends*, which are then converted into future bonus benefits. Focusing solely on the premium payment period (sometimes referred to as the *accumulation phase*) of the policy term we obtain some very general results on optimal dividend allocation. Apart from being of interest in their own right they allow of a simplification of the optimization problem, which is then approached using the martingale methodology in the sense that we consider the problem as a static optimization problem and use methods from convex analysis to obtain a bonus strategy that is optimal in terms of our objective. This analysis constitutes the first of the abovementioned two steps.

As mentioned, perfect hedging of the optimal payment streams is impossible due to the policy risk. Consequently, the company is left with a genuine risk that we propose to address by quadratic hedging methods known from the theory of risk minimization in incomplete markets. This consitutes the second step.

A related problem in a Markovian model is studied in Steffensen (2004), where dynamic programming techniques are applied. He works with the surplus process only (using a different notion of surplus than the traditional "actuarial" one, see Norberg (1999)) and focuses on dividend and investment strategies optimizing the bonus benefits, allowing for path-dependent utility in a certain form but not for dividend schemes that impose a genuine risk for the company. We allow for path-dependent utility in a general form as well as dividend schemes imposing a genuine risk. Furthermore, our approach is not based on a specific surplus or wealth process and thus appears to be less restrictive. We take into account the total payments and consider an objective in a general form that can be specialized, e.g. to involving the total pension benefits or the bonus benefits only.

Although our main focus is on optimization we also provide a generalized and detailed analysis of the total financial impact that a life insurance policy has on the company, without ignoring the policy risk.

The chapter is organized as follows. In Section 4.2 we present the model of the financial market and the insurance policy. We formulate the optimization problem in Section 4.3 and determine optimal dividend strategies in Section 4.4. We discuss two hedging approaches in Section 4.5. Section 4.6 is dedicated to an example covering the combined results to a large extent, and Section 4.7 concludes. Finally, Section 4.8 contains some discussions, and Section 4.9 compares our approach with the one taken in Steffensen (2004).

4.2 The model

A. Prefatory remarks on notation etc.
We consider throughout a time interval $[0, T']$ for some $T' \in (0, \infty)$. We shall have occasion also to include time $0-$, interpreted as the time point immediately before 0; this should cause no confusion as the expressions involved will have an obvious interpretation and therefore need no further explanation.

All random variables are defined on a complete and filtered probability space $(\Omega, \mathcal{F}, \mathbb{F} = (\mathcal{F}_t)_{t\in[0,T']}, \mathrm{P})$. We denote by $\mathcal{V}$ (resp. $\mathcal{A}$) the set of càdlàg, $\mathbb{F}$-adapted processes A with finite (resp. integrable) variation over $[0, T']$ and $A_{0-} \equiv 0$ (but, possibly, $A_0 \neq 0$). Clearly, $\mathcal{A} \subseteq \mathcal{V}$. For $A \in \mathcal{V}$ and an A-integrable process H, $\int_{[0,t]} H_s\, dA_s$ is interpreted as $H_0 \Delta A_0 + \int_{(0,t]} H_s\, dA_s = H_0 A_0 + \int_{(0,t]} H_s\, dA_s$.

The transpose of a vector or matrix A is denoted by A'. For a matrix A, $A^{i\cdot}$ denotes the i'th row, interpreted as a row vector. Relations are stated in weak form, i.e., "increasing" means "non-decreasing" etc.

B. The financial market.
We adopt a standard financial market model where all uncertainty is generated by a d-dimensional standard Brownian motion $W = (W^1, \ldots, W^d)'$, and where trading occurs without friction, i.e., there is unlimited supply of the available assets, no short-sale constraints, no trading costs etc. We shall make use of well-known fundamental results from mathematical finance without explicit references; the reader is referred to Karatzas and Shreve (1998), Ch. 1, for details and proofs.

The filtration generated by W and augmented by the P-null sets of $\mathcal{F}$ is denoted by $\mathbb{F}^W = (\mathcal{F}_t^W)_{t\in[0,T']}$. The market consists of a locally risk-free asset with price dynamics given by

$$\frac{dS_t^0}{S_t^0} = r_t\, dt,\ S_0^0 = 1,$$

i.e., $S_t^0 = e^{\int_0^t r_s\, ds}$, $0 \le t \le T'$, and d risky assets with price processes $S^1, \ldots, S^d$ given by the dynamics

$$\frac{dS_t^i}{S_t^i} = (r_t + \sigma_t^{i\cdot} \lambda_t)\, dt + \sigma_t^{i\cdot}\, dW_t,\ S_0^i > 0,\ i = 1, \ldots, d. \tag{2.1}$$

The (short) interest rate process r, the "market price of risk" vector process $\lambda = (\lambda^1, \ldots, \lambda^d)'$, and the volatility matrix process $\sigma = (\sigma^{ij})_{1\le i,j\le d}$ are the coefficients of the model. We assume that they are progressively measurable with respect to $\mathbb{F}^W$ and that r is integrable and all components of λ and σ are square integrable with respect to the Lebesgue measure, almost surely.

We assume that σ is non-singular for Lebesgue-a.e. $t \in [0, T']$, a.s., and that the process $\Lambda = (\Lambda_t)_{t\in[0,T']}$ given by

$$\Lambda_t = \exp\left(-\int_0^t \lambda_s' \, dW_s - \frac{1}{2}\int_0^t \|\lambda_s\|^2 \, ds\right), \; 0 \le t \le T',$$

is a square-integrable martingale.

Remark 4.2.1 Under suitable assumptions on admissibility of portfolio processes and contingent claims this market is arbitrage free and complete. In the following, such assumptions will be imposed when appropriate.

C. The insurance policy and its contractual payments.
We consider a life insurance policy issued at time 0 and terminating at time T'. We shall make use of well-known fundamentals and common terminology of participating life and pension insurance without explaining all details. Although our approach is somewhat different, we refer readers, to whom this is unfamiliar territory, to e.g. Norberg (1999, 2001) or Steffensen (2001, 2004), where further details and discussions on issues such as reserves, surplus, dividends and bonus are provided.

The (uncertain) course of life of the insured under the policy is represented by a stochastic *policy state process* $Z = (Z_t)_{t\in[0,T']}$ with a finite state space $\mathcal{Z} = \{0, \ldots, k\}$. We assume that Z and W are independent and for simplicity that $Z_0 \equiv 0$. The filtration generated by Z and augmented by the P-null sets of $\mathcal{F}$ is denoted by $\mathbb{F}^Z = (\mathcal{F}_t^Z)_{t\in[0,T']}$. We further assume that $Z \in \mathcal{A}$, so that it has a compensator.

Upon issue of the policy, the company and the insured agree upon a stream of premiums and a stream of *contractual* or *guaranteed* benefits. By counting premiums as negative and benefits as positive these payment streams can be merged into a single stochastic payment process $\widehat{B}$ counting the accumulated payments from the company to the insured (see e.g. Norberg (1999)). We assume for simplicity that $\widehat{B}$ is driven by Z (i.e., $\mathbb{F}^Z$-adapted) and further that $\widehat{B} \in \mathcal{V}$.

Apart from the stated technical assumptions on Z and $\widehat{B}$ we impose no particular restrictions on the distribution of Z or on the form of $\widehat{B}$, and our model is therefore quite general. To avoid confusion we stress that $\widehat{B}$ is a fixed $\mathbb{F}^Z$-adapted process that will not be subject to optimization in this chapter.

D. Valuation.
At any time $t \in [0, T']$ the future contractual payments constitute a liability on the part of the company. Valuation of this liability is taken to be market based and carried out on the basis of a probability measure Q, defined on $(\Omega, \mathcal{F}_{T'})$, and the available information at time t. Mathematically, this information is given by

$\mathcal{F}_t$; $\mathbb{F}$ is taken to be the filtration generated by W and Z and augmented by the P-null sets of $\mathcal{F}$, in particular $(\Omega, \mathcal{F}, \mathbb{F}, \mathrm{P})$ then satisfies the "usual conditions" of right-continuity and completeness. We assume that $\mathrm{Q} \sim \mathrm{P}$ and, more specifically, that

$$\frac{d\mathrm{Q}}{d\mathrm{P}} = \Lambda_{T'}\Gamma_{T'},$$

where $\Gamma = (\Gamma_t)_{t\in[0,T']}$ is a strictly positive square-integrable $\mathbb{F}^Z$-adapted martingale.

This means that purely financial quantities are evaluated at market value, and it is well-known (Girsanov's Theorem) that $W^{\mathrm{Q}} = (W^{\mathrm{Q},1}, \ldots, W^{\mathrm{Q},d})'$ given by

$$W_t^{\mathrm{Q}} = W_t + \int_0^t \lambda_s \, ds, \; t \in [0, T'],$$

is a standard d-dimensional Wiener process under Q. Note also that $d\mathrm{Q}/d\mathrm{P} \in L^2(\Omega, \mathcal{F}_{T'}, \mathrm{P})$, and that the independence between Z and W is preserved under Q.

We assume throughout that

$$\mathrm{E}\left(\int_{[0,T']} \frac{1}{S_s^0} \left| d\widehat{B}_s \right| \right)^2 < \infty, \tag{2.2}$$

and define the *market reserve* at time $t \in [0, T']$ corresponding to the contractual payment process $\widehat{B}$ as

$$\widehat{V}_t = S_t^0 \, \mathrm{E}^{\mathrm{Q}} \left(\int_{(t,T']} \frac{1}{S_s^0} \, d\widehat{B}_s \,\middle|\, \mathcal{F}_t \right). \tag{2.3}$$

Note that the reserve is in respect of strictly future payments; this implies that $\widehat{V}$ is also right-continuous.

Remark 4.2.2 We have made no particular assumptions about the (marginal) distribution of Z under Q, denoted by Q_Z, which calls for a comment: Since $\widehat{V}_t$ has been labelled as the market reserve, Q_Z should, ideally, be determined on the basis of market prices. This may not be fully possible due to the lack of a (sufficiently) liquid market, and Q_Z could instead be an estimated distribution based on historical data, possibly adjusted to allow for a risk loading. The actual form of Q_Z is not important for our purposes. However, for clarity it may be nice to think of the standard model in multi-state life insurance mathematics (see e.g. Hoem (1969) or Norberg (1991)), where Z is a continuous-time Markov chain. In simple special cases one can obtain closed-form expressions for the *state-wise (market) reserves* and thus for $\widehat{V}_t$.

Remark 4.2.3 The condition (2.2) is sufficient to ensure that the reserve is well defined. Furthermore, square integrability is necessary for the methods employed in Section 4.5 and thus for the overall approach in this chapter.

In participating life insurance it is usually required that the contractual payments (as given by $\widehat{B}$) be determined in accordance with the actuarial *equivalence principle* applied under prudent conditions (on interest, mortality etc.), which lead to comfortable safety loadings on the premiums, see, e.g. Norberg (1999). However, we only impose the (weaker) condition

$$\widehat{V}_{0-} \leq 0, \tag{2.4}$$

that is, the initial market reserve must be negative. In Section 4.4 we shall require that the inequality (2.4) be strict, but this is not needed in the general setup.

E. Bonus benefits and fair contracts.
In addition to the guaranteed benefits the insured is entitled to *bonus benefits* because the policy would otherwise be unfair (due to (2.4)). However, these benefits are typically not stipulated in the policy; they are in fact determined by the company concurrently during $[0, T']$ and depend in general on the development of the financial market. This chapter deals with optimization of the *bonus strategy.*

Certain principles must be adhered to. The bonus payments must be non-anticipative, i.e., they must always be paid out on the basis of the information at hand, that is, the observed history of the involved processes. Mathematically, this means that the payment process $\widetilde{B}$ representing the accumulated bonus benefits must be $\mathbb{F}$-adapted. Naturally, we must have $\widetilde{B}_{0-} \equiv 0$. Furthermore, bonus payments cannot be reclaimed once they have been paid out, so $\widetilde{B}$ must be increasing. We also assume that $\widetilde{B} \in \mathcal{V}$ (i.e., $\widetilde{B}_{T'} < \infty$, a.s.) and that

$$\mathrm{E}\left(\int_{[0,T']} \frac{1}{S_s^0}\, d\widetilde{B}_s\right)^2 < \infty. \tag{2.5}$$

Remark 4.2.4 Remark 4.2.3 (with suitable modifications) obviously also applies to condition (2.5).

We require the policy, viewed in its entirety, to form a *fair game* in the sense that it must not create a risk-free profit (or loss) for the company. To this end we impose the condition that

$$\mathrm{E}^{\mathrm{Q}}\left(\int_{[0,T']} \frac{1}{S_t^0} \left(d\widehat{B}_t + d\widetilde{B}_t\right)\right) = 0. \tag{2.6}$$

Since the contractual payments are specified in the policy and therefore cannot be altered, this can also be interpreted as a condition on the bonus benefits, as is seen explicitly from the equivalent condition

$$\mathrm{E}^{\mathrm{Q}}\left(\int_{[0,T']} \frac{1}{S_t^0}\, d\widetilde{B}_t\right) = -\widehat{V}_{0-}, \tag{2.7}$$

obtained from (2.6) by use of (2.3).

The fairness condition says nothing about the form of $\widetilde{B}$ and in that sense it is therefore quite weak; it should be taken only as necessary in order to avoid unfair insurance schemes. It is certainly not sufficient to ensure that $\widetilde{B}$ is reasonable; one can easily construct rather obscure bonus schemes that satisfy (2.6) but are neither reasonable nor fair in the usual sense of the word. On the other hand it is perhaps a strict condition from a practical point of view, as it represents a somewhat idealized insurance market. For a further discussion on the principle underlying (2.6) and a comparison with a more traditional "actuarial" principle we refer to Steffensen (2001).

F. The total market reserve.
Being $\mathbb{F}$-adapted, any *given* bonus process $\widetilde{B}$ can be treated as if it were a contractual obligation. Therefore we can define the *total market reserve* at time $t \in [0, T']$,

$$V_t = S_t^0 \,\mathrm{E}^{\mathrm{Q}} \left(\int_{(t,T']} \frac{1}{S_s^0} \left(d\widehat{B}_s + d\widetilde{B}_s \right) \,\middle|\, \mathcal{F}_t \right). \tag{2.8}$$

By (2.6) we must have $V_{0-} = 0$.

G. Examples.
We provide a few examples that may be nice to keep in mind.

Example 4.2.5 Suppose that under Q, Z is a continuous-time Markov chain that admits deterministic transition intensity functions (measurable and positive, of course) denoted by $\mu^{ef} = (\mu_t^{ef})_{t\in[0,T']}$, $e \neq f \in \mathcal{Z}$. Suppose further that $\widehat{B}$ has the form

$$d\widehat{B}_t = \sum_{e\in\mathcal{Z}} 1_{(Z_{t-}=e)}\, d\widehat{B}_t^e + \sum_{e,f\in\mathcal{Z}} \widehat{b}_t^{ef}\, dN_t^{ef},$$

where, for each $e, f \in \mathcal{Z}$, $\widehat{B}^e$ and $\widehat{b}^{ef}$ are deterministic functions (and $\widehat{B}^e \in \mathcal{V}$), and N^{ef} is the counting process counting the number of jumps from e to f, $e \neq f \in \mathcal{Z}$. This is a standard model in multi-state life insurance mathematics, see e.g. Hoem (1969) or Norberg (1991).

Then the market reserve at time $t \in [0, T']$, $\widehat{V}_t$, is given by

$$\widehat{V}_t = \sum_{e\in\mathcal{Z}} 1_{(Z_t=e)}\, \widehat{V}_t^e,$$

where the $\widehat{V}_t^e$ are the state-wise reserves, defined for each $e \in \mathcal{Z}$ by

$$\widehat{V}_t^e = S_t^0 \int_{(t,T']} \mathrm{E}^{\mathrm{Q}} \left(\frac{1}{S_s^0} \,\middle|\, \mathcal{F}_t \right) \sum_{f\in\mathcal{Z}} \mathrm{Q}\left(Z_s = f \,|\, Z_t = e \right) \left(d\widehat{B}_s^f + \sum_{g\in\mathcal{Z},\, g\neq f} \mu_s^{fg} \widehat{b}_s^{fg}\, ds \right).$$

Note that the state-wise reserves are $(\mathcal{F}_t^W)_{t\in[0,T']}$-adapted processes.

Example 4.2.6 (Example 4.2.5 continued). Let $\mathcal{Z} = \{0, 1\}$, where 0 and 1 are interpreted as the states "alive" and "dead", respectively, so that $\mu^{10} \equiv 0$. Assume further (for simplicity) that $\widehat{B}^1 \equiv 0$, i.e., there are no payments after the time of death. Then $\widehat{V}^1 \equiv 0$, and

$$\widehat{V}_t^0 = S_t^0 \int_{(t,T']} \mathrm{E}^{\mathrm{Q}} \left(\frac{1}{S_s^0} \middle| \mathcal{F}_t \right) e^{-\int_t^s \mu_u^{01}\, du} \left(d\widehat{B}_s^0 + \mu_s^{01} \widehat{b}_s^{01}\, ds \right),\ t \in [0, T'].$$

Here we put $e^{-\int_t^s \mu_u^{01}\, du} = 0$ if $\int_t^s \mu_u^{01}\, du = \infty$.

Example 4.2.7 (Example 4.2.6 continued). Let $\mu^{01} \equiv 0$, i.e., no uncertainty regarding Z is taken into account. This corresponds to the situation usually considered in the literature on investment and consumption. Then $\widehat{V}^1$ is irrelevant, and

$$\widehat{V}_t^0 = S_t^0 \int_{(t,T']} \mathrm{E}^{\mathrm{Q}} \left(\frac{1}{S_s^0} \middle| \mathcal{F}_t \right) d\widehat{B}_s^0,\ t \in [0, T'],$$

i.e., the reserve is just the market value of the (future part of the) payment stream $\widehat{B}^0$, which is deterministic and thus in particular "purely financial".

H. Some general remarks on the model.
Virtually all traditional life insurance products meet the assumption imposed in Paragraph C that $\widehat{B}$ be $\mathbb{F}^Z$-adapted. However, our framework also covers unit-linked products (with contractual payments linked to some financial index), even unit-linked products with $\mathbb{F}^Z$-adapted (e.g. state-dependent) guarantees. It is possible — albeit at the expense of more complex notation — to extend all results in this chapter to a more general situation with $\mathbb{F}$-adapted guarantees so as to allow, e.g., for products guaranteeing that the investments will outperform some stochastic benchmark strategy.

In practice a policy often provides the insured with certain intervention options, but we ignore this feature in order to ease the presentation. Once again we note that it is possible to include such options in a consistent manner that leaves the results of this chapter practically unaffected.

The financial market model is quite general in the sense that all its coefficients are random, but since life insurance policies are typically long-term contracts, the implied market completeness is admittedly somewhat unrealistic. However, more realistic long-term market models as such are beyond the aim and scope of the present chapter.

4.3 The optimization problem

A. The dividend method.
In this section we address the problem of determining optimal bonus strategies. However, instead of allowing for all strategies satisfying (2.6) and (2.7) we restrict

ourselves to a particular — yet fairly general — class of strategies to be characterized below. This is due to a widely established practice but it will also facilitate the optimization problem considerably.

The bonus payments are usually not determined directly, but through allocation of *dividends* and subsequent conversion(s) of dividends into future bonus payments according to a certain *conversion procedure* laid down by the company. Thus, it is the dividends (and the conversion procedure) rather than the bonus payments that are subject to direct control, and we adopt this method here. The conversion procedure can be quite complex, but we shall focus on a method that can be characterized as a variant of a bonus scheme known as *additional benefits*.

More specifically, we first assume that $\widehat{B}$ can be decomposed as

$$\widehat{B} = \widehat{B}^p + \widehat{B}^r,$$

where $\widehat{B}^p, \widehat{B}^r \in \mathcal{V}$, and where $\widehat{B}^r$ is increasing and satisfies $\widehat{B}^r_t \equiv 0$, $\forall t \in [0,T)$, for some fixed $T \in (0,T']$. We think of T as the time of retirement of the insured (hence the superscript r), which therefore marks the end of the premium payment period. The (major part of the) contractual retirement or pension benefits from time T onwards are accounted for by $\widehat{B}^r$, whereas $\widehat{B}^p$ accounts for premiums and any remaining part of the benefits not accounted for by $\widehat{B}^r$, i.e., term insurance, disability insurance etc. We shall sometimes refer to $\widehat{B}^p$ and $\widehat{B}^r$ as the *premium (payment) stream* and the *pension* or *retirement (benefit) stream*, respectively.

Obviously, (2.2) is satisfied with $\widehat{B}$ replaced by $\widehat{B}^p$ or $\widehat{B}^r$, since it holds for $\widehat{B}$. We define $\widehat{V}^p$ and $\widehat{V}^r$ as the market reserves corresponding to $\widehat{B}^p$ and $\widehat{B}^r$, respectively, (as in (2.3)) so that $\widehat{V} = \widehat{V}^p + \widehat{V}^r$.

As mentioned in the introduction we shall only consider optimal strategies during the premium payment period, and we therefore assume that dividends can only be allocated during $[0,T]$. Next, we assume that the company uses all dividends allocated during $[0,T)$ to purchase additional units of the pension benefit stream as specified by $\widehat{B}^r$ and add these to the original one so as to proportionally raise the guaranteed future pension benefits. A dividend lump sum allocated at time T, on the other hand, can be split into two parts, one that goes to increase the future pension benefits further, and one that is paid out to the insured.

In mathematical terms, we let D_t and K_t represent, respectively, the accumulated dividends and the accumulated number of guaranteed units of the basic payment process $\widehat{B}^r$ at time t, $t \in [0,T']$. As was the case for the bonus process we must have $D_{0-} \equiv 0$ and thus $K_{0-} \equiv 1$, and both D and K must be increasing. Furthermore, D, and thus also K, must be $\mathbb{F}$-adapted. We assume that $D, K \in \mathcal{V}$.

Since all dividends must be allocated by time T we must have

$$D_t = D_T,\ \forall t \in [T,T'], \tag{3.1}$$

$$K_t = K_T,\ \forall t \in [T,T']. \tag{3.2}$$

As for the specific conversion procedure we adopt a very natural one, by which the number of additional units of guaranteed pension benefits added in a small time interval equals the amount of allocated dividends divided by the present value of a single unit. More precisely, we put

$$dK_t = \frac{dD_t}{\widehat{V}^r_t},\ t \in [0, \tau),$$

where τ is the $(\mathcal{F}^Z_t)_{t\in[0,T]}$-stopping time

$$\tau = \inf\{t \in [0, T'] : \widehat{V}^r_t = 0\} \wedge T.$$

We need the restriction $t \in [0, \tau)$ because it obviously makes no sense to add additional units of $\widehat{B}^r$ if $\widehat{V}^r_t = 0$. If $\tau < T$ we must have $dD_t = dK_t = 0$ for $\tau \leq t < T$. In integral form we have

$$K_t = 1 + \int_{[0,t]} \frac{1}{\widehat{V}^r_s}\, dD_s,\ t \in [0, T). \tag{3.3}$$

As mentioned, any remaining dividends to be allocated at time T, ΔD_T, can be used to additionally increase the future pension benefits by adding ΔK_T units of $\widehat{B}^r$ (unless $\tau < T$) and/or to pay out an additional bonus lump sum (positive, of course), which consequently must equal

$$\Delta D_T - \Delta K_T \widehat{V}^r_T, \tag{3.4}$$

the dividend lump sum minus the value of the additional units of future pension benefits. Note that if a lump sum payment at time T is included in the pension benefits $\widehat{B}^r$, i.e., if $\widehat{B}^r_T = \Delta\widehat{B}^r_T > 0$, this payment is not affected by an increase in the future pension benefits made at time T, as "future" is to be interpreted in the strict sense.

To make things clear, the bonus payment process is given by

$$\begin{aligned} d\widetilde{B}_t &= 0,\ 0 \leq t < T, \\ \Delta\widetilde{B}_T &= (K_{T-} - 1)\,\Delta\widehat{B}^r_T + \Delta D_T - \Delta K_T \widehat{V}^r_T, \qquad (3.5) \\ d\widetilde{B}_t &= (K_T - 1)\, d\widehat{B}^r_t,\ T < t \leq T'. \qquad (3.6) \end{aligned}$$

Methods similar to the one outlined here are widely used in practice. Although this is mainly due to practical issues, it can also be supported theoretically inasmuch as $\widehat{B}^r$ — considered as the policyholder's desired payment profile — is maintained (see also the discussion in Paragraph C below). Of course, the possibility of an additional lump sum bonus payment at time T, which is included for the purpose of generality, may admittedly disturb this feature of the method. More general methods could be considered, but only at the expense of added complexity.

With a fixed conversion procedure (as the one described above) we consider the bonus payments process $\widetilde{B}$ as being secondary in the overall decision problem in the sense that it is the dividend process D that is immediately controllable. However, also ΔK_T is directly controllable, subject to the restriction

$$\Delta K_T \leq \frac{\Delta D_T}{\widehat{V}^r_T}, \tag{3.7}$$

which is due to the fact that the additional bonus lump sum at time T, given by (3.4), must be positive.

B. A relation for the dividend process.
Inserting (3.5) and (3.6) in (2.7) leads to the equation

$$\mathrm{E}^{\mathrm{Q}}\left(\int_{[0,T']} \frac{1}{S^0_t}\,(K_{t-}-1)\,d\widehat{B}^r_t + \frac{1}{S^0_T}\left(\Delta D_T - \Delta K_T \widehat{V}^r_T\right)\right) = -\widehat{V}_{0-}. \tag{3.8}$$

Now, using (3.2), and the fact that $\widehat{B}^r_t \equiv 0$, $\forall t \in [0,T)$, we have

$$\int_{[0,T']} \frac{1}{S^0_t}\,(K_{t-}-1)\,d\widehat{B}^r_t = (K_{T-}-1)\int_{[T,T']} \frac{1}{S^0_t}\,d\widehat{B}^r_t + \Delta K_T \int_{(T,T']} \frac{1}{S^0_t}\,d\widehat{B}^r_t.$$

By (3.3) and the law of iterated expectations (recall (2.3)) we can thus write (3.8) as

$$\mathrm{E}^{\mathrm{Q}}\left(\frac{1}{S^0_T}\widehat{V}^r_{T-}\int_{[0,T)} \frac{1}{\widehat{V}^r_t}\,dD_t + \frac{1}{S^0_T}\,\Delta D_T\right) = -\widehat{V}_{0-}.$$

Since D is increasing and fulfills (3.1), and $(\widehat{V}^r_t/S^0_t)_{t\in[0,T)}$ is a positive martingale under Q, we arrive at

$$\mathrm{E}^{\mathrm{Q}}\left(\int_{[0,T']} \frac{1}{S^0_t}\,dD_t\right) = -\widehat{V}_{0-}, \tag{3.9}$$

which corresponds to (2.6). In particular, the dividend and bonus processes must have the same initial market value, which is quite natural.

With a slight abuse of terminology we shall refer to a $(D,\Delta K_T)$-pair satisfying (3.7) and (3.9) (and the regularity conditions imposed in Paragraph A) as a *dividend strategy*, and we denote by $\mathcal{D}$ the set of dividend strategies. For a given $(D,\Delta K_T)\in\mathcal{D}$ the corresponding K and $\widetilde{B}$ are determined through (3.3) and (3.5). If another dividend strategy is considered all quantities pertaining to it will be equipped with the same superscript in an obvious way, e.g. $D', \Delta K'_T, K', \widetilde{B}'$, and so on.

C. Objectives.
Our aim is to find optimal dividend and investment strategies from the point of view of the insured, i.e., strategies that maximize the expected total utility of the payments. We take an approach based on the basic idea of the martingale method of optimal consumption and investment, which is a two-stage procedure that (in principle, at least) works as follows: First, combinations of consumption processes and terminal wealths are considered as elements of a suitable linear space, and an

optimal element among those that satisfy a natural budget constraint is identified (by use of methods from convex analysis). Second, an investment strategy that replicates the optimal consumption process/terminal wealth-pair is determined.

At this point we disregard the problem of finding a replicating strategy (which is impossible anyway; we turn to this issue in Section 4.5) and focus on finding optimal dividend strategies. We assume throughout this section that the inequality (2.4) is strict. Otherwise the whole issue of dividend optimization would be void because, as can be seen from (3.9), we would then be forced to put $D_t \equiv 0$, $\forall t \in [0,T]$.

We shall work with a rather special definition of a utility function, which will be motivated below.

Definition 4.3.1 A *utility function* is a function $u : \mathbb{R} \to [-\infty, \infty)$ such that

$$x_0 := \inf A \in [0, \infty), \tag{3.10}$$

where $A = \{x \in \mathbb{R} \,|\, u(x) > -\infty\}$, and such that on A, u is *either* constant and equal to 0 *or* strictly concave and increasing with $\lim_{x\to\infty} u'(x) = 0$.

We consider a preference structure given by

$$\Phi(D, \Delta K_T) = U\left(Y', \Delta\widehat{B}^r_T + \Delta\widetilde{B}_T\right) + U\left(Y'', K_T\right), \tag{3.11}$$

where Y' and Y'' are $\mathcal{F}^Z_T$-measurable random variables and $U : \mathbb{R}^2 \to [-\infty, \infty)$ is a measurable function such that, for each $y \in \mathbb{R}$, $U(y, \cdot) : \mathbb{R} \to [-\infty, \infty)$ is a utility function. Our objective is

$$\sup_{(D,\Delta K_T)\in\mathcal{D}^A} \mathrm{E}\left(\Phi(D, \Delta K_T)\right),$$

where $\mathcal{D}^A \subseteq \mathcal{D}$ denotes the set of *admissible* dividend strategies, i.e.,

$$\mathcal{D}^A = \left\{(D, \Delta K_T) \in \mathcal{D} \,\middle|\, \mathrm{E}\left(\left[U\left(Y', \Delta\widehat{B}^r_T + \Delta\widetilde{B}_T\right)\right]^- + \left[U\left(Y'', K_T\right)\right]^-\right) < \infty\right\},$$

where $[\cdot]^- = -\min(0, \cdot)$ is the negative part.

Some technical conditions are imposed below, but first we comment on the objective. The purpose of Y' and Y'' is to allow the preferences to be policy-path dependent. We think of the preference structure as being parameterized by Y' and Y'', and it is thus quite general (an example with two possible specifications is given below in Paragraph D). In particular, the insured may wish not to receive anything in some policy states in order to increase the benefits in others — the former then correspond to constant utility functions.

The utilities of the lump sum payment at time T and the pension benefit stream $\widehat{B}^r$ are assumed to be additive. However, the utility of the latter is measured entirely through K_T. One could have formulated the objective in terms of some utility functional (of the benefit stream), but this is a very delicate issue. The standard type of objective used in optimal consumption/investment problems in continuous time (as originally formulated in Merton (1969, 1971)) relies on the somewhat questionable assumption of time-additivity of utility. Other types of objectives with features such as local time substitution, habit formation, etc. have been proposed (see, e.g. Karatzas and Shreve (1998), Sect. 3.11 for references). By the dividend method considered in this chapter the policyholder's desired benefit profile is maintained, and we therefore completely avoid having to deal with this issue.

Now, some technical conditions. Put

$$C = \{y \in \mathbb{R} : U(y, \cdot) \text{ is constant on } A(y)\}, \tag{3.12}$$

where $A(y)$ is defined as A in Definition 4.3.1. To make sure the objective is sensible we assume that

$$\mathrm{P}\left(Y' \in C, Y'' \in C\right) < 1.$$

On the other hand we assume that $Y'' \in C$ if $\widehat{V}_T^r = 0$ (this assumption is well posed since $\{\widehat{V}_T^r = 0\} \in \mathcal{F}_T^Z$) because in that case the insured does not care about the value of K_T.

It is convenient also to define $\mathrm{E}\left(\Phi(D, \Delta K_T)\right)$ for inadmissible strategies, so we set $\mathrm{E}\left(\Phi(D, \Delta K_T)\right) = -\infty$ for any dividend strategy $(D, \Delta K_T) \in \mathcal{D} \setminus \mathcal{D}^A$. Note that, at this stage, we have not imposed conditions ensuring that

$$\sup_{(D,\Delta K_T) \in \mathcal{D}^A} \mathrm{E}\left(\Phi(D, \Delta K_T)\right) < \infty \tag{3.13}$$

nor that $\mathcal{D}^A \neq \emptyset$.

D. Example.
We provide an example of the preference structure that, although still fairly general, is suitable in most situations.

Example 4.3.2 We can work with a preference structure that corresponds to state-wise utility (as in Steffensen (2004)), which is often natural, e.g. in the situation of Examples 4.2.5, 4.2.6, and 4.2.7. To this end, it is easy to construct U such that

$$\mathrm{E}\left(\Phi(D, \Delta K_T)\right) = \sum_{e \in \mathcal{Z}} \mathrm{E}\left(1_{(Z_T = e)}\left[u^e\left(\Delta \widehat{B}_T^r + \Delta \widetilde{B}_T\right) + v^e\left(K_T\right)\right]\right) \tag{3.14}$$

or

$$\mathrm{E}\left(\Phi(D, \Delta K_T)\right) = \sum_{e \in \mathcal{Z}} \mathrm{E}\left(1_{(Z_T = e)}\left[u^e\left(\Delta \widetilde{B}_T\right) + v^e\left(K_T - 1\right)\right]\right), \tag{3.15}$$

where $u^e, v^e : \mathbb{R} \to [-\infty, \infty)$, $e \in \mathcal{Z}$, are utility functions. Here, (3.14) measures utility of the total payments, whereas (3.15) measures utility of the bonus payments.

4.4 General results on dividend optimization

A. A special class of dividend strategies.
In this paragraph we present a first result, which is based on a simple observation. Assume that $\mathcal{D}^A \neq \emptyset$ and consider an arbitrary admissible dividend strategy $(D, \Delta K_T) \in \mathcal{D}^A$. Since it is $\mathbb{F}$-adapted, the same goes for the corresponding proportionality factor process K, as already noted. In particular, D_T and K_T are $\mathcal{F}_T$-measurable. Therefore we can construct another admissible dividend strategy $(D', \Delta K'_T)$ by setting

$$\begin{aligned} D'_t &\equiv 0,\ t \in [0,T), \\ \Delta K'_T &= K_T - 1, \\ \Delta \widetilde{B}'_T &= \Delta \widetilde{B}_T. \end{aligned}$$

It follows immediately that $K'_t \equiv 1,\ t \in [0,T)$, and, after a few simple calculations, that

$$D'_T = \Delta D'_T = (K_{T-} - 1)\widehat{V}^r_{T-} + \Delta D_T.$$

We have $\Delta D'_T \neq \Delta D_T$ and $D'_T \neq D_T$ in general, but, by construction, $(D, \Delta K_T)$ and $(D', \Delta K'_T)$ yield exactly the same payment streams, so in particular we have $\mathrm{E}\left(\Phi(D', \Delta K'_T)\right) = \mathrm{E}\left(\Phi(D, \Delta K_T)\right)$. We formulate this result as a proposition. Let

$$\mathcal{D}^A_0 = \{(D, \Delta K_T) \in \mathcal{D}^A \mid D_t \equiv 0,\ \forall t \in [0,T)\}.$$

Proposition 4.4.1 *For any $(D, \Delta K_T) \in \mathcal{D}^A$ there exists a $(D', \Delta K'_T) \in \mathcal{D}^A_0$ that yields the same expected utility, i.e.,*

$$\mathrm{E}\left(\Phi(D', \Delta K'_T)\right) = \mathrm{E}\left(\Phi(D, \Delta K_T)\right).$$

Because of this result we need only consider $\mathcal{D}^A_0$ in our search for optimal strategies. However, we cannot rule out the possibility that in general $\mathcal{D}^A \setminus \mathcal{D}^A_0$ may also contain optimal strategies.

B. Optimal dividend strategies in $\mathcal{D}^A_0$.
Our aim in this paragraph is to find an optimal dividend strategy in $\mathcal{D}^A_0$. For $(D, \Delta K_T) \in \mathcal{D}^A_0$, (3.5) becomes

$$\Delta \widetilde{B}_T = \Delta D_T - \Delta K_T \widehat{V}^r_T,$$

and (3.9) may therefore be written as

$$\mathrm{E}^{\mathrm{Q}}\left(\frac{1}{S^0_T}\left(\Delta \widetilde{B}_T + \Delta K_T \widehat{V}^r_T\right)\right) = -\widehat{V}_{0-},$$

or equivalently

$$\widehat{V}_{0-} + \mathrm{E}\left(H_T\left(\Delta \widetilde{B}_T + \Delta K_T \widehat{V}^r_T\right)\right) = 0, \tag{4.1}$$

where $H = (H_t)_{t\in[0,T']}$ is the state price deflator process,

$$H_t = \frac{\Lambda_t \Gamma_t}{S_t^0},\ t \in [0, T'].$$

With a slight abuse of notation regarding Φ the optimization problem can now be formulated in terms of the $\mathcal{F}_T$-measurable random variables $\Delta\widetilde{B}_T$ and ΔK_T as

$$\sup_{(\Delta\widetilde{B}_T, \Delta K_T)} \mathrm{E}\left(\Phi(\Delta\widetilde{B}_T, \Delta K_T)\right),$$

where

$$\Phi(\Delta\widetilde{B}_T, \Delta K_T) = U\left(Y', \Delta\widehat{B}_T^r + \Delta\widetilde{B}_T\right) + U\left(Y'', 1 + \Delta K_T\right), \tag{4.2}$$

subject to the constraint (4.1) and

$$\begin{aligned} \Delta\widetilde{B}_T &\geq 0, \quad \text{a.s.}, &(4.3)\\ \Delta K_T &\geq 0, \quad \text{a.s.} &(4.4) \end{aligned}$$

In order to enforce the constraints (4.3)-(4.4) we introduce the generalized utility function $\Upsilon : \mathbb{R}^2 \times [0, \infty) \to [-\infty, \infty)$ given by

$$\Upsilon(y, x, z) = \begin{cases} U(y, x), & x \geq z, \\ -\infty, & x < z, \end{cases}$$

and — once again with a slight abuse of notation — replace the expression for Φ in (4.2) by

$$\Phi(\Delta\widetilde{B}_T, \Delta K_T) = \Upsilon\left(Y', \Delta\widehat{B}_T^r + \Delta\widetilde{B}_T, \Delta\widehat{B}_T^r\right) + \Upsilon\left(Y'', 1 + \Delta K_T, 1\right). \tag{4.5}$$

From the theory of convex optimization under constraints (see e.g. Holmes (1975) or Kreyszig (1978)) it can be deduced that the problem reduces to the unconstrained maximization problem

$$\sup_{(\Delta\widetilde{B}_T, \Delta K_T)} \mathrm{E}\left(L(\Delta\widetilde{B}_T, \Delta K_T, \xi)\right),$$

where

$$L(\Delta\widetilde{B}_T, \Delta K_T, \xi) = \Phi(\Delta\widetilde{B}_T, \Delta K_T) - \xi\left[\widehat{V}_{0-} + H_T\left(\Delta\widetilde{B}_T + \Delta K_T \widehat{V}_T^r\right)\right],$$

and $\xi > 0$ is a *Lagrange/Kuhn-Tucker multiplier.* We have

$$
\begin{aligned}
L(\Delta\widetilde{B}_T, \Delta K_T, \xi) &= \Upsilon\left(Y', \Delta\widehat{B}^r_T + \Delta\widetilde{B}_T, \Delta\widehat{B}^r_T\right) - \xi H_T \Delta\widetilde{B}_T \\
&\quad + \Upsilon\left(Y'', 1 + \Delta K_T, 1\right) - \xi H_T \Delta K_T \widehat{V}^r_T \\
&\quad - \xi \widehat{V}_{0-} \\
&= \Upsilon\left(Y', \Delta\widehat{B}^r_T + \Delta\widetilde{B}_T, \Delta\widehat{B}^r_T\right) - \xi H_T\left(\Delta\widehat{B}^r_T + \Delta\widetilde{B}_T\right) \\
&\quad + \Upsilon\left(Y'', 1 + \Delta K_T, 1\right) - \xi H_T \widehat{V}^r_T (1 + \Delta K_T) \\
&\quad - \xi\left(\widehat{V}_{0-} - H_T \Delta\widehat{B}^r_T - H_T \widehat{V}^r_T\right) \\
&\leq \widetilde{\Upsilon}\left(Y', \xi H_T, \Delta\widehat{B}^r_T\right) \qquad (4.6) \\
&\quad + \widetilde{\Upsilon}\left(Y'', \xi H_T \widehat{V}^r_T, 1\right) \\
&\quad - \xi\left(\widehat{V}_{0-} - H_T \Delta\widehat{B}^r_T - H_T \widehat{V}^r_T\right),
\end{aligned}
$$

where $\widetilde{\Upsilon}$ is the generalized *convex dual* of Υ given by

$$\widetilde{\Upsilon}(y, w, z) = \sup_{x \in \mathbb{R}} \left(\Upsilon(y, x, z) - wx\right), \ (y, w, z) \in \mathbb{R}^2 \times [0, \infty).$$

As functions of $\xi > 0$ our candidate optimal choices of $\Delta\widetilde{B}_T$ and ΔK_T, $\Delta\widetilde{B}^*_T$ and ΔK^*_T, become (recall (3.10) and (3.12))

$$\Delta\widetilde{B}^*_T(\xi) = \begin{cases} \left(x_0(Y') - \Delta\widehat{B}^r_T\right)^+, & Y' \in C, \\ \left((U')^{-1}(Y', \xi H_T) - \Delta\widehat{B}^r_T\right)^+, & Y' \notin C, \end{cases} \tag{4.7}$$

$$\Delta K^*_T(\xi) = \begin{cases} \left(x_0(Y'') - 1\right)^+, & Y'' \in C, \\ \left((U')^{-1}(Y'', \xi H_T \widehat{V}^r_T) - 1\right)^+, & Y'' \notin C, \end{cases} \tag{4.8}$$

where $(U')^{-1} = (U')^{-1}(y, \cdot)$ is the inverse of $U' = U'(y, \cdot)$, defined for $y \notin C$ and extended if necessary from its natural domain $(0, U'(x_0(y)+))$ to $(0, \infty)$ by $(U')^{-1}(y, w) = x_0(y)$, $w \geq U'(x_0(y)+)$.

Indeed, a few calculations show that the expressions (4.7)-(4.8) are sufficient for the inequality (4.6) to become an equality. However, only (4.7) and the expression for $\Delta K^*_T(\xi)$ when $Y'' \notin C$ are also necessary. If $Y'' \in C$ one has to take the value of $\widehat{V}^r_T$ into account because if $\widehat{V}^r_T = 0$ it is only necessary that $\Delta K^*_T(\xi) \geq (x_0(Y'') - 1)^+$, i.e., $\Delta K^*_T(\xi)$ can then be chosen arbitrarily in $[(x_0(Y'') - 1)^+, \infty)$.

We need some technical assumptions in order to actually prove that our candidates are admissible and optimal. We define the function $\psi : (0, \infty) \to (0, \infty]$ by

$$\psi(\xi) = \mathrm{E}\left(H_T\left(\Delta\widetilde{B}^*_T(\xi) + \Delta K^*_T(\xi)\widehat{V}^r_T\right)\right), \ \xi > 0. \tag{4.9}$$

Assumption 4.4.2 *For every $\xi > 0$,*

$$\psi(\xi) < \infty.$$

This assumption holds under certain conditions on U and H, but we shall not go into a detailed study here. If it holds, then ψ is easily seen to be strictly decreasing and fulfill $\psi(\xi) \to \infty$ for $\xi \searrow 0$.

Assumption 4.4.3 *The function ψ satisfies*

$$\lim_{\xi \to \infty} \psi(\xi) < -\widehat{V}_{0-}.$$

This assumption is necessary to ensure that the problem is well posed; it is easily verified that $\mathcal{D}^A = \emptyset$ if it does not hold.

Assumption 4.4.4 *There exists a function $F : \mathbb{R} \to [0, \infty)$ such that*

$$\Delta \widehat{B}_T^r \leq F(Y'),\ 1 \leq F(Y''), \text{ a.s.,}$$

and

$$\mathrm{E}\left(HF(Y') + [U(Y', F(Y'))]^- + HF(Y'') + [U(Y'', F(Y'')]^-\right) < \infty.$$

In combination, Assumptions 4.4.3 and 4.4.4 ensure the existence of admissible strategies.

Proposition 4.4.5 *Under Assumptions 4.4.2, 4.4.3, and 4.4.4, the dividend strategy given by $(\Delta \widetilde{B}_T^*(\xi^*), \Delta K_T^*(\xi^*))$, where $\xi^* > 0$ is uniquely determined by $\psi(\xi^*) = -\widehat{V}_{0-}$, is admissible and optimal. If (3.13) also holds, then this strategy is the only optimal one in $\mathcal{D}_0^A$ (except for the abovementioned arbitrariness regarding $\Delta K_T^*(\xi)$ on the set $(\widehat{V}_T^r = 0)$).*

Proof. Under Assumptions 4.4.2 and 4.4.3, the strategy $(\Delta \widetilde{B}_T^*(\xi^*), \Delta K_T^*(\xi^*))$ is well defined. To see that it is admissible (i.e., belongs to $\mathcal{D}^A$), consider the strategy

$$(\Delta \widetilde{B}_T, \Delta K_T) := (L(Y') - \Delta \widehat{B}_T^r, L(Y'') - 1).$$

From the analysis above we have

$$\Upsilon\left(Y', \Delta \widehat{B}_T^r + \Delta \widetilde{B}_T^*(\xi^*), \Delta \widehat{B}_T^r\right) - \xi^* H_T \left(\Delta \widehat{B}_T^r + \Delta \widetilde{B}_T^*(\xi^*)\right) = \widetilde{\Upsilon}\left(Y', \xi^* H_T, \Delta \widehat{B}_T^r\right) \tag{4.10}$$

and

$$\begin{aligned} \widetilde{\Upsilon}\left(Y', \xi^* H_T, \Delta \widehat{B}_T^r\right) &\geq \Upsilon\left(Y', \Delta \widehat{B}_T^r + \Delta \widetilde{B}_T, \Delta \widehat{B}_T^r\right) - \xi^* H_T \left(\Delta \widehat{B}_T^r + \Delta \widetilde{B}_T\right) \quad (4.11) \\ &= \Upsilon\left(Y', F(Y'), \Delta \widehat{B}_T^r\right) - \xi^* H_T F(Y') \\ &= U(Y', F(Y')) - \xi^* H_T F(Y'). \end{aligned}$$

Thus,

$$\left[\Upsilon\left(Y', \Delta\widehat{B}_T^r + \Delta\widetilde{B}_T^*(\xi^*), \Delta\widehat{B}_T^r\right)\right]^- \leq \left[U(Y', F(Y'))\right]^- + \xi^* H_T F(Y'),$$

so by Assumption 4.4.4 we have

$$\mathrm{E}\left(\left[U(Y', \Delta\widehat{B}_T^r + \Delta\widetilde{B}_T^*(\xi^*))\right]^-\right) < \infty.$$

Similar calculations show that

$$\mathrm{E}\left(\left[U(Y'', 1 + \Delta K_T^*(\xi^*))\right]^-\right) < \infty,$$

and we conclude that $(\Delta\widetilde{B}_T^*(\xi^*), \Delta K_T^*(\xi^*)) \in \mathcal{D}^A$.

Now, consider an arbitrary admissible strategy $(\Delta\widetilde{B}_T', \Delta K_T')$. Using (4.10) and (4.11) with $(\Delta\widetilde{B}_T, \Delta K_T)$ replaced by $(\Delta\widetilde{B}_T', \Delta K_T')$, and the fact that (4.1) must be satisfied by $(\Delta\widetilde{B}_T^*(\xi^*), \Delta K_T^*(\xi^*))$ as well as $(\Delta\widetilde{B}_T', \Delta K_T')$, it is easily seen that

$$\mathrm{E}\left(\Phi(\Delta\widetilde{B}_T^*(\xi^*), \Delta K_T^*(\xi^*))\right) \geq \mathrm{E}\left(\Phi(\Delta\widetilde{B}_T', \Delta K_T')\right),$$

i.e., $(\Delta\widetilde{B}_T^*(\xi^*), \Delta K_T^*(\xi^*))$ is optimal.

Finally, the inequality (4.11) is strict unless $\Delta\widetilde{B}_T = \Delta\widetilde{B}_T^*(\xi^*)$, and the corresponding inequality involving Y'' is strict unless $\Delta K_T = \Delta K_T^*(\xi^*)$ (except on $(\widehat{V}_T^r = 0)$, where the condition is $\Delta K_T \geq \Delta K_T^*(\xi^*)$). Thus, the asserted uniqueness in $\mathcal{D}_0^A$ regarding the optimal strategy is valid if (3.13) also holds. □

Proposition 4.4.5 provides optimal strategies, albeit only in a somewhat semi-explicit fashion. If $Y' \notin C$, the total lump sum payment at time T is given by

$$\begin{aligned}\Delta\widehat{B}_T^r + \Delta\widetilde{B}_T^*(\xi) &= \Delta\widehat{B}_T^r + \left((U')^{-1}(Y', \xi^* H_T) - \Delta\widehat{B}_T^r\right)^+ \\ &= (U')^{-1}(Y', \xi^* H_T) + \left(\Delta\widehat{B}_T^r - (U')^{-1}(Y', \xi^* H_T)\right)^+.\end{aligned}$$

It is seen to resemble the combined payoff from a zero-coupon bond and a European call option on the *unconstrained* optimal (total) payment (which is obtained by putting $\Delta\widehat{B}_T^r \equiv 0$) with strike price equal to the payoff from the zero-coupon bond, or, equivalently, the combined payoff from the unconstrained optimal (total) payment and a corresponding European put option. Thus, it generalizes the well-known structure of the optimal wealth obtained in purely financial optimization problems with *portfolio insurance* constraining the terminal wealth to be greater than or equal to some $K \geq 0$, which can be a deterministic constant as in, e.g. Cox and Huang (1989) ($K = 0$) and Grossman and Zhou (1996) ($K \geq 0$), or a general stochastic benchmark as in, e.g. Tepla (2001). Similar remarks are valid also for the optimal total number of units of pension benefits, $1 + \Delta K_T^*(\xi^*)$.

Remark 4.4.6 If (3.13) does not hold, there may be several optimal strategies, which is an unsatisfactory situation. We do not provide general conditions ensuring that (3.13) holds, but we stress that it should be checked in applications of Proposition 4.4.5.

C. Optimal dividend strategies in $\mathcal{D}^A$.
The remark made after Proposition 4.4.1 stating that $\mathcal{D}^A \setminus \mathcal{D}_0^A$ may contain optimal strategies in general, i.e., that it may not be strictly suboptimal in general to allocate dividends before T, leads to the question of when this is actually the case. The following corollary settles this issue.

Corollary 4.4.7 *Under Assumptions 4.4.2, 4.4.3, and 4.4.4, a dividend strategy* $(D, \Delta K_T) \in \mathcal{D}^A$ *is optimal if*

$$\begin{aligned} K_T &= 1 + \Delta K_T^*(\xi^*) \text{ on } (\widehat{V}_T^r > 0), && (4.12) \\ \Delta D_T &= \Delta \widetilde{B}_T^*(\xi^*) - (K_{T-} - 1)\Delta \widehat{B}_T^r + \Delta K_T \widehat{V}_T^r, && (4.13) \end{aligned}$$

and (4.12)-(4.13) are necessary conditions for optimality if (3.13) also holds. Furthermore, (4.12) can be met if and only if

$$\mathrm{P}\left((\widehat{V}_T^r > 0) \cap (K_t > 1 + \Delta K_T^*(\xi^*)) \,\middle|\, \mathcal{F}_t \right) = 0, \text{ a.s., } \forall t \in [0, T]. \quad (4.14)$$

Proof. Under Conditions (4.12) and (4.13) the bonus payments exactly match those obtained by using the optimal strategy in $\mathcal{D}_0^A$ given in Proposition 4.4.5. For the bonus pension benefits this follows immediately from (4.12), which easily translates into a condition on D and ΔK_T by use of (3.3). For the lump sum bonus at time T it easily follows from (3.5). The first two assertions are then easy consequences of Propositions 4.4.1 and 4.4.5.

The last assertion follows from the fact that K must be increasing. □

Condition (4.14) tells us to what extent it is possible to allocate dividends during $[0, T)$ and still obtain an optimal strategy. As long as the pension benefits are non-null with strictly positive probability (conditionally, given all current information), the total number of guaranteed units of the pension benefits must be smaller than $(1 + \Delta K_T^*(\xi^*))$ on $(\widehat{V}_T^r > 0)$ (except possibly on a null-set), that is, smaller than the conditional essential infimum of $(1 + \Delta K_T^*(\xi^*))$ given $(\widehat{V}_T^r > 0)$ and all current information. It can be shown that unless $H_T \widehat{V}_T^r$ is bounded from above (conditionally), which is typically not the case, this conditional essential infimum is 1, and it is thus strictly suboptimal to allocate dividends during $[0, T)$.

D. Examples.
The following example specifies a particular preference structure featuring constant relative risk aversion in a generalized form. We put up the optimal $\mathcal{D}_0^A$-strategy.

Example 4.4.8 Suppose U is given by

$$U(y,x) = \alpha(y)u^{(\gamma(y))}(x - x_0(y)),\ (y,x) \in \mathbb{R}^2,$$

where $\alpha(y) \in [0,\infty)$, and $u^{(\gamma)}$ is a CRRA utility function with relative risk aversion coefficient $\gamma \geq 0$, i.e.,

$$u^{(\gamma)}(x) = \begin{cases} \frac{x^{1-\gamma}}{1-\gamma}, & \gamma \neq 1, \\ \log(x), & \gamma = 1, \end{cases}$$

for $x > 0$, $u^{(\gamma)}(0) = \lim_{x\searrow 0} u^{(\gamma)}(x)$, and $u^{(\gamma)}(x) = -\infty$ for $x < 0$. Then $C = \{y \in \mathbb{R} : \alpha(y) = 0\}$, and

$$\Delta\widetilde{B}^*_T(\xi) = \begin{cases} \left(x_0(Y') - \Delta\widehat{B}^r_T\right)^+, & \alpha(Y') = 0, \\ \left((\xi H_T/\alpha(Y'))^{-1/\gamma(Y')} + x_0(Y') - \Delta\widehat{B}^r_T\right)^+, & \alpha(Y') > 0, \end{cases} \tag{4.15}$$

$$\Delta K^*_T(\xi) = \begin{cases} (x_0(Y'') - 1)^+, & \alpha(Y'') = 0, \\ \left((\xi H_T\widehat{V}^r_T/\alpha(Y''))^{-1/\gamma(Y'')} + x_0(Y'') - 1\right)^+, & \alpha(Y'') > 0. \end{cases} \tag{4.16}$$

4.5 Quadratic hedging strategies

A. The investment issue.
While the previous section dealt with optimization of the dividend strategy, this section deals with optimization of the company's investment strategy in relation to the policy. These issues are intimately related; in fact, in purely financial problems of this type the investment issue and the consumption/wealth issue are essentially two sides of the same story because the optimal investment strategy, by definition, is the one that yields the optimal consumption/wealth process. In the problem studied here, however, there is an unhedgeable risk source, Z, and therefore the overall optimization problem faced by the company is not settled by the results of the previous section.

Motivated by this informal discussion we proceed to consider the company's investment problem from scratch, taking the bonus payment process $\widetilde{B}$ as given. It is assumed that $\widetilde{B}$ satisfies (2.7), but for the purpose of generality we do not assume that $\widetilde{B}$ is optimal with respect to our objective (although it may be, of course) as this is unnecessary for the analysis.

B. Risk minimization.
Our objective will be to find strategies that hedge the payments as closely as possible in some sense. To this end we shall consider quadratic hedging approaches, by which the aim is risk minimization. More precisely, we consider the Q-*risk-minimizing* and P-*locally risk-minimizing* strategies. Our aim is to make a brief comparison of the two strategies and their implications. Thus, we shall not go into a detailed study of the general theory but just give a sketchy presentation that is sufficient to get an understanding of the ideas. For details and discussions we refer

to Schweizer (2001) and its references (see in particular Møller (2001) for numerous applications in insurance).

Let us first present the ideas. A risk-minimizing strategy is an investment strategy that, at any time during the policy term, minimizes the expected value of the squared discounted total remaining net costs pertaining to the policy. However, risk-minimizing strategies only exist in general when the discounted price processes are (local) martingales, so this forces us to work under Q (which explains the term Q-*risk-minimizing*). A locally risk-minimizing strategy is an investment strategy that, at any time during the policy term, minimizes the expected value of the squared discounted net costs incurred over the next (infinitesimally) small time interval, loosely speaking. Such a strategy exists in general, so we can work under P, which is perhaps more natural from a theoretical point of view. This also explains the term P-*locally risk-minimizing*.

Now, the total payment process pertaining to the policy is given by $B = \widehat{B} + \widetilde{B}$, and the value of the discounted (net) payments from the company to the policyholder in $[0, t]$ is

$$A_t = \int_{[0,t]} \frac{1}{S_s^0} \, dB_s, \; t \in [0, T'].$$

In order to hedge the policy payments as closely as possible the company employs an investment strategy formalized by a stochastic process $(\pi_t^0, \pi_t)_{t \in [0,T']}$, where π^0 is real-valued and $\pi = (\pi^1, \ldots, \pi^d)'$ is $\mathbb{R}^d$-valued, and both are assumed to be progressively measurable with respect to $\mathbb{F}$. For $i = 0, \ldots, d$ and $t \in [0, T']$, π_t^i represents the discounted amount invested in the i'th asset at time t. Thus, the discounted total value of the investment portfolio at time $t \in [0, T']$ is given by

$$X_t(\pi^0, \pi) = \sum_{i=0}^{d} \pi_t^i.$$

Since $d(S_t^0 / S_t^0) = 0$, and

$$d(S_t^i / S_t^0) = S_t^i / S_t^0 \left(\sigma_t^{i\cdot} \lambda_t \, dt + \sigma_t^{i\cdot} \, dW_t \right) = S_t^i / S_t^0 \sigma_t^{i\cdot} \, dW_t^{\mathrm{Q}},$$

for $i = 1, \ldots, d$, the discounted *gains process*, measuring the accumulated discounted investment gains of the portfolio, is given by

$$G_t(\pi^0, \pi) = \int_0^t \pi_s' \sigma_s \, dW_s^{\mathrm{Q}}, \; t \in [0, T'].$$

Remark 4.5.1 This choice of terminology may, unfortunately, be somewhat misleading: In general, $G_t(\pi^0, \pi)$ is *not* equal to the discounted value of the accumulated investment gains over $(0, t]$ (this is only the case in general for self-financing investment strategies starting at 0). The term *gains process* seem to be standard in the theory of quadratic hedging approaches (see, e.g. Schweizer (2001)), where one always works with discounted price processes. Since we have worked with undiscounted price processes so far, we have added the qualifier "discounted" here in order to emphasize that G is not the (undiscounted) gains process.

Let $L^2(W)$ denote the set of progressively measurable d-dimensional processes $\theta = (\theta_t)_{t\in[0,T']}$ such that $\mathrm{E}\left(\int_0^{T'} \|\theta_t\|^2\, dt\right) < \infty$. We impose the restriction that the company is only allowed to use investment strategies (π^0, π) satisfying

$$\sigma'\pi \in L^2(W), \tag{5.1}$$

$$\mathrm{E}\left(\int_0^{T'} |\pi_t'\sigma_t\lambda_t|\, dt\right)^2 < \infty,$$

and

$$\mathrm{E}\left(X_t(\pi^0,\pi)\right)^2 < \infty,\ \forall t \in [0,T']. \tag{5.2}$$

Since $d\mathrm{Q}/d\mathrm{P} \in L^2(\Omega, \mathcal{F}_{T'}, \mathrm{P})$, the condition (5.1) ensures in particular that the gains process, $G(\pi^0,\pi)$, is a $(\mathrm{Q},\mathbb{F})$-martingale, so arbitrage is ruled out.

Now we define the *cost process*, which measures the discounted accumulated total net costs and thus the total financial impact that the policy has on the company. It is given by

$$C_t(\pi^0,\pi) = A_t + X_t(\pi^0,\pi) - G_t(\pi^0,\pi),\ t \in [0,T'],$$

i.e., the (discounted) accumulated policy payments plus the (discounted) value of the investment portfolio minus the (discounted) accumulated investment gains. It is important to note that the investment strategy (π^0,π) is not assumed to be self-financing; if this were the case we would have $X_t - G_t = X_0$, $\forall t \in [0,T']$. We also emphasize that $X_t(\pi^0,\pi)$ does not represent the company's total assets, but only the assets allocated to cover the (future) net payments pertaining to the policy.

At time T', when the contract is finally settled, this investment portfolio is liquidated and added to the company's other assets, i.e.,

$$X_{T'}(\pi^0,\pi) = 0,\ \text{a.s.} \tag{5.3}$$

Now, the aim is to minimize the (squared) fluctuations of C in some sense. As already mentioned, a Q-risk-minimizing (resp. P-locally risk-minimizing) strategy is one that, at any time $t \in [0,T']$, minimizes the expected squared total discounted net costs over the remaining period $(t,T']$ (resp. over an (infinitesimally) small interval $(t,t+dt)$).

Square integrability (under P) of the components of C is ensured by the assumptions made so far. When we work under Q we need to assume that (2.2), (2.5), (5.1), and (5.2) all hold with $\mathrm{E}(\cdot)$ replaced by $\mathrm{E}^{\mathrm{Q}}(\cdot)$.

In both cases the sought strategy can be obtained from the Galtchouk-Kunita-Watanabe decomposition of the so-called *intrinsic value process*, calculated under a suitable equivalent martingale measure. In the case of Q-risk-minimizing strategies this measure is Q itself, and in the case of P-locally risk-minimizing strategies it is the so-called *minimal* martingale measure $\check{\mathrm{P}}$, which in our model is given by

$$\frac{d\check{\mathrm{P}}}{d\mathrm{P}} = \Lambda_{T'},$$

that is, $\check{\mathrm{P}}|_{\mathcal{F}^W_T} = \mathrm{Q}|_{\mathcal{F}^W_T}$ and $\check{\mathrm{P}}|_{\mathcal{F}^Z_T} = \mathrm{P}|_{\mathcal{F}^Z_T}$.

For $\mathrm{R} \in \{\mathrm{Q}, \check{\mathrm{P}}\}$, the intrinsic value process $M(\mathrm{R}) = (M_t(\mathrm{R}))_{t\in[0,T']}$ is given by

$$M_t(\mathrm{R}) = \mathrm{E}^{\mathrm{R}}\left(A_{T'} \,|\, \mathcal{F}_t \right) = \mathrm{E}^{\mathrm{R}}\left(\int_{[0,T']} \frac{1}{S^0_s}\, dB_s \,\middle|\, \mathcal{F}_t \right),\ t \in [0, T'].$$

Being a square-integrable $(\mathbb{F}, \mathrm{R})$-martingale, it admits a unique Galtchouk-Kunita-Watanabe decomposition, i.e., there exists a $g(\mathrm{R}) = (g^1(\mathrm{R}), \ldots, g^d(\mathrm{R}))' \in L^2(W^{\mathrm{Q}})$ and a zero-mean square-integrable martingale $L(\mathrm{R})$ orthogonal to $\int \theta\, dW^{\mathrm{Q}}$ for any $\theta \in L^2(W^{\mathrm{Q}})$, such that

$$M_t(\mathrm{R}) = \mathrm{E}^{\mathrm{R}}\left(A_{T'} \,|\, \mathcal{F}_0 \right) + \int_0^t g'_s(\mathrm{R})\, dW^{\mathrm{Q}}_s + L_t(\mathrm{R}), \quad \text{a.s.}, \forall t \in [0, T']. \tag{5.4}$$

Now, the Q-risk-minimizing strategy $(\pi^0(\mathrm{Q}), \pi(\mathrm{Q}))$ and the P-locally risk-minimizing strategy $(\pi^0(\check{\mathrm{P}}), \pi(\check{\mathrm{P}}))$ are given by

$$\pi_t(\mathrm{R}) = \left(g'_t(\mathrm{R})\sigma_t^{-1}\right)', \tag{5.5}$$

$$\pi^0_t(\mathrm{R}) = M_t(\mathrm{R}) - A_t - \sum_{i=1}^{d} \pi_t, \tag{5.6}$$

with $\mathrm{R} = \mathrm{Q}$ and $\mathrm{R} = \check{\mathrm{P}}$, respectively. This follows from Schweizer (2001), Theorems 2.4 and 3.5 (although the latter only guarantees that $(\pi^0(\check{\mathrm{P}}), \pi(\check{\mathrm{P}}))$ is P-*pseudo-locally risk-minimizing*).

Note that $M_0(\mathrm{Q}) = \mathrm{E}^{\mathrm{Q}}\left(A_{T'} \,|\, \mathcal{F}_0 \right) = 0$ by (2.6), whereas, in general, $M_0(\check{\mathrm{P}}) = \mathrm{E}^{\check{\mathrm{P}}}\left(A_{T'} \,|\, \mathcal{F}_0 \right) \neq 0$. The total market reserve V can be written as

$$V_t = S^0_t\, \mathrm{E}^{\mathrm{Q}}\left(\int_{(t,T']} \frac{1}{S^0_s}\, dB_s \,\middle|\, \mathcal{F}_t \right) = S^0_t\,(M_t(\mathrm{Q}) - A_t). \tag{5.7}$$

Thus, at any time $t \in [0, T']$ the value of the Q-risk-minimizing portfolio is equal to the total market reserve, and in particular it is 0 at time $0-$. In comparison, the value of the P-locally risk-minimizing portfolio is given by $S^0_t\,(M_t(\check{\mathrm{P}}) - A_t)$, which in general differs from V_t unless $\mathrm{Q} = \check{\mathrm{P}}$, in particular its value is in general different from 0 at time $0-$. Therefore, the Q-risk-minimizing strategy may be considered more natural from a practical point of view.

Usually, the valuation measure Q has a (strictly positive) built-in risk loading due to the non-hedgeable policy risk, which means that

$$V_t = S^0_t\, \mathrm{E}^{\mathrm{Q}}\left(\int_{(t,T']} \frac{1}{S^0_s}\, dB_s \right) > S^0_t\, \mathrm{E}^{\check{\mathrm{P}}}\left(\int_{(t,T']} \frac{1}{S^0_s}\, dB_s \right),\ \text{a.s.}, \forall t \in [0-, T']. \tag{5.8}$$

In this case the Q-risk-minimizing strategy is thus also the more prudent of the two.

From (5.7) it is further seen that

$$d\left(V_t/S_t^0\right) = dM_t(Q) - dA_t = \pi_t'(\mathrm{Q})\sigma_t\, dW_t^{\mathrm{Q}} + dL_t(\mathrm{Q}) - 1/S_t^0\, dB_t,$$

so if the company employs the Q-risk-minimizing strategy, then at any time $t \in [0, T']$ the position in the risky assets, $\pi_t'(\mathrm{Q})$, is chosen such that the part of the changes in the reserve caused by fluctuations in W in a small time interval $(t, t+dt)$ is hedged by the investment gains in that interval. It is intuitively clear (and it has been proved under certain regularity conditions in Møller (2001); see also Example 4.6.1 in Section 4.6 below) that this amounts to choosing $\pi_t'(\mathrm{Q})$ as if the remaining payment stream $(dB_s)_{s\in(t,T']}$ were replaced by the conditional expected remaining payment stream (with respect to Q), given all current information and the (future) information in $\mathbb{F}^W$, i.e.,

$$\left(\mathrm{E}^{\mathrm{Q}}\left(dB_s \,|\, \mathcal{F}_t \vee \mathcal{F}_s^W\right)\right)_{s\in(t,T']}.$$

In other words, $\pi_t'(\mathrm{Q})$ is obtained by "integrating out" the remaining uncertainty regarding Z and then forming the initial portfolio of the hedging strategy of this artificial payment stream — such a strategy exists because the artificial payment stream is $(\mathcal{F}_s^W)_{s\in(t,T']}$-adapted. In particular, $\mathbb{F}^Z$-adapted payments (e.g. the contractual ones given by $\widehat{B}$), must be backed by zero-coupon bonds. Furthermore, explicit expressions for the Q-risk-minimizing portfolio strategy when $\widetilde{B}$ has the optimal form obtained in Section 4.4 can be obtained in many special cases of interest, see, e.g. Cox and Huang (1989), Grossman and Zhou (1996), and Tepla (2001). Note, however, that the artificial payment streams change constantly, so the portfolio must be updated accordingly.

We also see that the reserve, which is a liability that "belongs" to the insured, is maintained by adding investment gains from the corresponding assets, subtracting payments to the insured and then adding (the undiscounted value of) $dL_t(\mathrm{Q})$; this last term must be financed by the equity of the company.

Similar observations can be made for the P-locally risk-minimizing strategy. In particular the portfolio is in that case chosen so as to hedge the fluctuations with respect to W of what could appropriately be called the $\check{\mathrm{P}}$-reserve.

As for the cost process, we find by use of (5.4) and (5.6) that

$$\begin{aligned} C_t(\pi^0(\mathrm{R}), \pi(\mathrm{R})) &= A_t + X_t(\pi^0(\mathrm{R}), \pi(\mathrm{R})) - G_t(\pi^0(\mathrm{R}), \pi(\mathrm{R})) \\ &= M_0(\mathrm{R}) + L_t(\mathrm{R}),\ t \in [0, T'], \end{aligned}$$

i.e., C is a $(\mathbb{F}, \mathrm{R})$-martingale as it should be (Schweizer (2001)). If the company follows the Q-risk-minimizing strategy, the cost process will thus be the zero-mean $(\mathbb{F}, \mathrm{Q})$-martingale $L(\mathrm{Q})$, but it will not be a martingale in general under the real measure P. In particular it will have a systematic drift term representing the policy risk loading. In comparison, if the company follows the P-locally risk-minimizing strategy, the cost process will be a martingale under P, but with a non-zero initial value. More explicit results in this direction are presented in Section 4.6.

Remark 4.5.2 If Z also has integrable variation under Q, then, for $\mathrm{R} \in \{\mathrm{Q}, \check{\mathrm{P}}\}$, it follows from Jacod and Shiryaev (2003), Theorem 4.29, that $L_t(\mathrm{R})$ has the form

$$L_t(\mathrm{R}) = \int_{[0,t]} h_s(\mathrm{R})\, d(Z - Z^{\mathrm{R}})_s,\ 0 \le t \le T',$$

for some predictable process $h(\mathrm{R})$, where Z^{R} is the R-compensator of Z. In practice, life insurance companies typically have a large number (say n) of policies, and apart from systematic risks such as increasing longevity, which is currently a major issue in pension insurance, it is reasonable to assume that the individual policy processes $Z^1, \dots, Z^n$ are mutually independent. The corresponding loss processes $L^1, \dots, L^n$ are then mutually uncorrelated (but not independent!) martingales; this forms the foundation for the fundamental diversification argument in actuarial mathematics.

4.6 A worked-out example

We dedicate this entire section to an example with Z in its perhaps simplest non-trivial and fairly realistic form, that is, an intensity-driven survival model. Still, it is probably the most important special case, and sufficient to illustrate our results to a great extent.

Example 4.6.1 Suppose we are in the situation of Example 4.2.6. Let T^{01} denote the time of death, and assume that $\mathrm{Q}\left(T^{01} > T\right) > 0$. It is well-known (e.g. Jacod and Shiryaev (2003), Theorem 5.43) that Γ has the form

$$\Gamma_t = \begin{cases} e^{\int_0^t \kappa_s \mu_s^{01}\, ds}, & Z_t = 0, \\ e^{\int_0^{T^{01}} \kappa_s \mu_s^{01}\, ds} \left(1 + \kappa_{T^{01}}\right)^{-1}, & Z_t = 1, \end{cases}$$

for some deterministic, measurable function $\kappa : [0, T'] \to (-1, \infty)$ such that

$$\int_0^t \mu_s^{01}\, ds < \infty \Rightarrow \int_0^t \mu_s^{01}(1 + \kappa_s)\, ds < \infty,\ t \in [0, T'],$$

and Z admits the mortality intensity $(\mu_t^{01}(1 + \kappa_t))_{t \in [0,T']}$ under P. Put $\Gamma_T^0 = e^{\int_0^T \kappa_s \mu_s^{01}\, ds}$.

We begin with a discussion of the Q-risk-minimizing and P-locally risk-minimizing strategies and thus take the total payment process $B = \widehat{B} + \widetilde{B}$ as given at this point (as in Section 4.5). In this example B has a particularly simple form, and it can be shown that the last two components of (5.4) are given by

$$g_t(\mathrm{R}) = \frac{1_{(Z_t = 0)}}{\mathrm{R}\,(Z_t = 0)} \eta_t(\mathrm{R}),\ \mathrm{R} = \mathrm{Q}, \check{\mathrm{P}}, \tag{6.1}$$

$$L_t(\mathrm{Q}) = \int_{[0,t]} \frac{1_{(Z_{s-} = 0)}}{S_s^0} (b_s^{01} - V_s^0) \left(dN_s^{01} - \mu_s^{01}\, ds\right), \tag{6.2}$$

$$L_t(\check{\mathrm{P}}) = \int_{[0,t]} \frac{1_{(Z_{s-} = 0)}}{S_s^0} (b_s^{01} - V_s^0) \left(dN_s^{01} - \mu_s^{01}(1 + \kappa_s)\, ds\right), \tag{6.3}$$

$t \in [0, T']$, where $\eta(\mathrm{R})$ is the integrand in the stochastic representation $M_t^W(\mathrm{R}) = M_0(\mathrm{R}) + \int_0^t \eta_s'(\mathrm{R})\, dW_s^{\mathrm{Q}}$ of the $(\mathbb{F}^W, \mathrm{R})$-martingale $M^W(\mathrm{R}) = (M_t^W(\mathrm{R}))_{t\in[0,T']}$ defined by $M_t^W(\mathrm{R}) = \mathrm{E}^{\mathrm{R}}\left(A_{T'} \,|\, \mathcal{F}_t^W\right)$, $t \in [0, T']$, b_t^{01} denotes the total amount payable upon death at time t, and V_t^0 denotes the total market reserve if the insured is alive at time t. The Q-risk-minimizing strategy and the P-locally risk-minimizing strategy are given by (5.5)-(5.6) with $g(\mathrm{R})$ given by (6.1).

In agreement with the observations made in the general case in Section 4.5, we see that according to the Q-risk-minimizing strategy the company should always invest as if the remaining total policy payments were replaced by the conditional expected payments (wrt. Q) given all current information and the (future) information in $\mathbb{F}^W$, that is, as if the remaining uncertainty with respect to Z was "integrated out". The same goes for the P-locally risk-minimizing strategy except for the fact that P (or $\check{\mathrm{P}}$) should be used instead of Q. In particular we note that $\mathbb{F}^Z$-adapted components of the payment process should be backed by zero-coupon bonds and that the investment portfolio should be liquidated at the time of death of the insured.

The dynamics of the (discounted) total market reserve can be written as

$$d(V_t/S_t^0) = g_t(\mathrm{Q})dW_t^{\mathrm{Q}} - \frac{1_{(Z_{t-}=0)}}{S_t^0}(b_t^{01} - V_t^0)\mu_t^{01}\, dt - \frac{1_{(Z_{t-}=0)}}{S_t^0} V_t^0\, dN_t^{01} - \frac{1_{(Z_{t-}=0)}}{S_t^0}\, dB_t^0.$$

The first term is matched by the (discounted) return on the portfolio if the Q-risk-minimizing strategy is used, the second term represents subtraction of the actuarial *risk premium*, the third term accounts for the change in the reserve upon death, and the fourth term constitutes the state-wise payments. Note that if the P-locally risk-minimizing strategy is used, the (discounted) gains on the portfolio will not match the first term (unless $\mathrm{Q} = \check{\mathrm{P}}$).

Now, from the previous section we know that the cost process associated with the Q-risk-minimizing strategy is the zero-mean $(\mathbb{F}, \mathrm{Q})$-martingale $L(\mathrm{Q})$, which, however, is not a martingale under P. We can now elaborate on these observations. From (6.2)-(6.3) we get

$$L_t(\mathrm{Q}) = L_t(\check{\mathrm{P}}) - \int_{[0,t]} \frac{1_{(Z_{s-}=0)}}{S_s^0}(V_s^0 - b_s^{01})\mu_s^{01}\kappa_s\, ds,\ t \in [0, T'],$$

and since $\check{\mathrm{P}}|_{\mathcal{F}_{T'}^Z} = \mathrm{P}|_{\mathcal{F}_{T'}^Z}$, it is seen that the cost process has a non-zero compensator under P (unless $\kappa \equiv 0$). As mentioned in the previous section this represents the risk loading: In any small time interval $(t, t+dt)$ the company makes the systematic profit $1_{(Z_{t-}=0)}(V_t^0 - b_t^{01})\mu_t^{01}\kappa_t\, dt$ (which is positive if, e.g. $\mathrm{sgn}(\kappa_t) = \mathrm{sgn}(V_t^0 - b_t^{01}),\ \forall t \in [0, T']$) to compensate for the genuine risk it has taken on. Accordingly, κ can be interpreted as the company's assessment of a fair price of mortality risk, cf. Remark 4.2.2). In comparison, if the company follows the P-locally risk-minimizing strategy, the cost process will be the $(\mathbb{F}, \mathrm{P})$-martingale $L(\check{\mathrm{P}})$, which in general has non-zero mean. In the case with a strictly positive risk loading we have $C_0 = M_0(\check{\mathrm{P}}) < 0$ (this follows from (5.8) and (2.6)). Thus, the company initiates the policy with a strictly negative cost (i.e., a strictly positive surplus) but makes no expected profits during the term of the policy.

We now turn our attention to the issue of optimality of the (bonus) payments, and to this end we adopt the preference structure of Example 4.4.8. We suppose further that $\alpha(Y') = \alpha(Y'') = 1_{(Z_T=0)}$ and $\gamma(Y') = \gamma(Y'') = \gamma$ for some $\gamma > 0$, so that

$$\Delta\widetilde{B}^*_T(\xi) = \begin{cases} 0, & Z_T = 1, \\ \left((\xi H_T)^{-1/\gamma} + x_0(Y') - \Delta\widehat{B}^{r,0}_T\right)^+, & Z_T = 0, \end{cases} \tag{6.4}$$

$$\Delta K^*_T(\xi) = \begin{cases} 0, & Z_T = 1, \\ \left((\xi H_T \widehat{V}^{r,0}_T)^{-1/\gamma} + x_0(Y'') - 1\right)^+, & Z_T = 0, \end{cases} \tag{6.5}$$

where $\Delta\widehat{B}^{r,0}_T$ and $\widehat{V}^{r,0}_T$ are, respectively, the lump sum paid out at time T and the reserve at time T if the insured is alive.

We consider two cases corresponding to the preference structures in Example 4.3.2, that is, in Case 1 (resp. Case 2) we measure utility of the total payments (as in (3.14)) (resp. bonus payments (as in (3.15))).

Case 1: Here, $x_0(Y') = x_0(Y'') = 0$, and

$$\Delta\widetilde{B}^*_T(\xi) = 1_{(Z_T=0)} \left[\left(\xi \frac{\Lambda_T \Gamma^0_T}{S^0_T}\right)^{-1/\gamma} - \Delta\widehat{B}^{r,0}_T\right]^+, \tag{6.6}$$

$$\Delta K^*_T(\xi) = 1_{(Z_T=0)} \left[\left(\xi \frac{\Lambda_T \Gamma^0_T \widehat{V}^{r,0}_T}{S^0_T}\right)^{-1/\gamma} - 1\right]^+. \tag{6.7}$$

Case 2: Here, $x_0(Y') = \Delta\widehat{B}^{r,0}_T$, $x_0(Y'') = 1$, and

$$\Delta\widetilde{B}^*_T(\xi) = 1_{(Z_T=0)} \left(\xi \frac{\Lambda_T \Gamma^0_T}{S^0_T}\right)^{-1/\gamma}, \tag{6.8}$$

$$\Delta K^*_T(\xi) = 1_{(Z_T=0)} \left(\xi \frac{\Lambda_T \Gamma^0_T \widehat{V}^{r,0}_T}{S^0_T}\right)^{-1/\gamma}. \tag{6.9}$$

In both cases $\Delta\widetilde{B}^*_T(\xi)$ and $\Delta K^*_T(\xi)$ are products of the indicator $1_{(Z_T=0)}$ and a purely financial (i.e., $\mathcal{F}^W_T$-measurable) random variable. Letting ψ^i denote the function ψ given by (4.9) in Case i, i = 1,2, we have

$$\psi^i(\xi) = \mathrm{Q}\left(T^{01} > T\right) \mathrm{E}^{\mathrm{Q}}\left(\frac{\Phi^i(\xi)}{S^0_T}\right), \ \xi > 0,$$

where

$$\Phi^1(\xi) = \left[\left(\xi\frac{\Lambda_T\Gamma_T^0}{S_T^0}\right)^{-1/\gamma} - \Delta\widehat{B}_T^{r,0}\right]^+ + \left[\left(\xi\frac{\Lambda_T\Gamma_T^0\widehat{V}_T^{r,0}}{S_T^0}\right)^{-1/\gamma} - 1\right]^+,$$

$$\Phi^2(\xi) = \left(\xi\frac{\Lambda_T\Gamma_T^0}{S_T^0}\right)^{-1/\gamma} + \left(\xi\frac{\Lambda_T\Gamma_T^0\widehat{V}_T^{r,0}}{S_T^0}\right)^{-1/\gamma}.$$

Note that $\Phi^1(\xi)$ and $\Phi^2(\xi)$ are $\mathcal{F}_T^W$-measurable. Assumption 4.4.2 is met if and only if

$$\mathrm{E}^{\mathrm{Q}}\left(\Phi^i(\xi)/S_T^0\right) < \infty, \ \forall\xi > 0, \tag{6.10}$$

and for this to hold it is easily seen to be sufficient (and obviously necessary) that $\mathrm{E}^{\mathrm{Q}}\left(\Phi^i(\xi)/S_T^0\right) < \infty$ for *some* $\xi > 0$. This condition also implies that Assumption 4.4.3 holds, and Assumption 4.4.4 holds due to (2.2). Therefore, the assertions in Proposition 4.4.5 and Corollary 4.4.7 hold if (6.10) holds.

In Case 1 the optimal (bonus) payment structure is formed by the product of $1_{(Z_T=0)}$ and European call options on two financial derivatives,

$$\mathcal{X}'(\xi) := (\xi\Lambda_T\Gamma_T^0/S_T^0)^{-1/\gamma},$$

and

$$\mathcal{X}''(\xi) := (\xi\Lambda_T\Gamma_T^0\widehat{V}_T^{r,0}/S_T^0)^{-1/\gamma},$$

with strike prices $\Delta\widehat{B}_T^{r,0}$ and 1, respectively. According to the analysis above the company's investment portfolio at time $t \in [0,T]$ should, if the insured is still alive, be formed as if the actual payment streams were replaced by the artificial ones given by

$$e^{-\int_t^s \mu_\tau^{01}\,d\tau}(d\widehat{B}_s^0 + \mu_s^{01}\widehat{b}_s^{01}\,ds), \ s \in (t, T'],$$

and

$$e^{-\int_t^T \mu_\tau^{01}\,d\tau}(\mathcal{X}'(\xi^*) - \Delta\widehat{B}_T^{r,0})^+\,d1_{[T,T']}(s)$$
$$+ (\mathcal{X}''(\xi^*) - 1)^+\,1_{(T,T']}(s)e^{-\int_t^s \mu_\tau^{01}\,d\tau}\,(d\widehat{B}_s^{r,0} + \mu_s^{01}\widehat{b}_s^{r,01}\,ds), \ s \in (t, T'],$$

corresponding to the contractual and the bonus payments, respectively (here we have assumed for simplicity that the Q-risk-minimizing strategy is used). The point is, of course, that these artificial payment streams are adapted to $(\mathbb{F}_s^W)_{t\le s\le T'}$ when viewed at time t. In particular, the one corresponding to the contractual payments is deterministic and should therefore, as already mentioned, be backed by zero-coupon bonds. The one corresponding to the bonus payments can be hedged by appropriate positions in the abovementioned European call options. We emphasize that the artificial payments streams themselves change constantly, of course, so the portfolios must be updated accordingly.

The only random factors in $\mathcal{X}'$ and $\mathcal{X}''$ are $(\Lambda_T/S_T^0)^{-1/\gamma}$ and $(\Lambda_T\widehat{V}_T^{r,0}/S_T^0)^{-1/\gamma}$, respectively, and explicit hedging portfolios can be obtained in many cases of interest. We shall not pursue this here, however; the interested reader is referred to the literature. Note that since the portfolio hedging a European call option typically involves a short position in zero-coupon bonds, the overall portfolio need not be dominated by bonds.

In Case 2 the optimal (bonus) payment structure is formed by the product of $1_{(Z_T=0)}$ and two financial derivatives of the exact same forms as $\mathcal{X}'$ and $\mathcal{X}''$. However, the value of ξ ensuring fairness, ξ^*, will of course be smaller than in Case 1. Also here, an investment in zero-coupon bonds is made as if the contractual payment stream were replaced by its artificial counterpart. The investment strategy corresponding to the bonus payments is determined without consideration for the contractual payments, it has the same form as if one were simply to invest the amount $-\widehat{V}_{0-}^0$ so as to maximize

$$\mathrm{E}\left(1_{(Z_T=0)}\left[u^{(\gamma)}(\Delta\widetilde{B}_T)+u^{(\gamma)}(\Delta K_T)\right]\right).$$

We end this example by applying the last assertion of Corollary 4.4.7. A few simple calculations show that as long as the insured is alive, early dividend allocation can be made during $[0,T)$ without loss of utility to the extent that

$$\mathrm{P}\left(K_t \le 1+\left[\left(\xi\Lambda_T\Gamma_T^0\widehat{V}_T^{r,0}/S_T^0\right)^{-1/\gamma}-1\right]^+ \middle| \mathcal{F}_t\right)=1,\ \text{a.s.},\ t\in[0,T),$$

in Case 1, and

$$\mathrm{P}\left(K_t \le 1+\left(\xi\Lambda_T\Gamma_T^0\widehat{V}_T^{r,0}/S_T^0\right)^{-1/\gamma} \middle| \mathcal{F}_t\right)=1,\ \text{a.s.},\ t\in[0,T),$$

in Case 2. In particular, it must be the case that the conditional essential infimum of S_T^0/Λ_T is strictly greater than 0. Early dividend allocation is thus strictly suboptimal if S_T^0/Λ_T can get arbitrarily close to 0 (given the current information), which is the typical situation.

4.7 Conclusions

We shall briefly summarize and comment on the main results of this chapter. We have shown that with a preference structure given by (3.11), dividend allocation during the accumulation phase is in general suboptimal under dividend schemes as the one studied in this chapter, where dividends are converted into future bonus benefits. The natural explanation is that allocation of dividends has no immediate value for the insured; only the bonus benefits paid out at a later stage do. And for any optimal strategy under which early dividends can be allocated with no loss of utility there is a strategy with no early dividend allocation that is just as good. Furthermore, the suboptimality is often strict because the optimal (bonus) payment structure cannot be met if dividends have been allocated prematurely. Early dividend allocation leads to more fixed investments in zero-coupon bonds and thus restricts the freedom that the company has regarding investments on behalf of the insured.

Nevertheless, early dividend allocation is often encountered in practice. There may be several practical concerns playing a role in this regard, e.g. tradition, legislation, simplicity etc. Furthermore, the fairness constraint (2.6) is not (yet!) fulfilled in practice, and it should also be admitted that our bonus conversion scheme does not work exactly as in practice. However, although some degree of relevance can be attached to these issues, they do not seem sufficient to reject our result that early dividend allocation is in general suboptimal. It should be mentioned, though, that in recent years many companies have introduced variations of the traditional dividend schemes that are more flexible and less vulnerable to criticism and adverse financial scenarios.

We have also shown how the theory of risk minimization in incomplete markets can be used to find optimal hedging strategies for the company, aiming at minimizing the squared fluctuations of the total net costs, either globally or locally. It was found that the value of the company's investment portfolio associated with the policy should always be equal to the total market reserve (or the total $\check{P}$-reserve in the case of local risk minimization). This is of course quite natural. However, the important point is that it serves as a building block in the foundation on which most papers on the investment aspect of optimal pension fund management rest, in the sense that it provides a *criterion-based* argument that one can, to some extent, ignore the individual policy risk, and thus supports the usual (implicit) diversification argument.

4.8 Discussions and generalizations

A. On the dividend and bonus schemes.
An important part of the overall problem treated in this chapter has been optimization of the bonus strategy. As explained in Section 4.3 we have used a certain method, by which the bonus payment stream as such is not determined directly, but through the dividend stream and a certain, fixed conversion procedure. As mentioned, this is similar to a traditional method that has been widely used in practice (at least until recently, where quite a few new methods have been introduced). However, there are also dissimilarities, and we therefore discuss our approach in more detail.

By the mentioned traditional method, dividends allocated to the individual policyholder are converted into (future) bonuses by a proportional increase of *all* future guaranteed benefits to the policyholder. This means that not only the guaranteed pension benefits, but also insured sums to be paid out, e.g. upon death of the insured, are increased. By the method employed in this chapter it is only what we have referred to as the *pension benefits* that are (proportionally) increased when dividends are allocated. Moreover, dividends are typically allocated throughout the entire policy term in practice, whereas we have assumed that no dividends are allocated after the time of retirement (of course, other interpretations of T than as "the time of retirement" are possible, though). However, these differences are of minor importance: Although this chapter does not contain results on optimal dividend allocation for the purpose of increasing insurance benefits to be paid out at random times before the time of retirement (such as a sum upon death), our results concerning allocation of dividends during the premium payment period for the purpose of increasing the pension benefits are certainly valid (under all other modelling assumptions, of course) and important in their own right, since the (value of the) pension benefits typically outweigh the (value of the) insurance benefits in pension schemes.

Another dissimilarity lies in the calculation of the increase of the (future) guaranteed benefits triggered by the dividend allocations. As explained in Section 4.3 the dividends dD_t allocated in a small time interval around t lead to an increase in the future benefits through an increase dK_t in the number of units of the basic future guaranteed benefit stream. By the method used in this chapter, dK_t is determined from dD_t according to the equation

$$dK_t = \frac{dD_t}{\widehat{V}_t^r}.$$

The increase in the number of units equals the added amount of dividends divided by the value of a single unit, $\widehat{V}_t^r$. The traditional method is similar, but the value of a single unit is calculated under technical, prudent assumptions (constituting the actuarial so-called *first order basis*), that is, $\widehat{V}_t^r$, which in our model is the market reserve corresponding to a single unit of the future benefits, is replaced by the technical corresponding *first order reserve*, which, by statute, is larger (the first order basis may have to be adjusted over time to ensure this). This means that for a given value of dD_t, dK_t is smaller by the traditional method than by the method we have employed.

However, this is actually just a technicality: At the end of the day it is K, rather than D, that matters, which is also reflected by our objective (4.2). In other words, what is important is the number of units, dK_t, that is added, not the way this number of units is calculated. Thus, using the traditional method would not change our results. In fact, we could have dropped the process D from the model altogether and simply worked directly with K instead.

B. Intervention options.

It was claimed in Paragraph 4.2.H that certain intervention options could be included in the problem setup without affecting the results. We shall now briefly explain how this could be done.

The intervention options in question are the *surrender* and the *free policy* options, which are, of course, options held by the policyholder (in financial terms, the policyholder has the *long* position). Both are, quite naturally, included automatically in most life insurance or pension policies. The surrender option gives the policyholder the right to cancel the policy prematurely; the company then pays out a surrender value. The free policy option gives the policyholder the right to stop paying premiums (before the end of the planned premium payment period); the policy then remains in force, but the (total) future benefit stream is adjusted.

There are no uniformly used principles governing the calculation of the surrender value and the adjustment of the future benefits connected to the surrender and free policy option, respectively. We shall not discuss the issue in detail; for a general and detailed treatment of these options we refer to Steffensen (2002) and the references therein. Here we just propose a simple (and natural) principle, namely that the abovementioned calculations and adjustments should always be based on the current market value of (the remaining payments connected to) the policy.

The total market reserve, defined by (2.8), therefore plays an important role, as it is the current market value of the total future benefits minus the contractual future premiums. If the surrender option is exercised, then the company should simply pay out an amount equal to the market reserve. If the free policy policy option is exercised, then the company should simply reduce the total future benefits so as to make the market value of the total future benefit equal to the current market reserve (there are several ways to do this in general, but we shall not discuss them here).

Remark 4.8.1 In the presence of these intervention options it is of course necessary to assume that the total market reserve is nonnegative throughout the policy term. This is a standard assumption in life insurance, both in theory and practice, but it is only necessary because of the intervention options. We have not imposed it in the (otherwise general) model of this chapter, where intervention options were not included.

It is intuitively clear that with this approach to the intervention options, the qualitative features of the overall problem of this chapter are not affected, since the financial balance between the company and the policyholder is not disturbed if one of the options is exercised. It is also a very natural approach, and it makes it meaningless for the policyholder to speculate against the company in order to exercise (one of) the options when the time is right, so to speak. If, for example, the surrender value to be paid out exceeds the market value of the future payments, then the policyholder can make a profit (in terms of market values) by exercising his surrender option (Steffensen (2002) studies this aspect in particular). However, this opportunity is not (and has never been) the purpose of the surrender option.

C. Generalization of the market model.
This paragraph contains a discussion of some possible generalizations of the financial market model. We shall focus mostly on ideas and principles and thus leave out technicalities.

Let us begin by noting that a generalization of the Brownian motion-driven market model employed in this chapter to a general semimartingale market model as in Schweizer (2001) is fairly straightforward as long as the financial market is assumed complete. The basic setup is easily generalized, Proposition 4.4.1 goes through without changes, and the optimization methods used in Section 4.4 are well established for the general semimartingale model (Kramkov and Schachermayer (1999)). As for the quadratic hedging approaches, certain regularity conditions may have to be imposed in order to ensure the existence of P-locally risk-minimizing strategies (see Schweizer (2001)), but we shall not go into details on this.

Much more complex issues arise if the assumption about completeness of the financial market is dropped. Let us briefly recapitulate the basic structural model assumptions: We assumed the financial market in itself to be complete, by taking the market price of risk process λ to be given. In particular this implied the existence of a unique equivalent martingale measure for the discounted marketed asset price processes. In contrast, the insurance policy risk was considered to be (completely) unhedgeable, but it was assumed that the company had chosen a pricing measure for the policy risk (given by the martingale Γ).

The assumption of completeness of the financial market is rather strong, in particular since we are dealing with life and pension insurance with (typically) very long-term contracts. In the following discussion we consider a relaxation of the assumption that the financial market is complete, but we keep the assumption regarding the pricing of policy risk. For simplicity we also stick to the general Brownian motion framework, which is rich enough for this discussion.

When the financial market is incomplete the set of equivalent local martingale measures is no longer a singleton (but we assume that it is nonempty). In other words, there is more than one market price of risk process λ for which the dynamics of the risky asset price processes satisfy (2.1). Therefore, the basic fairness constraint (2.6) is no longer well specified, and the issue of fairness becomes an open question. To ease the discussion we shall employ a somewhat inaccurate terminology by referring to an equivalent probability measure Q with a Radon-Nikodym derivative of the form

$$\frac{d\mathrm{Q}}{d\mathrm{P}} = \Gamma(T)\widetilde{\Lambda}(T),$$

where $\widetilde{\Lambda}$ is the density process of an equivalent local martingale measure for the financial market, i.e., an $\mathbb{F}^W$-martingale with $\mathrm{E}\left(\widetilde{\Lambda}(T)\right) = 1$, as an *equivalent local martingale measure*, although we stress that the policy risk is still considered unmarketed and therefore unhedgeable.

Now, it is clearly a necessary requirement for fairness that (2.6) must hold for at least one equivalent local martingale measure. A natural question, therefore, is whether this condition is also sufficient. In general that would depend on the (attitudes towards risk of the) company and the policyholder: If it were possible for them to come to an agreement on some bonus payment process such that (2.6) were satisfied just for *some* equivalent local martingale measure, then, of course, this could not be considered to be unfair. However, this hardly corresponds to the normal situation in practice, where the policyholder only specifies a certain stream of desired guaranteed benefits, agrees to pay a certain stream of premiums, and otherwise leaves the decisions regarding the bonus payment process more or less to the company.

We propose here, as a (sort of) fairness constraint, that the optimization should be performed over the class of bonus payment processes satisfying

$$\mathrm{E}^{\mathrm{Q}}\left(\int_{[0,T']}\frac{1}{S_t^0}\left(d\widehat{B}_t+d\widetilde{B}_t\right)\right)\leq 0,\ \forall \mathrm{Q}\in\mathbb{P}, \tag{8.1}$$

where $\mathbb{P}$ denotes the set of equivalent local martingale measures in the above sense. If the payment processes were purely financial and thus did not depend on the policy process Z, then (8.1) would be a natural budget constraint. Since the price of the policy risk (the risk associated with Z) is given, the budget constraint should simply carry over to the general case. However, we stress that (8.1) is not a general fairness constraint in itself; it is a part of the fairness constraint that the company should optimize the bonus payment process.

Apart from being fairly justifiable in itself, this constraint would also facilitate the application of the established methodology for portfolio optimization in incomplete markets (see e.g. Karatzas et al. (1991), He and Pearson (1991b), and Kramkov and Schachermayer (1999)) to the dividend/bonus optimization problem of the present chapter. We shall not go into details, but the idea is that it should be possible to formulate a generalized dual problem (as in the mentioned papers) by simply letting the domain of the dual problem be the set $\mathbb{P}$ of equivalent local martingale measures in the above sense. This, in turn, should lead to the existence and characterization of an optimal bonus/dividend strategy (at least under the mild assumption that the contractual payment processes be bounded, see Cvitanić et al. (2001)).

For a given bonus/dividend strategy (optimal or not) the company would still be left with an unhedgeable total payment process, and to handle this problem we would argue that quadratic hedging approaches would still be quite natural.

Formalizing these ideas would be an interesting topic for future research.

Let us end this discussion with the important observation that the message of Proposition 4.4.1 does not depend on the assumption of market completeness. This is not obvious, because the market reserve, which is based on the unique equivalent martingale measure in the complete-market model, is used in the conversion from dividends into future bonuses. However, as mentioned in Paragraph A above, we could work directly with K rather than D as the controlled process. This would mean that the market reserve would no longer be involved, and the proposition would hold true. Thus, whether the market is complete or not, it is always sufficient to look for an optimal dividend strategy in the class $\mathcal{D}_0^A$ of admissible dividend strategies satisfying $D_t \equiv 0$, $\forall t\in[0,T)$ or, equivalently, $K_t \equiv 1$, $\forall t \in [0,T)$.

4.9 A comparison with Steffensen (2004)

A. Outline and purpose.
As mentioned in the introduction to this chapter, Steffensen (2004) studies a related problem. In particular, he also considers optimization of the bonus payments through allocation of dividends but obtains explicit solutions in this regard only in a few special cases based on rather strong conditions on the life insurance policy and the dividend/bonus scheme. The purpose of this paragraph is to analyze Steffensen's problem with the results of this chapter in mind. More precisely, we shall see if our result concerning optimal dividend allocation, namely that dividends need not be allocated before the time of retirement under certain (relatively mild) conditions on the policy and the dividend/bonus scheme, carry over and lead to an explicit solution to Steffensen's problem under simimlar conditions.

In contrast to what one might expect, it turns out that a certain amount of dividends *should* be allocated before the time of retirement according to the model studied here, but this does not make the situation less interesting. We provide a self-contained account of the relevant parts of the general problem; to some extent it differs from the presentation in Steffensen (2004).

B. Problem setup.
The basic model is a special case of the model in this chapter. As the title of Steffensen's paper suggests, the financial market is as in Merton (1969); it is obtained in our model by setting $d = 1$ (i.e., just a single Brownian motion and a single risky asset) and letting the market coefficients r, λ, and σ, be constant. The model of the insurance policy and the contractual payments are of the same type as the one in Example 4.2.5, with the (minor) restriction that the state-wise payment functions $\widehat{B}^e$, $e \in \mathcal{Z}$, must be piece-wise absolutely continuous and have a finite number of jumps.

The combined model is thus Markovian, and as in Steffensen (2004) we take a dynamic programming approach to the optimization problem considered below. We refer to Appendix B for an introduction to this issue and for further details.

The so-called *free reserve* process $X = (X_t)_{t\in[0,T']}$, the exact definition of which follows below, plays the role of the wealth process in consideration. Initially (i.e., at time $0-$), the free reserve is, by definition, the difference between the total market reserve (cf. (2.8)) and the market reserve corresponding to the guaranteed payments (cf. (2.3)), that is,

$$X_{0-} = V_{0-} - \widehat{V}_{0-}, \tag{9.1}$$

which is nonnegative by (2.4). We assume here that it is strictly positive. The free reserve is invested in the financial market according to an adapted *stock proportion* process $w = (w_t)_{t\in[0,T']}$, i.e., w_t denotes here the *proportion* of X_t invested in the risky asset at time t, and it finances the dividends allocated to the policyholder (in a sense that is made precise below). An important principle of the model is that all guaranteed payments are assumed to be taken care of independently, i.e., they have no effect on the free reserve process.

Now, in order to define X in the case that is relevant here (Steffensen considers another case as well), we introduce a certain fixed "payment process" A, which is a non-decreasing, non-null, $\mathbb{F}^Z$-adapted process of the same form as the contractual payment process $\widehat{B}$. This process is interpreted as the policyholder's desired bonus payment profile, and allocated dividends are thus used throughout to purchase units of the (remaining part of the) payment process A (as described in more detail below). In this respect the framework is quite general as it allows for various specifications of A. In particular, in the model of this chapter we assumed in Section 4.3 that dividends were used to purchase additional units of the retirement benefit stream; this is just a special case corresponding to $A = \widehat{B}^r$. The same goes for the "traditional" method mentioned in Paragraph 4.8.A, which, with a slightly sloppy notation, corresponds to $dA = (d\widehat{B})^+$.

More precisely, we shall assume that dividends allocated at time t lead to an (additional) bonus payment process, which is proportional to

$$\left(\sum_{e \in \mathcal{Z}} 1_{(Z_t = e)} \, \Delta A_t^e + A_s - A_t \right)_{s \in [t, T']}, \tag{9.2}$$

where the notation corresponds to that of Example 4.2.5. It is important to note here that if a state-wise lump sum (which is not triggered by jumps of Z) is to be paid at time t according to A, then it is taken into account here. This approach distinguishes itself from the one taken in Steffensen (2004), where only the strictly future bonus payments are affected by dividend allocations. See Remark 4.9.1 below for a motivation of our approach, which we believe is more natural.

Corresponding to the process A we define the quantities

$$\begin{aligned}
\widehat{V}_t^A &= \mathrm{E}\left(\int_{(t,T']} e^{-r(t-s)} \, dA_s \,\middle|\, \mathcal{F}_t \right) + \sum_{e \in \mathcal{Z}} 1_{(Z_t = e)} \, \Delta A_t^e, \\
V_t^{*A} &= \mathrm{E}^{\mathrm{P}^*}\left(\int_{(t,T']} e^{-r^*(t-s)} \, dA_s \,\middle|\, \mathcal{F}_t \right) + \sum_{e \in \mathcal{Z}} 1_{(Z_t = e)} \, \Delta A_t^e,
\end{aligned}$$

for $t \in [0, T']$. Here, $r^* \in \mathbb{R}$ is a technical interest rate, and P^* is a technical probability measure equivalent to P; in combination they constitute the actuarial first order basis.

The first terms (i.e., the conditional expectations) in the definitions of $\widehat{V}_t^A$ and V_t^{*A} are the market reserve and the *first order reserve*, respectively, at time t, corresponding to (the "strictly future" remaining part of) A. Corresponding to the way in which bonus payments are added upon dividend allocation (explained below), we have added state-wise lump sums paid at time t, so $\widehat{V}_t^A$ and V_t^{*A} can be interpreted as reserves for the payment process (9.2). Note that due to the Markov assumption on Z and the structure of A one can replace $\mathcal{F}_t$ with Z_t.

Now, X is defined as the solution to the SDE

$$dX_t = (r + \lambda\sigma w_t)X_t\,dt + \sigma w_t X_t\,dW_t - 1_{(t<\tau)}\frac{\widehat{V}_t^A}{V_t^{*A}}\,dD_t - 1_{(t\geq\tau)}dD_t, \tag{9.3}$$

(with the initial condition (9.1)), where τ is the stopping time

$$\tau = \inf\{t \in [0, T'] : V_t^{*A} = 0\} \wedge T'.$$

According to (9.3), upon allocation of dividends of the amount dD_t before τ, the term

$$\frac{\widehat{V}_t^A}{V_t^{*A}}\,dD_t \tag{9.4}$$

is subtracted from the free reserve. The factor $\widehat{V}_t^A/V_t^{*A}$ appears because we work here (as in Steffensen (2004)) with the "traditional" calculation method described in Paragraph 4.8.A above: The (additional) number of units of A obtained with the amount dD_t is thus given by $dK_t = dD_t/V_t^{*A}$. The net (market) cost of allocating dD_t is given by $\widehat{V}_t^A\,dK_t$, which equals (9.4). Once the dividends dD_t have been allocated, the corresponding bonus payments, which are given by dK_t units of the process defined in (9.2), are considered as guaranteed and therefore, in keeping with the abovementioned principle, assumed to be taken care of independently. This is why a subtraction of the term (9.4) from the free reserve is made although nothing is actually paid out.

Remark 4.9.1 As explained above, our approach differs from the one taken by Steffensen (2004), because he assumes that allocated dividends only affect the strictly future payments. We believe our approach to be more natural, because it makes it possible to allocate dividends so as to increase an anticipated state-wise lump sum bonus payment to be made at time t, say, without increasing the bonus payments paid out continuously immediately before t, by simply putting $dD_t > 0$. This is particularly relevant for the terminal lump sum at time T'.

For $t > \tau$ we have $dA_t \equiv 0$, i.e., the desired bonus payment profile contains no further payments. This means that conversion of dividends into bonuses as specified by A becomes meaningless. At this point we therefore assume that dividends allocated after τ are simply paid out directly.

As in Steffensen (2004), we shall consider a stock proportion/dividend-pair (w, D) to be *admissible* if the corresponding free reserve process X is well defined (by the SDE given by (9.1) and (9.3), of course) and nonnegative. We denote by $\mathcal{C}$ the set of all such pairs.

Obviously, by the nonnegativity constraint on X, a necessary requirement for admissibility of (w, D) is that

$$1_{(t<\tau)}\frac{\widehat{V}_t^A}{V_t^{*A}}\,dD_t + 1_{(t\geq\tau)}dD_t \leq X_{t-}, \tag{9.5}$$

in particular $dD_t = 0$ if $X_{t-} = 0$.

Remark 4.9.2 What is less obvious is that the constraint (9.5) is in fact also sufficient to ensure that X is well defined and nonnegative under the very mild assumption that

$$\inf\left\{t \in [0, T'] : \int_0^t w^2(s)\,ds = \infty\right\} > 0, \text{ a.s.} \tag{9.6}$$

We refer to Karatzas and Shreve (1998), Section 6.2, for elaborations on this (note, however, that they actually do not even assume (9.6), at least not explicitly, but it is a necessary condition). With this observation, the set $\mathcal{C}$ is characterized explicitly rather than implicitly.

To ensure *fairness*, Steffensen imposes the requirement that the free reserve must eventually be "emptied", i.e., D must be such that

$$X_{T'} = 0. \tag{9.7}$$

To enforce this requirement we shall for simplicity assume that the entire free reserve is simply paid out at time τ, i.e., that

$$dD_\tau = X_{\tau-},$$

so that $X_t \equiv 0$, $\forall t \in [\tau, T']$. In particular, no dividends are paid out (strictly) after τ.

It can be shown that the requirement (9.7) makes the total payment process comply with our fairness constraint (2.6) (or, equivalently, (2.7)). However, it is a stronger requirement than the one imposed by (2.6) because (9.7) makes dividend redistribution across different policies inadmissible. The bonus payments resulting from the conversion of dividends into bonuses may lead to redistributions, though.

Remark 4.9.3 It is interesting to note that the relation (3.9) involving the dividend process does *not* hold in the present model. This is due to the abovementioned fact that the calculation method employed when converting dividends into bonuses is based on the technical "reserve" V^{*A} rather than the market "reserve" $\widehat{V}^A$.

C. Objective.
Now, let $u : [0, \infty) \to [-\infty, \infty)$ be a CRRA utility function, i.e.,

$$u(x) = x^\gamma/\gamma, \; x > 0,$$

and $u(0) = \lim_{x \searrow 0} u(x)$, for some $\gamma < 1$ (for $\gamma = 0$, $u(x) = \log(x)$, $x > 0$). With a denoting the continuously-paid rate (defined between jumps of A), and a^{ef}, $e, f \in \mathcal{Z}$, $e \neq f$, denoting the (deterministic) lump sum payments upon jumps of Z, we consider the objective

$$\sup_{(w,D)\in\mathcal{C}} \mathrm{E}\left(U^c_{[0,T']} + U^j_{[0,T']}\right),$$

where, for $t \in [0,T']$,

$$
\begin{aligned}
U^c_{[t,T']} &= \int_{[t,T']} 1_{(a_s>0)} u(K_s a_s)\, ds, \\
U^j_{[t,T']} &= \int_{(t,T']} \sum_{e,f\in\mathcal{Z}} 1_{(a_s^{ef}>0)} u\left(K_{s-} a_s^{ef}\right) dN_s^{ef} \\
&\quad + \sum_{s\in[t,T'],\, e\in\mathcal{Z}} 1_{(Z_s=e)} 1_{(\Delta A_s^e>0)} u(K_s \Delta A_s^e).
\end{aligned}
$$

The terms $U^c_{[0,T']}$ and $U^j_{[0,T']}$ measure the total utility of the continuous bonus payments and the lump sum bonus payments, respectively. The indicators $1_{(a_s>0)}$, $1_{(a_s^{ef}>0)}$, and $1_{(\Delta A_s^e>0)}$ appear in order to ensure that the objective is well posed for $\gamma \le 0$, where $u(0) = -\infty$.

Remark 4.9.4 It is tacitly assumed that the expectation in the objective is well defined for all admissible (w, D)-pairs. Otherwise the admissibility condition should be strenghtened so as to ensure this to be the case for admissible pairs. It is also assumed that the value function is finite. We do not elaborate on these assumptions.

We define correspondingly the state-wise value functions $\Phi^e : [0,T'] \times [0,\infty)^2 \to \mathbb{R}$ by

$$
\Phi^e(t,x,k) = \sup_{(w,D)\in\mathcal{C}} \mathrm{E}_{(t,e,x,k)} \left(U^c_{[t,T']} + U^j_{[t,T']} \right),\ (t,x,k) \in [0,T'] \times [0,\infty)^2,
$$

where $\mathrm{E}_{(t,e,x,k)}(\cdot) = \mathrm{E}(\cdot \mid Z_t = e, X_{t-} = x, K_{t-} = k)$, and, with a slight abuse of notation, $\mathcal{C}$ is now the set of admissible strategies in respect of the initial state (e,t,x,k).

It is important to note that the utility of a lump sum state-wise payment made at time t, ΔA_t^e, is included in $U^j_{[t,T']}$. This is due to our setup, by which it is possible to increase a state-wise lump sum payment at time t by making a lump sum dividend allocation at time t. This also has the consequence that we can expect Φ^e to be left-continuous in t at the fixed jump times of A^e.

D. Variational inequalities.
As in the problem considered in Chapter 3, the HJB equation of dynamic programming becomes a system of variational inequalities for the state-wise value functions, which are here given by (with the generic argument (t,x,k) skipped for notational convenience)

$$
0 \ \ge\ \frac{\partial \phi^e}{\partial k} - \frac{\partial \phi^e}{\partial x} \widehat{V}_t^{A,e}, \tag{9.8}
$$

$$
\begin{aligned}
0 \ \ge\ & 1_{(a_t^e>0)} u(k a_t^e) + \frac{\partial \phi^e}{\partial t} + \sup_{w\in\mathbb{R}} \left(\frac{\partial \phi^e}{\partial x} (r + \lambda\sigma w) x + \frac{1}{2} \frac{\partial^2 \phi^e}{\partial x^2} \sigma^2 w^2 x^2 \right) \\
& + \sum_{f\in\mathcal{Z}\setminus\{e\}} \lambda_t^{ef} R^{ef},
\end{aligned} \tag{9.9}
$$

for $e \in \mathcal{Z}$, where the λ^{ef} denote the transition intensity, the R^{ef} are given by

$$R^{ef}(t,x,k) = \phi^f(t,x,k) - \phi^e(t,x,k) + 1_{(a_t^{ef}>0)} u(a_t^{ef}k),$$

and the notation is otherwise obvious. For each $e \in \mathcal{Z}$, one of the above inequalities must hold with equality. We have used the notation ϕ rather than Φ in order to make clear the value function Φ need *not* be a solution to this system.

We refer to Steffensen (2004) for a heuristic derivation of this system and an abstract characterization of the (expected) form of the optimal dividend strategy. The boundary conditions are given by

$$\phi^e(T',x,k) = 1_{(\Delta A^e_{T'}>0)} u(\Delta A^e_{T'}k + x)),\ (x,k) \in [0,\infty)^2. \tag{9.10}$$

Moreover, at each $(e,t) \in \mathcal{Z} \times [0,T)$ for which $\Delta A^e_t > 0$, the value function is only left-continuous in t. This means that we also have side conditions, which can be seen to be given (implicitly) by

$$\phi^e(t,x,k) = \sup_{\Delta k \in [0, x\widehat{V}_t^{A,e}]} \Big[u((k+\Delta k)\Delta A^e_t) + \phi^e(t+, x - \Delta k/\widehat{V}_t^{A,e}, k + \Delta k) \Big].$$

E. A special case.
Steffensen (2004) obtains explicit solutions to the problem in the special two-state cases of a life annuity and a term insurance on an infinite time horizon under the assertion of constant mortality intensity.

We shall here consider the general case of a multi-state policy with arbitrary contractual payment processes and transition intensities (adhering to the imposed regularity assumptions, of course). However, as in the problem otherwise considered in this chapter we assume that all bonus payments are paid out in the time interval $[T,T']$ for some fixed $T \in (0,T')$. In the present framework this amounts to the assumption that $A_{T-} \equiv 0$. Of course we think of the case where $A = \widehat{B}^r$. Moreover, we impose again the restriction that all dividends must be allocated by time T, which in the present setting means that we must have $X_T = 0$, almost surely. Consequently we have

$$K_t \equiv K_T,\ \forall t > T. \tag{9.11}$$

This changes the setup slightly, because all control terminates at time T, so we must replace T' by T in the boundary condition.

Under these restrictions the state-wise value functions become

$$\Phi^e(t,x,k) = \sup_{(w,D)\in\mathcal{C}} \mathrm{E}_{(t,e,x,k)} \left(U^c_{[T,T']} + U^j_{[T,T']} \right),\ (t,x,k) \in [0,T] \times [0,\infty)^2.$$

Remark 4.9.5 The remaining analysis below is carried under the assumption that $\gamma \neq 0$. The analysis for $\gamma = 0$ is similar.

By inserting the expression for u and using (9.11) we get

$$\Phi^e(t,x,k) = \sup_{(w,D)\in\mathcal{C}} \mathrm{E}_{(t,e,x,k)}\left(K_T^\gamma/\gamma\left(\widetilde{U}^c_{[T,T']} + \widetilde{U}^j_{[T,T']}\right)\right),$$

where

$$\begin{aligned}
\widetilde{U}^c_{[T,T']} &= \int_{[T,T']} 1_{(a_s>0)} a_s^\gamma \, ds, \\
\widetilde{U}^j_{[T,T']} &= \int_{(T,T']} \sum_{e,f\in\mathcal{Z}} 1_{(a_s^{ef}>0)} \left(a_s^{ef}\right)^\gamma dN_s^{ef} + \sum_{s\in[T,T'],\, e\in\mathcal{Z}} 1_{(Z_s=e)} 1_{(\Delta A_s^e>0)} (\Delta A_s^e)^\gamma.
\end{aligned}$$

Further, using the rule $\mathrm{E}(\cdot) = \mathrm{E}(\mathrm{E}(\cdot \,|\, \mathcal{F}_T))$ we get

$$\Phi^e(t,x,k) = \sup_{(w,D)\in\mathcal{C}} \mathrm{E}_{(t,e,x,k)}\left(K_T^\gamma/\gamma \sum_{f\in\mathcal{Z}} 1_{(Z_T=f)} J^f\right),$$

where

$$J^f = \mathrm{E}\left(\widetilde{U}^c_{[T,T']} + \widetilde{U}^j_{[T,T']} \,\middle|\, Z_T = f\right), \; f\in\mathcal{Z}.$$

The value function has been simplified considerably, but it is still quite difficult to make a qualified guess at its form. To ease things further we therefore assume that there is a subset $\mathcal{Z}' \subseteq \mathcal{Z}$ such that for $e\in\mathcal{Z}'$, $\widehat{V}_T^{A,e}$ and J^e do not depend on the state e, i.e., for $e\in\mathcal{Z}'$,

$$\begin{aligned}
\widehat{V}_T^{A,e} &= \widehat{V}_T^{A,\mathcal{Z}'}, \\
J^e &= J^{\mathcal{Z}'},
\end{aligned}$$

and that for $e\in\mathcal{Z}\setminus\mathcal{Z}'$, $\widehat{V}_T^{A,e} = 0$ (of course we assume $\widehat{V}_T^{A,\mathcal{Z}'} > 0$). This assumption is perhaps not as strong as it may seem: No loss of generality is implied in the single-state (purely financial) and classical two-state policy models. Moreover, in the three-state disability model (illustrated in Steffensen (2004)) the contractual payments are often constructed in such a way that the retirement benefits only depend on whether the insured is alive or not (disability may simply imply a reduction of the premiums), and if the same goes for the mortality intensity (often assumed in practice), then the imposed assumption is met.

F. Solution proposal and analysis.
We shall now see if the result on dividend allocation from the main model of this chapter (as mentioned in Paragraph A above) carries over to the present setting. If it were optimal to hold back dividends until T and then simply convert the obtained wealth X_{T-} into bonus payments at time T (provided that $Z_T\in\mathcal{Z}'$), then the value function would be given by

$$\Phi^e(t,x,k) = J^{\mathcal{Z}'}\mathrm{P}\left(Z_T\in\mathcal{Z}' \,\middle|\, Z_t = e\right)\psi(t,x,k),\; (t,x,k)\in[0,\infty)^2, \tag{9.12}$$

where

$$\begin{aligned}\psi(t,x,k) &= \sup_{(w,D^0)\in\mathcal{C}} \mathrm{E}_{(t,x,k)}\left(u\left(k+\Delta K_T\right)\right)\\ &= \sup_{(w,D^0)\in\mathcal{C}} \mathrm{E}_{(t,x)}\left(\left(k+X_{T-}/\widehat{V}_T^{A,\mathcal{Z}'}\right)^{\gamma}/\gamma\right)\\ &= \left(\widehat{V}_T^{A,\mathcal{Z}'}\right)^{-\gamma} \sup_{(w,D^0)\in\mathcal{C}} \mathrm{E}_{(t,x)}\left(\left(k\widehat{V}_T^{A,\mathcal{Z}'}+X_{T-}\right)^{\gamma}/\gamma\right). \qquad (9.13)\end{aligned}$$

Here, D^0 denotes the fixed outlined dividend strategy (exercised from time t), i.e., $dD_s^0 \equiv 0$ for $t \leq s < T$, and $dD_T = \Delta D_T = X_{T-}V_T^{*A,\mathcal{Z}'}/\widehat{V}_T^{A,\mathcal{Z}'}$.

Now, the idea is to simply check whether the function defined by the expression on the right-hand side of (9.12) satisfies the HJB variational inequalities, and to this end we shall derive a semi-explicit expression for ψ. We first note that ψ is the value function corresponding to another problem, namely a pure investment problem, in which k just appears as a parameter in the utility function (this explains why k dropped out of the subscript of the expectation operator above). This problem can be solved for each starting point (t,x) by use of the martingale method. We directly apply Karatzas and Shreve (1998), Theorem 3.7.6, from which we have that the optimal attainable terminal wealth, starting from the initial point (t,x), is given by

$$X_{T-} = \xi^{(t,x,k)} := \left(h(t,x,k)\left(e^{-r(T-t)}\Lambda_T/\Lambda_t\right)^{1/(\gamma-1)} - kV_T^{*A,\mathcal{Z}'}\right)^{+},$$

where $\Lambda_s = \exp\left(-\lambda W_s - \lambda^2 s/2\right)$, $s \in [0,T]$, and $h(t,x,k)$ is determined in such a way that $\xi^{(t,x,k)}$ has the market value x at time t. Now, $\xi^{(t,x,k)}$ has the form of a European call option with strike price $kV_T^{*A,\mathcal{Z}'}$ written on the contingent claim

$$Y_T := h(t,x,k)\left(e^{-r(T-t)}\Lambda_T/\Lambda_t\right)^{1/(\gamma-1)}.$$

The price process of this claim has the dynamics

$$dY_s = Y_s\left(r+\lambda^2/(1-\gamma)\right)ds + Y_s\,\lambda/(1-\gamma)\,dW_s,\ t<s\leq T, \qquad (9.14)$$

starting at $Y_t = h(t,x,k)m(t)$, where

$$m(t) = \exp\left(\frac{\gamma r(T-t)}{1-\gamma} + \frac{\gamma\lambda^2(T-t)}{2(1-\gamma)^2}\right).$$

It is seen from (9.14) that, in combination with the bank account, this claim forms another Black-Scholes market, so we must have

$$F\left(t, h(t,x,k)m(t), k\widehat{V}_T^{A,\mathcal{Z}'}\right) = x,$$

where $F(t,s,q)$ denotes the (well-known) Black-Scholes price at time t of a European call option with strike price q on a stock with price s (at time t) and volatility coefficient $\lambda/(1-\gamma)$, given by

$$F(t,s,q) = sN(d_1(t,s/q)) - e^{-r(T-t)}qN(d_2(t,s/q)),\ (t,s,q) \in [0,T) \times (0,\infty)^2,$$

where N is the standard normal distribution function, and

$$\begin{aligned} d_1(t,y) &= \frac{1-\gamma}{\lambda\sqrt{T-t}}\left(\log y + \left(r + \frac{\lambda^2}{2(1-\gamma)^2}\right)(T-t)\right), \\ d_2(t,y) &= d_1 - \lambda/(1-\gamma)\sqrt{T-t}. \end{aligned}$$

Since we may have $k=0$, we extend the definition of F to $[0,T) \times (0,\infty) \times [0,\infty)$ in the natural way by setting $F(t,s,0) = s,\ (t,s) \in [0,T] \times (0,\infty)$. Then, for each fixed $(t,q) \in [0,T) \times [0,\infty)$ the function $s \mapsto F(t,s,q)$, $s \in (0,\infty)$, is strictly increasing and maps $(0,\infty)$ onto itself, so it has an inverse, which also maps $(0,\infty)$ onto itself. With a somewhat sloppy notation we may therefore define the function F^{-1} as the unique function satisfying

$$F(t, F^{-1}(t,x,q), q) = x,\ (t,x,q) \in [0,T) \times (0,\infty) \times [0,\infty).$$

We can interpret $F^{-1}(t,s,q)$ as the price of the stock at time t corresponding to the price x of a European call option with strike price q. We then have

$$h(t,x,k) = F^{-1}\left(t,x,k\widehat{V}_T^{A,\mathcal{Z}'}\right)/m(t),\ (t,x,k) \in [0,T) \times (0,\infty) \times [0,\infty).$$

Now, inserting $\xi^{(t,x,k)}$ on the right-hand side of (9.13) yields

$$\begin{aligned} \psi(t,x,k) &= \left(\widehat{V}_T^{A,\mathcal{Z}'}\right)^{-\gamma} \mathrm{E}_{(t,x)}\left(\left(k\widehat{V}_T^{A,\mathcal{Z}'} + \xi^{(t,x,k)}\right)^{\gamma}/\gamma\right) \\ &= \frac{1}{\gamma}\left(h(t,x,k)/\widehat{V}_T^{A,\mathcal{Z}'}\right)^{\gamma} \mathrm{E}_{(t,x)}\left(1_G\left(e^{-r(T-t)}\Lambda_T/\Lambda_t\right)^{\gamma/(\gamma-1)}\right) \\ &\quad + \frac{k^\gamma}{\gamma}\mathrm{P}_{(t,x)}\left(G^C\right), \end{aligned}$$

where

$$G = \left(h(t,x,k)\left(e^{-r(T-t)}\Lambda_T/\Lambda_t\right)^{1/(\gamma-1)} > k\widehat{V}_T^{A,\mathcal{Z}'}\right).$$

Some straightforward calculations yield

$$\begin{aligned} \psi(t,x,k) &= \frac{1}{\gamma}\left(h(t,x,k)/\widehat{V}_T^{A,\mathcal{Z}'}\right)^{\gamma} m(t)N\left(\left[r + \frac{\lambda^2}{2(1-\gamma)}\right]\frac{\sqrt{T-t}}{|\lambda|} - d(t,x,k)\right) \\ &\quad + \frac{k^\gamma}{\gamma}N(d(t,x,k)), \end{aligned}$$

where

$$d(t,x,k) = \frac{1-\gamma}{|\lambda|\sqrt{T-t}}\left(\log k + \log \widehat{V}_T^{A,\mathcal{Z}'} - \log h(t,x,k)\right).$$

For $k = 0$ all expressions are understood in the limiting sense; in particular, $d(t,x,0) = -\infty$, and

$$\psi(t,x,0) = \frac{1}{\gamma}\left(\frac{x}{m(t)\widehat{V}_T^{A,\mathcal{Z}'}}\right)^{\gamma} m(t).$$

Now, by partial differentiation with respect to x and k we find that for any $(t,x) \in [0,t) \times (0,\infty)$,

$$\frac{\partial \psi}{\partial x}(t,x,k) = A(t,x)\frac{m(t)}{\widehat{V}_T^{A,\mathcal{Z}'}}\frac{\partial h}{\partial x}(t,x,k) + o(k),$$
$$\frac{\partial \psi}{\partial k}(t,x,k) = A(t,x)\frac{m(t)}{\widehat{V}_T^{A,\mathcal{Z}'}}\frac{\partial h}{\partial k}(t,x,k) + o(k),$$

for $k \searrow 0$, where

$$A(t,x) = \left(x/\left(m(t)\widehat{V}_T^{A,\mathcal{Z}'}\right)\right)^{\gamma-1}.$$

By using the fact that

$$F(t,s,q) = qh(t,s/q),\ \forall(t,s,q) \in [0,T) \times (0,\infty)^2,$$

where

$$h(t,y) = yN(d_1(t,y)) - e^{-r(T-t)}N(d_2(t,y)),\ (t,y) \in [0,T) \times (0,\infty),$$

it is seen that

$$F^{-1}(t,x,q) = qh^{-1}(t,x/q),\ (t,x,q) \in [0,T) \times (0,\infty)^2,$$

where $h^{-1}(t,\cdot)$ denotes the inverse of $h(t,\cdot)$. With this expression for F^{-1} it can be shown that

$$\begin{aligned}
\lim_{k\searrow 0}\frac{\partial h}{\partial x}(t,x,k) &= 1/m(t),\\
\lim_{k\searrow 0}\frac{\partial h}{\partial k}(t,x,k) &= e^{-r(T-t)}\widehat{V}_T^{A,\mathcal{Z}'}/m(t),
\end{aligned}$$

so we get

$$\begin{aligned}
\lim_{k\searrow 0}\left(\frac{\partial \psi}{\partial k}(t,x,k) - \frac{\partial \psi}{\partial x}(t,x,k)\widehat{V}_t^{A,e}\right) &= A(t,x)\left(e^{-r(T-t)} - \frac{\widehat{V}_t^{A,e}}{\widehat{V}_T^{A,\mathcal{Z}'}}\right)\\
&= A(t,x)e^{-r(T-t)}\left(1 - \mathrm{Q}_{(t,e)}\left(Z_T \in \mathcal{Z}'\right)\right)\\
&\geq 0. \qquad (9.15)
\end{aligned}$$

Here we have used the identity

$$\widehat{V}_t^{A,e} = e^{-r(T-t)} \mathrm{Q}_{(t,e)} \left(Z_T \in \mathcal{Z}' \right) \widehat{V}_T^{A,\mathcal{Z}'}.$$

G. Conclusion and discussion.
The inequality (9.15) is strict whenever $\mathrm{Q}_{(t,e)} \left(Z_T \in \mathcal{Z}' \right) < 1$, and in this case we thus have that if K_t is sufficiently small, then the inequality (9.8) is not satisfied. Therefore, we conclude that ϕ is *not* the value function for our original problem, so it *is* optimal to allocate dividends when K_t is sufficiently small (and $\mathrm{Q}_{(t,e)} \left(Z_T \in \mathcal{Z}' \right) < 1$), even though it leads to a reduction of the free reserve, and the dividends, once they have been allocated, cannot be reclaimed. In particular, since $K_0 = 0$ (in the present setting), it is in fact optimal to allocate a certain lump sum dividend at time 0.

Since allocation of dividends corresponds to buying and holding units of a fixed, nonnegative payment process, which is adapted to $\mathbb{F}^Z$ and thus independent of the financial market, this result might seem to contradict the well-known solution of the purely financial problem of maximizing the expected utility of the terminal wealth for a CRRA investor, which is characterized by a constant stock proportion and thus leads to a terminal wealth that is not bounded away from 0. However, in the purely financial problem, which in terms of the policy state process Z is characterized by $Z_t \equiv 0$, $\forall t \in [0, T']$, we have $\mathrm{Q}_{(t,e)} \left(Z_T \in \mathcal{Z}' \right) = 1$ (since, necessarily, $0 \in \mathcal{Z}'$), and (9.15) thus holds with equality. Thus, the inequality (9.8) holds (with equality) (and it can easily be verified that (9.9) also holds).

Still, it is somewhat counterintuitive that dividend allocation is optimal when K_t is small. Moreover, we note that the effect dies out when t is close to T; that is, for $t \to T$,

$$A(t,x) e^{-r(T-t)} \left(1 - \mathrm{Q}_{(t,e)} \left(Z_T \in \mathcal{Z}' \right) \right) \to 0.$$

This indicates that if t is close to T, then K_t should be smaller than if t is far away from T for dividend allocation to be optimal (i.e., the boundary between the "no-action region" and the "jump region" in the state space is decreasing in t (in the k-dimension)). In other words, dividend allocation may be optimal early on in the policy term, but if no dividends are allocated early on, then the optimal amount to allocate becomes smaller as the time passes.

What makes dividend allocation optimal when K_t is small in this problem is exactly the fact that the factor $\mathrm{Q}_{(t,e)} \left(Z_T \in \mathcal{Z}' \right)$, which is strictly smaller than 1, appears in the price paid at time t for a unit of the payment process A and thus makes it strictly smaller than the corresponding price of a fixed payment in the purely financial problem. Since dividend redistribution across policies is not allowed, the policyholder can only obtain non-financial gains by turning current wealth into units of the bonus payment process A. The above-mentioned time-dependence is due to the fact that $\mathrm{Q}_{(t,e)} \left(Z_T \in \mathcal{Z}' \right)$ is increasing in t, i.e., the later the dividends are allocated, the less is the non-financial gain.

We end this discussion with a comment on the seemingly contradictory results obtained in this section and otherwise in the chapter. There is no contradiction.

The difference in the results is due to the difference of the models. In the main model considered in this chapter there is no free reserve process, which is restricted to stay nonnegative. Therefore, although a similar dividend/bonus conversion procedure is employed, and the "price" at time t of a unit of the bonus payment process is therefore significantly lower early on in the policy term, there is no need to allocate dividends early because the dividend allocations that can be made later on are not bounded from above.

Chapter 5

On Optimal Long-term Investment in a Market Without Long-term Bonds

We study an optimal investment problem for a long-term investor in an incomplete financial market model, where the interest rate is stochastic, but where the investor's time horizon exceeds the maturities of all the initially available interest rate derivatives (e.g. bonds). New derivatives are issued along the way, but at initial prices that are affected by unhedgeable randomness. In the special case of an extended Vasicek term structure model we obtain an optimal investment strategy in explicit form, which corresponds to the natural generalization of the optimal strategy in a certain reduced complete market. We demonstrate that a similar result does not hold in an extended Cox-Ingersoll-Ross model, but we show how an optimal strategy can be obtained.

5.1 Introduction

The theory of optimization of investment (and consumption) strategies in continuous time dates back to the seminal papers by Merton (1969, 1971), where explicit results were obtained for the fairly broad class of HARA (hyperbolic absolute risk aversion) utility functions under the assumption that the (short) interest rate is constant (or deterministic). Merton (1973) extended the analysis to allow for stochastic interest rates, but obtained only semi-explicit results.

Merton's results were based on the methodology of dynamic programming, which relies in particular on the assumption of Markovian dynamics of all factors affecting the price processes of the available assets. Based on the connection between absence of arbitrage and existence of equivalent martingale measures established in Harrison and Kreps (1979) and Harrison and Pliska (1981, 1983), and the theory of convex analysis, another, arguably more powerful, approach to investment (and consumption) optimization, which allows for quite general (in particular non-

Markovian) price processes, was proposed and studied by Pliska (1986), Karatzas et al. (1987), and Cox and Huang (1989, 1991) under the assumption of complete markets. This approach, sometimes referred to as the *martingale approach*, was developed further and in particular generalized to the case of incomplete markets by Karatzas et al. (1991) and He and Pearson (1991a,b). However, although the results of these papers are quite strong in some respects, they only characterize the optimal portfolio processes in a semi-explicit fashion in general (involving the abstract integrand processes in certain martingale representations).

The theory of investment optimization problems with specific stochastic interest rate models has been developed quite recently. Sørensen (1999), Brennan and Xia (2000), Bouiler et al. (2001), Korn and Kraft (2001), Jensen and Sørensen (2001), and Bajeux-Besnainou et al. (2003) study pure investment problems under the assumption of a (nondegenerate) Gaussian term structure of interest (i.e., a Vasicek term structure (Vasicek (1977)) or some generalization of it), and Deelstra et al. (2000) assume a Cox-Ingersoll-Ross term structure (Cox et al. (1985)). Deelstra et al. (2003) generalize to a term structure that includes the Vasicek as well as the Cox-Ingersoll-Ross term structures as special cases.

A common feature of the models in these papers is that the markets are assumed to be complete. In particular, it is implied that zero-coupon bonds of any maturity are available in the market, and, moreover, these play an important role for the solutions. Canestrelli and Pontini (2000) study an incomplete market model with Vasicek dynamics of the short rate and otherwise risky price processes, none of which is assumed to have relative returns that are perfectly (negatively) correlated with the movements of the short rate. This implies, in contrast, that no zero-coupon bonds (of any maturity) are available.

For a long-term investor, such as a pension saver, the assumption of completeness over the full time horizon may be somewhat unrealistic, since the time horizon may exceed the maturity of the longest available bonds (or other interest rate derivatives) in real-world markets. On the other hand, the implied assumption of Canestrelli and Pontini (2000), that no zero-coupon bonds are available, is quite restrictive.

The purpose of this paper is to formulate and solve an optimal investment problem for a long-term investor who can invest in a market where the interest rate is stochastic, but where only short-term bonds are available. More precisely, it is assumed that the longest available interest rate derivatives (such as bonds) mature before the investor's time horizon. New interest rate derivatives (bonds) are issued along the way, however, but it is assumed that the initial prices of the new bonds are affected by unobservable (and thus unhedgeable) randomness.

The modelling of financial markets in which new interest rate derivatives are introduced as time passes (as is the case in practice) is an issue that has received very little attention in the literature. Sommer (1997) proposes a model in which the new derivatives are issued in a continuous fashion and studies the issues of pricing and hedging of contingent claims. Dahl (2005) proposes discrete-time and continuous-time models, where new derivatives are issued at fixed time points, and studies quadratic hedging strategies for (unattainable) contingent claims. Our model is a special case of the continuous-time model of Dahl (2005). To the author's knowledge, optimal investment strategies under the described market conditions have not been studied previously in the literature.

We consider the optimization problem of an investor who seeks to maximize expected utility of terminal wealth. In the general case we allow for constraints on the terminal wealth, that is, we allow for the situation where the terminal wealth must exceed some strictly positive value. Such constraints appear particularly in the investment problem of a pension fund that has issued long-term minimum benefit guarantees to its policyholders, as is the case with, e.g., most *participating* or *with-profit* pension schemes. Investment problems with terminal wealth constraints, also referred to as problems with *portfolio insurance*, are well understood as long as the financial market is complete, see, e.g. Basak (1995), Grossman and Zhou (1996), Korn (1997), Bouiler et al. (2001), and their references. However, the incomplete-market case is much more complicated, and to the author's knowledge, no explicit results exist in the literature in non-trivial cases.

The remaining part of the chapter is organized as follows. We put up the general model of an incomplete financial market in Section 5.2 and formulate the investor's optimization problem in Section 5.3. For convenience in the subsequent sections, and for completeness of the exposition, we treat the problem in a reduced, complete market in Section 5.4 in some detail. We address the incomplete-market problem in special cases in Sections 5.5 (an extended Vasicek model) and 5.6 (an extended Cox-Ingersoll-Ross model). However, only the unconstrained problems will be treated in detail.

5.2 The general market

A. Financial market model.
We consider a financial market in which an investor can trade continuously during $[0, T]$, where $T > 0$ is the investor's time horizon. We shall work within the general framework of Karatzas and Shreve (1998), Chapter 3, and make use of their results.

Let $(W_r, W_s, W_m)' = (W_r(t), W_s(t), W_m(t))'_{t\in[0,T]}$ be a standard Brownian motion in $\mathbb{R}^3$ defined on a probability space $(\Omega, \mathcal{F}, \mathrm{P})$. The following notation will be employed throughout: For any stochastic process $X = (X(t))_{t\in[0,T]}$ defined on $(\Omega, \mathcal{F}, \mathrm{P})$ we denote by $\mathbb{F}^X = (\mathcal{F}^X(t))_{t\in[0,T]}$ the augmentation (by the null-sets of $\mathcal{F}$) of the filtration generated by X.

The market offers three basic assets with strictly positive price (per-share) processes S_0, S_1, and S_2, with fixed initial values. The dynamics of the price processes are given by the SDE's

$$\begin{aligned}
dS_0(t)/S_0(t) &= r(t)\,dt, \\
dS_1(t)/S_1(t) &= (r(t) + \lambda_r(t)\sigma_{1r}(t))\ dt + \sigma_{1r}(t)\,dW_r(t) \\
dS_2(t)/S_2(t) &= (r(t) + \lambda_r(t)\sigma_{2r}(t) + \lambda_s\sigma_{2s})\ dt + \sigma_{2r}(t)\,dW_r(t) + \sigma_{2s}\,dW_s(t),
\end{aligned}$$

for $0 \leq t \leq T$, where $r, \lambda_r, \sigma_{1r}$ and σ_{2r} are $\mathbb{F}^{(W_r,W_s,W_m)'}$-progressively measurable processes satisfying

$$\int_0^T \left(|r(t)| + \lambda_r^2(t) + \sigma_{1r}^2(t) + \sigma_{2r}^2(t)\right) dt < \infty, \text{ a.s.},$$

and $\lambda_s, \sigma_{2s} \in \mathbb{R}$ are constants. The volatility matrix,

$$\begin{pmatrix} \sigma_{1r}(t) & 0 & 0 \\ \sigma_{2r}(t) & \sigma_{2s} & 0 \end{pmatrix}, \tag{2.1}$$

is assumed to have full rank, 2, (almost surely) for Lebesgue-almost-every $t \in [0, T]$.

The process r and the asset with price process S_0 are, of course, interpreted as the short interest rate process and the money market account, respectively. The asset with price process S_1 is an interest rate derivative (described in more detail below), and the asset with price process S_2 is interpreted as a stock (or a portfolio of stocks). The *money market (account)* is locally risk-free, the interest rate derivative is locally risky, but may be "terminally" risk-free (e.g. if it is a zero-coupon bond with maturity T), and the stock is a risky asset, whose price process is in general correlated with the interest rate process but also influenced by an independent (stock) risk source. The Brownian motions W_r and W_s are the sources of interest rate and stock risk, respectively (hence the subscripts), and the corresponding market prices of interest rate and stock risk are given by $\lambda_r(\cdot)$ and λ_s, respectively. Our focus is on interest rate risk rather than stock risk, and we have therefore, for simplicity, chosen to work only with a single stock and furthermore taken λ_s and σ_{2s} to be constant; this can obviously be generalized.

We shall immediately specify the model further in order to make it exhibit the special kind of risk that we have in focus. Let $T_1 < \ldots < T_n$ be fixed time points in $(0, T)$, and put $T_0 = 0$, $T_{n+1} = T$. We shall refer to the intervals $[T_i, T_{i+1})$, $i = 0, \ldots, n$, as sub-periods.

The following assumption will be in force throughout.

Assumption 5.2.1 *The processes r, σ_{1r}, and σ_{2r} are $\mathbb{F}^{(W_r,W_s)'}$-progressively measurable. The process λ_r has the form*

$$\lambda_r(t) = \mu_{(r,s)}(t, \xi(t)), \ \forall t \in [0, T],$$

where $\xi(\cdot)$ is a piece-wise constant $\mathbb{R}$-valued $\mathbb{F}^{W_m}$-adapted process with possible jumps only at $T_1, \ldots, T_n$, and, for each fixed $\xi_0 \in \mathbb{R}$, the process

$$\mu_{(r,s)}(\cdot, \xi_0) = (\mu_{(r,s)}(t, \xi_0))_{t \in [0,T]}$$

is $\mathbb{F}^{(W_r,W_s)'}$-progressively measurable.

The second statement basically means that the random behaviour of λ_r in each sub-period (T_{i-1}, T_i) is determined entirely by W_r and W_s, whereas it may be affected at each time T_i by a jump in the parameter process $\xi(\cdot)$, which in turn is determined by the independent process W_m.

The motivation for this specification is as follows. We imagine that the market incompleteness arises due to purely random "shocks" affecting the market price of interest rate risk, i.e., the process λ_r, at the fixed time points $T_1, \ldots, T_n$. The process W_m is thus considered to be observable only at these time points (and only through the jumps in ξ). We have in mind the following interpretation: Initially, and during $[0, T_1]$, the market price of interest rate risk is fully determined by the market, since there is an interest rate derivative (e.g. a zero-coupon bond) maturing at time T_1. At time T_1 the derivative matures, and a new one (e.g., a new zero-coupon bond) is issued, the opening price (per share) of which is affected by a random shock. This new derivative, in turn, is available during $(T_1, T_2]$ and thus determines the market price of interest rate risk during (T_1, T_2); it matures at time T_2, where a new derivative is issued, and so forth. We interpret the price process S_1 as the price process of an asset that is obtained by a "rolling-over" strategy in these derivatives sub-period by sub-period. The random shocks affecting the opening prices (per share) of the newly-issued derivatives is then modelled by random shocks to the market price of interest rate risk. However, during (T_{i-1}, T_i), i.e., between the random shocks, the market price of interest rate risk depends stochastically only on W_r and W_s, as is the case for the other market coefficients as well, so the market is "piece-wise complete".

We interpose a remark on terminology: In general the market does not offer (hedging possibilities for) a zero-coupon bond maturing at T. To avoid confusion we shall therefore refer to the asset with price process S_1 as an "interest rate derivative" rather than a bond. It may represent a bond in each sub-period (as explained above), but these bonds have different maturities, so using the term "bond" would be ambiguous. However, various zero-coupon bonds will play a role in the following, and to ease the presentation we shall simply refer to a zero-coupon bond with maturity $T' \in (0, T]$ as a T'-bond.

Modelling the market price of interest rate risk in this way may seem somewhat unmotivated, since there is no particular reason that it should be affected by unhedgeable risk sources only at the times $T_1, \ldots, T_n$. However, we aim for a fairly simple — and not necessarily perfectly realistic — model of the situation where the investment horizon is longer than the maturities of all the initially available interest rate derivatives (such as bonds). Single-factor term structure models typically imply completeness over the full time horizon in question and may therefore be quite unrealistic if the time horizon is long. Our model is arguably more realistic in this regard.

It should be noted that the interpretation can be relaxed: Nothing prevents the interest rate derivatives (bonds) belonging to the sub-periods from overlapping, i.e., the maturity of the derivative issued at time T_i may exceed T_{i+1}, $i = 1, \ldots, n$. Similarly, the time points $T_1, \ldots, T_n$ need not coincide with the issuance of new derivatives. The important feature of the model is that the process λ_r is affected by random shocks at times $T_1, \ldots, T_n$, and this need not have anything to do with specific events. We shall not discuss the model further, apart from the following technical remark.

Remark 5.2.2 The random shocks could as well be modelled by some $\mathbb{R}^n$-valued random vector independent of W_r and W_s, rather than through the independent Brownian motion W_m. The technical motivation for our somewhat special model is simply that it allows us to work within the Brownian motion framework, and it implies virtually no loss of generality.

B. A family of local martingales.
We introduce here a family of exponential local martingales, which will play an important role in the following (the paragraph can be skipped at the first reading). First, for $t \in [0, T]$, we set

$$\begin{aligned} Z_r(t) &= \exp\left(-\int_0^t \lambda_r(u)\, dW_r(u) - \frac{1}{2}\int_0^t \lambda_r^2(u)\, du\right), \\ Z_s(t) &= \exp\left(-\lambda_s W_s(t) - \frac{1}{2}\lambda_s^2 t\right). \end{aligned}$$

Next, letting $L(W_m)$ denote the set of $\mathbb{R}$-valued $\mathbb{F}^{(W_r,W_s,W_m)'}$-progressively measurable process $\lambda_m = (\lambda_m(t))_{t\in[0,T]}$ satisfying

$$\int_0^T \lambda_m^2(t)\, dt < \infty, \text{ a.s.},$$

we define for any $\lambda_m \in L(W_m)$ the process Z_{λ_m} by

$$Z_{\lambda_m}(t) = \exp\left(-\int_0^t \lambda_m(u)\, dW_m(u) - \frac{1}{2}\int_0^t \lambda_m^2(u)\, du\right), \ 0 \le t \le T.$$

The processes Z_r, Z_s, and Z_{λ_m} are exponential local martingales (not independent in general!), and for any $\lambda_m \in L(W_m)$, the process Z defined by

$$Z_{(r,s,\lambda_m)} = Z_r Z_s Z_{\lambda_m},$$

is also an exponential local martingale, and in particular a supermartingale.

To motivate the introduction of this family of exponential local martigales let us note that for any $\lambda_m \in L(W_m)$ such that $\mathrm{E}\left(Z_{(r,s,\lambda_m)}\right) = 1$, the probability measure Q_{λ_m} defined by

$$\mathrm{Q}_{\lambda_m}(A) = \mathrm{E}\left(1_A Z_{(r,s,\lambda_m)}\right), \ A \in \mathcal{F},$$

is an equivalent martingale measure for the market. However, in general we do not assume that $\mathrm{E}\left(Z_{(r,s,\lambda_m)}\right) = 1$ for any $\lambda_m \in L(W_m)$.

C. Investment strategies.
The investor's initial wealth is denoted by x_0 and assumed to be strictly positive. He can invest his wealth continuously and without frictions in the three assets according to any admissible investment strategy.

Definition 5.2.3 A *portfolio process* is an $\mathbb{F}^{(S_0,S_1,S_2)'}$-progressively measurable process $\pi = (\pi_1(\cdot), \pi_2(\cdot))'$ with values in $\mathbb{R}^2$ such that

$$\int_0^T \left((\pi_1(t)\sigma_{1r}(t))^2 + (\pi_2(t)\sigma_{2r}(t))^2 + \pi_2^2(t)\right) dt < \infty, \text{ a.s.} \tag{2.2}$$

Remark 5.2.4 Not all portfolio processes are considered to be admissible (an admissibility constraint will be imposed below).

The components $\pi_1(t)$ and $\pi_2(t)$ denote the amounts invested in the interest rate derivative and the stock, respectively, at time t, $t \in [0,T]$. The measurability condition imposed on $\pi_1(t)$ and $\pi_2(t)$ is important; it reflects the assumption that the investor can only observe the price processes of the assets and therefore must base his decisions on them. The flow of information from the market is represented by $\mathbb{F}^{(S_0,S_1,S_2)'}$. Clearly,

$$\mathcal{F}^{(S_0,S_1,S_2)'}(t) \subseteq \mathcal{F}^{(W_r,W_s,W_m)'}(t), \ \forall t \in [0,T],$$

and in general we have strict inclusion for $t \in (0,T]$.

We only consider strategies that are financed by the initial wealth x_0, so for a given portfolio process $\pi = (\pi_1, \pi_2)'$ the corresponding wealth process X^π develops according to

$$X^\pi(0) = x_0, \tag{2.3}$$

$$dX^\pi(t) = X^\pi(t)r(t)\,dt + dI^\pi(t), \tag{2.4}$$

where I^π is the process given by

$$I^\pi(0) = 0, \tag{2.5}$$

$$dI^\pi(t) = \pi_1(t)\sigma_{1r}(t)\,d\widetilde{W}_r(t) + \pi_2(t)\left(\sigma_{2r}(t)\,d\widetilde{W}_r(t) + \sigma_{2s}\,d\widetilde{W}_s(t)\right), \tag{2.6}$$

with

$$\widetilde{W}_r(t) = W_r(t) + \int_0^t \lambda_r(u)\,du, \tag{2.7}$$

$$\widetilde{W}_s(t) = W_s(t) + \lambda_s t. \tag{2.8}$$

The amount invested in the money market account at time $t \in [0,T]$ is given by $\pi_0(t) = X^\pi(t) - \pi_1(t) - \pi_2(t)$. We shall of course allow π to be specified through X as long as the resulting SDE (2.3)-(2.4) has a unique solution.

We impose the following admissibility constraint, which will be briefly discussed in the last part of Remark 5.4.2 below.

Definition 5.2.5 A portfolio process π is said to be *admissible* if the corresponding wealth process satisfies (almost surely)

$$X^\pi(t) \geq 0, \ \forall t \in [0, T]. \tag{2.9}$$

The set of admissible portfolio processes is denoted by $\mathcal{A}$.

It is straightforward to verify, by Itô's formula, that for any $\lambda_m \in L(W_m)$ and any $\pi \in \mathcal{A}'$, the *deflated* wealth process,

$$X^\pi S_0^{-1} Z_{(r,s,\lambda_m)},$$

is a nonnegative local martingale, and therefore a supermartingale. In particular, this implies that if the wealth hits zero, then it stays there, i.e., if

$$\tau_0^\pi := T \wedge \inf\{t \in [0, T] : X^\pi(t) = 0\}, \tag{2.10}$$

with $\inf \emptyset = \infty$, then

$$X^\pi(t) = 0, \ \forall t \in [\tau_0^\pi, T], \text{ a.s.}$$

This can be seen by the same argument as used in Karatzas and Shreve (1998), Remark 3.3.4.

Any admissible portfolio process $\pi \in \mathcal{A}$ is uniquely determined for $t < \tau_0^\pi$ by the corresponding *relative risk loadings process* $h^\pi = (h_r^\pi(t), h_s^\pi(t))'_{t\in[0,T]}$ given by

$$h_r^\pi(t) X^\pi(t) = \pi_1(t)\sigma_{1r}(t) + \pi_2(t)\sigma_{2r}(t), \tag{2.11}$$

$$h_s^\pi(t) X^\pi(t) = \pi_2(t)\sigma_{2s}, \tag{2.12}$$

for $t \in [0, \tau_0^\pi)$ (and defined arbitrarily for $t \in [\tau_0^\pi, T]$). It is seen from (2.4) and (2.6) that h_r^π and h_s^π are the relative *loadings* of the portfolio process π on the risk sources W_r and W_s, respectively, i.e.,

$$dX^\pi(t)/X^\pi(t) = r(t)\,dt + h_r^\pi(t)\,d\widetilde{W}_r(t) + h_s^\pi(t)\,d\widetilde{W}_s(t), \ 0 < t \leq T. \tag{2.13}$$

The uniqueness of the portfolio process corresponding to a given relative risk loadings process is due to the fact that the volatility matrix (2.1) has full rank.

5.3 The optimization problem

A. Utility functions and objective.
The preferences of the investor are determined by his utility function $U : \mathbb{R} \to [-\infty, \infty)$. We assume that, for some $\bar{x} \in [0, \infty)$ and $\gamma < (-\infty, 1)$, U is given either by

$$U(x) = u^{(\gamma)}(x - \bar{x}), \ x \in \mathbb{R}, \tag{3.1}$$

or by

$$U(x) = \begin{cases} u^{(\gamma)}(x), & x \geq \bar{x}, \\ -\infty, & x < \bar{x}, \end{cases} \tag{3.2}$$

where $u^{(\gamma)} : \mathbb{R} \to [-\infty, \infty)$ denotes the standard CRRA utility function with relative risk aversion coefficient $1-\gamma$, i.e.,

$$u^{(\gamma)}(x) = \begin{cases} x^\gamma/\gamma, & \text{if } \gamma \in (-\infty, 1) \setminus \{0\}, \\ \log(x), & \text{if } \gamma = 0, \end{cases}$$

for $x > 0$, $u^{(\gamma)}(x) = -\infty$ for $x < 0$, and $u^{(\gamma)}(0) = \lim_{x \searrow 0} u^{(\gamma)}(x)$. We think of $\bar{x}$ as a minimum permitted level of the terminal wealth; this constraint on the terminal wealth is sometimes referred to as *portfolio insurance*.

In the special case where $\bar{x} = 0$, the utility functions given by (3.1) and (3.2) coincide and equal $u^{(\gamma)}$. We refer to this as the *unconstrained* case.

We denote by $I : (0, \infty) \to [\bar{x}, \infty)$ the function

$$I(y) = \begin{cases} (U')^{-1}(y), & 0 < y < U'(\bar{x}+), \\ \bar{x}, & U'(\bar{x}+) \le y < \infty, \end{cases}$$

i.e., I is the inverse of the derivative U', extended if necessary to $(0, \infty)$. For U of the form (3.1) resp. (3.2), it is given by

$$I(y) = \bar{x} + y^{1/(\gamma-1)}, \ 0 < y < \infty, \tag{3.3}$$

resp.

$$I(y) = \max\left(\bar{x}, y^{1/(\gamma-1)}\right) = \bar{x} + \left(y^{1/(\gamma-1)} - \bar{x}\right)^+, \ 0 < y < \infty. \tag{3.4}$$

By maximizing the function $x \mapsto U(x) - yx$, $x > 0$, for fixed $y > 0$, it is easy to verify the important inequality

$$U(I(y)) \ge U(x) + y(I(y) - x), \ \forall x \ge 0, \ y > 0. \tag{3.5}$$

The investor has the objective of maximizing the expected utility of terminal wealth, and the value function V is thus defined as

$$V(x) := \sup_{\pi \in \mathcal{A}'} \mathrm{E}\left(U(X^\pi(T))\right),$$

where

$$\mathcal{A}' = \{\pi \in \mathcal{A} : \mathrm{E}\left(U(X^\pi(T))^-\right) < \infty\}.$$

Remark 5.3.1 The set $\mathcal{A}'$ is always non-empty in the unconstrained case, since the portfolio process given by $\pi_1 \equiv \pi_2 \equiv 0$ is seen to belong to $\mathcal{A}'$. However, it need not be non-empty in the general case with $\bar{x} > 0$. If it is empty, then the problem clearly has no solution.

5.4 The complete-market case

A. A reduced market.

In this section we consider the reduced market obtained by removing the risk source W_m, and thus taking ξ to be constant throughout $[0, T]$. Assumption 5.2.1 is then automatically strengthened to the assumption that λ_r itself is $\mathbb{F}^{(W_r,W_s)'}$-progressively measurable, which is then the case for all the market coefficients. Moreover, the third column of the volatility matrix (2.1) drops out, and the family of local martingales introduced in Paragraph 5.2.B becomes a singleton with the process $Z_r Z_s$ as its only member.

This means that the market is complete (when only $\mathcal{F}^{(W_r,W_s)'}(T)$-measurable contingent claims are allowed), and the state price density process H is given by

$$H(t) = S_0^{-1}(t) Z_r(t) Z_s(t),\ 0 \le t \le T.$$

We shall work in this section under the following very mild assumption.

Assumption 5.4.1 *The state price density process H satisfies*

$$\mathrm{E}\,(H(T)) < \infty.$$

Due to Assumption 5.4.1 it follows from Karatzas and Shreve (1998), Theorem 3.3.5, that it is possible to hedge a T-bond, starting with the initial wealth $\mathrm{E}\,(H(T))$. As we shall see, this particular interest rate derivative will play an important role in this section. In fact, one may assume without loss of generality that the interest rate derivative with price process S_1 is a T-bond; but this is not necessary.

Remark 5.4.2 The process $Z_r Z_s$ is a supermartingale, and it may in general be strict (i.e., not a martingale). In that case there is no equivalent martingale measure, and since the volatility matrix is non-singular (for Lebesgue-almost-every $t \in [0, T]$, a.s.) the market consequently admits arbitrage opportunities, even with *tame* portfolios, i.e., portfolios for which the discounted wealth process is uniformly bounded from below (see Levental and Skorohod (1995)). However, the admissibility constraint (2.9) prevents the investor from exploiting them (severely, at least): Since $X^\pi H$ a supermartingale for any $\pi \in \mathcal{A}$, we have the inequality

$$\mathrm{E}\,(X^\pi(T) H(T)) \le x_0,\ \forall \pi \in \mathcal{A}. \tag{4.1}$$

This calls for a discussion of the constraint (2.9). As is well known, in general financial market models it is necessary to require portfolio processes to be tame in order to avoid arbitrage via “doubling strategies” and similar “outrageous” portfolios (see, e.g. Karatzas and Shreve (1998), Ex. 1.2.3): If an equivalent martingale measure exists, then this requirement implies that the discounted wealth process corresponding to a tame portfolio is a supermartingale under the equivalent martingale measure, so that arbitrage is ruled out. In an optimization problem with a nonnegativity constraint on the terminal wealth (imposed either explicitly or implicitly through the investor's utility function) it is therefore necessary (and sufficient) to impose the nonnegativity constraint (2.9). When there is no equivalent martingale measure, this constraint cannot be justified by the same economic arguments. However, allowing the investor to employ any tame portfolio process would lead to an ill-posed optimization problem, because the arbitrage opportunities (which exist in the model in this case) could be exploited in any scale. We shall not discuss this further; we simply take (2.9) as a given constraint and accept the fact that our market model may allow arbitrage in some cases.

B. Solution.
Now, for notational convenience, put

$$x^* = \bar{x}\mathrm{E}\left(H(T)\right), \tag{4.2}$$

(recall Assumption 5.4.1), and consider the function $\mathcal{X} : (0, \infty) \to (x^*, \infty]$ given by

$$\mathcal{X}(y) = \mathrm{E}\left(H(T)I(yH(T))\right),\ 0 < y < \infty.$$

Corresponding to (3.1) and (3.2), we have

$$\mathcal{X}(y) = x^* + y^{1/(\gamma-1)}\mathrm{E}\left(H(T)^{\gamma/(\gamma-1)}\right),\ 0 < y < \infty, \tag{4.3}$$

and

$$\mathcal{X}(y) = x^* + \mathrm{E}\left(H(T)\left((yH(T))^{1/(\gamma-1)} - \bar{x}\right)^+\right),\ 0 < y < \infty, \tag{4.4}$$

respectively. In both cases it is easily seen that $\mathcal{X}(y) < \infty$, $\forall y \in (0, \infty)$, if and only if

$$\mathrm{E}\left(H(T)^{\gamma/(\gamma-1)}\right) < \infty, \tag{4.5}$$

and that $\mathcal{X}$ is strictly decreasing and maps $(0, \infty)$ onto (x^*, ∞) in this case. From Karatzas and Shreve (1998), Theorem 3.7.6, we now have that if (4.5) holds, and $x_0 \in (x^*, \infty)$, then there exists an optimal portfolio $\hat{\pi} \in \mathcal{A}'$, and the corresponding optimal terminal wealth is given by

$$X^{\hat{\pi}}(T) = I(H(T)\mathcal{Y}(x_0)), \tag{4.6}$$

where $\mathcal{Y} : (x^*, \infty) \to (0, \infty)$ is the inverse of $\mathcal{X}$, i.e., $\mathcal{X}(\mathcal{Y}(x)) = x$, $\forall x \in (0, \infty)$. Corresponding to (3.1) and (3.2), we have

$$X^{\hat{\pi}}(T) = \bar{x} + (H(T)\mathcal{Y}(x_0))^{1/(\gamma-1)}$$

and

$$X^{\hat{\pi}}(T) = \bar{x} + \left((H(T)\mathcal{Y}(x_0))^{1/(\gamma-1)} - \bar{x}\right)^+,$$

respectively (note that $\mathcal{Y}$ is different in the two cases, of course).

Remark 5.4.3 For completeness we briefly consider the case $x_0 \leq x^*$ (note that we must have $x^* > 0$ and $\bar{x} > 0$ in this case since $x_0 > 0$ by assumption):

If $x_0 = x^*$, then it follows from (4.1) that the only way to ensure that the terminal wealth $X^\pi(T)$ satisfies $X^\pi(T) \geq \bar{x}$, a.s., is to put $X^\pi(T) = \bar{x}$, that is, to buy (and hold) a T-bond with payoff $\bar{x}$. In the case of (3.1) this is then optimal if and only if $\gamma \in (0,1)$; if $\gamma \leq 0$, then there is no solution (i.e., $\mathcal{A}' = \emptyset$). In the case of (3.2) this is optimal for any $\gamma < 1$.

If $x_0 < x^*$, then no solution exists (i.e., $\mathcal{A}' = \emptyset$) in either case ((3.1) or (3.2)) for any $\gamma \in (-\infty, 1)$.

Remark 5.4.4 The solution to the problem as given by (4.6) is satisfactory only if

$$\mathrm{E}\left(U(X^{\hat{\pi}}(T))\right) < \sup_{x \in \mathbb{R}} U(x) = \begin{cases} \infty, & \gamma \in [0,1), \\ 0, & \gamma \in (-\infty, 0), \end{cases} \tag{4.7}$$

and the optimal investment strategy is in particular unique in this case. It is easily seen that (4.7) is valid for $\gamma \in (-\infty, 0)$. It need not hold for $\gamma = 0$, so although (4.5) obviously is satisfied in this case so that an optimal strategy exists, it need not be unique. For $\gamma \in (0,1)$ it is straightforward to verify that (4.7) holds if and only if (4.5) holds, so in this case the solution is satisfactory, and the optimal strategy is unique under condition (4.5).

Now, assume that condition (4.5) holds, and that $x_0 > x^*$. From (4.3) and (4.4) it transpires that in both cases the optimal portfolio can be viewed as a two-component portfolio: In both cases a T-bond with payoff $\bar{x}$ should be bought at time 0 and held throughout $[0, T]$. As for the remaining wealth, which initially is given by $x_0 - x^*$, the investment strategies are different, but the contingent claim $Y := H(T)^{1/(\gamma-1)}$ and the constant $\alpha := \mathcal{Y}(x_0)^{1/(\gamma-1)}$ (which is different in the two cases!) play a role in both of them:

In the case of (4.3), resp. (4.4), the investor should buy α units of Y, resp. a European call option on αY with strike price $\bar{x}$; in both cases it is easily shown that this can be done at the price $x_0 - x^*$. In the unconstrained case the optimal terminal wealth is proportional to Y, that is, the optimal strategy is to buy as many units of Y as possible and hold them until time T.

Note that if the contingent claim Y is not readily available in the market it can of course be replicated using the basic assets.

5.5 An extended Vasicek term structure

A. Model specification.

In this section we consider an extended Vasicek (1977) term structure model. The interest rate process is assumed to be given by the dynamics

$$dr(t) = \kappa(\theta - r(t))\, dt - \sigma_r\, dW_r(t),\ 0 < t \leq T,$$

where $\kappa > 0$ and $\theta, \sigma_r \in \mathbb{R}$ are constants. We shall make use of the well-known identity

$$\int_0^T r(t)\,dt = (r(0) - \theta)B(t,T) + \theta T - \sigma_r \int_0^T B(t,T)\,dW_r(t), \tag{5.1}$$

where $B(\cdot, T) : [0, T] \to \mathbb{R}$ is given by

$$B(t,T) = \frac{1 - e^{-\kappa(T-t)}}{\kappa},\ t \in [0,T].$$

We assume in this section that the market price of interest rate risk, λ_r, is constant in each sub-period (in the terminology of Assumption 5.2.1 one can thus set $\mu_{(r,s)}(\cdot, \xi) \equiv \xi$ and let ξ be the market price of interest rate risk). Moreover, we assume that the volatility of the interest rate derivative is deterministic (but possibly time-dependent), and, finally, that σ_{2r} is constant. This means that the only random behaviour of the market coefficients may take place at the times $T_1, \ldots, T_n$ and only affects λ_r.

B. The complete-market case.
In this paragraph we briefly consider the reduced complete-market case of Section 5.4, where W_m is removed from the model. This means that λ_r is constant throughout $[0, T]$ in this paragraph. We may and shall assume that the interest rate derivative is a T-bond. The volatility of the T-bond is then given by

$$\sigma_{1r}(t) = \sigma_r B(t,T),\ 0 \le t \le T.$$

The unconstrained problem has been solved explicitly by Sørensen (1999), Brennan and Xia (2000), Korn and Kraft (2001) (in the case $\gamma \in (0,1)$), and Bajeux-Besnainou et al. (2003). The constrained problem given by the utility function (3.1) has been solved explicitly by Bouiler et al. (2001) and Deelstra et al. (2003).

For $t \in [0, T]$ we have

$$Z_r(t) = \exp\left(-\lambda_r W_r(t) - \lambda_r^2 t/2\right).$$

The processes Z_r and Z_s are here independent martingales, and $Z_r Z_s$ is therefore also a martingale. In particular, there is an equivalent martingale measure (with Radon-Nikodym derivative given by $Z_r(T)Z_s(T)$), and the model precludes arbitrage opportunities. Moreover, by use of (5.1) it is easily seen that the (unique) state price density H is given by

$$\begin{aligned} H(T) &= e^{-\int_0^T r(t)\,dt} Z_r(T) Z_s(T) \\ &= \exp\left(-\int_0^T (\lambda_r - \sigma_r B(t,T))\,dW_r(t)\right) \exp\left(-\lambda_r^2 T/2 - C\right) Z_s(T), \end{aligned}$$

where

$$C = (r(0) - \theta)B(t,T) + \theta T. \tag{5.2}$$

Thus, $H(T)$ has a log-normal distribution, and Assumption 5.4.1 is clearly satisfied. Similarly, condition (4.5) is immediately seen to hold, so there exists an optimal portfolio $\hat{\pi} \in \mathcal{A}'$ with corresponding terminal wealth given by (4.6). Finally, it is easily verified that (4.7) holds for $\gamma = 0$ and thus for all $\gamma \in (-\infty, 1)$.

In the unconstrained case the optimal investment strategy, $\hat{\pi}$, is given explicitly in terms of the corresponding risk loadings process $\hat{h}$ by

$$\hat{h}_r(t) = \frac{\lambda_r - \gamma\sigma_r B(t,T)}{1-\gamma}, \tag{5.3}$$

$$\hat{h}_s(t) = \frac{\lambda_s}{1-\gamma}, \tag{5.4}$$

for $t \in [0,T]$ (see, e.g. Bouiler et al. (2001)). We refer to the abovementioned references for a closer analysis of the optimal portfolio process.

In the general constrained cases (where $\bar{x} > 0$) the optimal strategy can, as explained in Section 5.4, be viewed as a two-component strategy involving $\bar{x}$ T-bonds and a sub-portfolio corresponding either to the optimal unconstrained portfolio or to a European call option with strike price $\bar{x}$ on the optimal unconstrained terminal wealth. We shall not go further into detail here.

C. The incomplete-market case, unconstrained problem.
We now turn our attention to the incomplete market of Paragraph A, and we first consider the unconstrained problem (with $\bar{x} = 0$). We shall immediately state and prove a result that yields the optimal portfolio process (see Remark 5.5.3 for some comments).

Proposition 5.5.1 *The portfolio process $\hat{\pi}$ defined by*

$$\hat{\pi}_1(t) = \frac{X^{\pi}(t)}{1-\gamma}\left(\frac{\lambda_r(t)\sigma_{2s} - \lambda_s\sigma_{2r}}{\sigma_{2s}\sigma_r B(t,T)} - \gamma\right), \tag{5.5}$$

$$\hat{\pi}_2(t) = \frac{X^{\pi}(t)}{1-\gamma}\frac{\lambda_s}{\sigma_{2s}}, \tag{5.6}$$

$t \in [0,T]$, *is optimal in the extended, incomplete Vasicek market.*

Remark 5.5.2 From (2.11)-(2.12) it is easily verified that the risk loadings process corresponding to $\hat{\pi}$ is as in (5.3)-(5.4), with λ_r replaced by $\lambda_r(t)$ for each $t \in [0,T]$.

Note that the result is quite general in the sense that no particular assumptions (e.g. distributional) about the jumps of λ_r at $T_1, \ldots, T_n$ have been imposed.

Proof. In terms of the risk loadings process corresponding to $\hat{\pi}$ (see Remark 5.5.2), the wealth process $X^{\hat{\pi}}$ is given by

$$\begin{aligned} X^{\hat{\pi}}(t) &= x_0 \exp\left(\int_0^t \left(r(u)\,du + \frac{\lambda_r(u) - \gamma\sigma_r B(u,T)}{1-\gamma}\,d\widetilde{W}_r(u) + \frac{\lambda_s}{1-\gamma}\,d\widetilde{W}_s(u)\right)\right) \\ &\quad \times \exp\left(-\int_0^t \frac{1}{2}\left(\frac{\lambda_r(u) - \gamma\sigma_r B(u,T)}{1-\gamma}\right)^2 du - \frac{\lambda_s^2 t}{2(1-\gamma)^2}\right), \end{aligned} \tag{5.7}$$

for every $t \in [0, T]$, as can be seen by use of (2.13).

We first argue that $\hat{\pi} \in \mathcal{A}$. It is straightforward to verify that $\hat{\pi}$ satisfies (2.2) and (2.9). Moreover, $\hat{\pi}$ is $\mathbb{F}^{(S_0,S_1,S_2)'}$-progressively measurable. This is not obvious, since, for $i = 1, \ldots, n$, $\hat{\pi}(T_i)$ depends on $\lambda_r(T_i)$, which in turn depends on the history of the unobservable process W_m. However, it follows from the fact that $\mathbb{F}^{(S_0,S_1,S_2)'}$ is right-continuous and $\lambda(T_i)$ is $\mathcal{F}^{(S_0,S_1,S_2)'}(T_i + \epsilon)$-measurable for each $\epsilon > 0$ (the last assertion, in turn, follows from the structure of the market). Note that, in contrast, $\hat{\pi}(T_i)$ is in general not measurable with respect to $\sigma((S_0(t), S_1(t), S_2(t)); t \in [0, T_i])$, the natural σ-algebra generated by the price processes over $[0, T_i]$.

By use of (5.1), (5.2), and some straightforward calculations, we obtain

$$\begin{aligned} X^{\hat{\pi}}(T) &= x_0 e^C \exp\left(\int_0^T \frac{\lambda_r(t) - \sigma_r B(t,T)}{1-\gamma}\, d\widetilde{W}_r(t) + \frac{\lambda_s}{1-\gamma}(W_s(T) + \lambda_s T)\right) \\ &= x_0 e^C \exp\left(\int_0^T \frac{\lambda_r(t) - \sigma_r B(t,T)}{1-\gamma}\, d\widetilde{W}_r(t)\right) Z_s(T)^{1/(\gamma-1)} e^{\lambda_s^2 T/(2(1-\gamma))}. \end{aligned}$$

Now, consider the process H defined by

$$H(t) = S_0^{-1}(t) Z_r(t) Z_s(t),\ 0 \le t \le T. \tag{5.8}$$

It belongs to the family of exponential local martingales of Paragraph 5.2.B; it is obtained by setting $\lambda_m \equiv 0$. We have, using (3.3) (or (3.4)) with $\bar{x} = 0$,

$$\begin{aligned} I(H(T)) &= H(T)^{1/(\gamma-1)} \\ &= (S_0^{-1}(T) Z_r(T))^{1/(\gamma-1)} Z_s(T)^{1/(\gamma-1)} \\ &= e^{C/(1-\gamma)} Z_s(T)^{1/(\gamma-1)} \\ &\quad \times \exp\left(\int_0^T \frac{\lambda_r(t) - \sigma_r B(t,T)}{1-\gamma}\, dW_r(t) + \frac{\lambda_r^2(t)}{2(1-\gamma)}\, dt\right), \end{aligned}$$

where (5.1) has been used, once again, in the last equality.

By inserting (2.7) in the expression for $X^{\hat{\pi}}(T)$ and performing a few calculations we see that

$$X^{\hat{\pi}}(T) = I(H(T))M, \text{ a.s.}, \tag{5.9}$$

where

$$M = x_0 e^{C\gamma/(\gamma-1) + \lambda_s^2 T/(2(1-\gamma))} \exp\left(\int_0^T \frac{\lambda_r^2(t)/2 - \lambda_r(t)\sigma_r B(t,T)}{1-\gamma}\, dt\right).$$

We note that M is a strictly positive, $\mathcal{F}^{W_m}(T)$-measurable random variable.

Now, let π be an arbitrary strategy in $\mathcal{A}'$ (recall Remark 5.3.1), with corresponding terminal wealth denoted by $X^\pi(T)$. From (3.5), (5.9), and the obvious identity $I(H(T))M = I(H(T)M^{\gamma-1})$, we have (almost surely)

$$\begin{aligned} U(X^{\hat{\pi}}(T)) &= U(I(H(T))M) = U\left(I\left(H(T)M^{\gamma-1}\right)\right) \\ &\ge U\left(X^\pi(T)\right) + H(T)M^{\gamma-1}\left[I\left(H(T)M^{\gamma-1}\right) - X^\pi(T)\right] \\ &= U\left(X^\pi(T)\right) + H(T)M^{\gamma-1}\left[X^{\hat{\pi}}(T) - X^\pi(T)\right]. \end{aligned}$$

From (5.7) and (5.8) it is fairly straightforward to verify, by Itô's formula, that conditionally, given $\mathcal{F}^{W_m}(T)$, $X^{\hat{\pi}}H = (X^{\hat{\pi}}(t)H(t))_{t\in[0,T]}$ is an $\mathbb{F}^{(W_r,W_s)'}$-martingale (almost surely), because λ_r is a square-integrable deterministic process conditionally (almost surely). Therefore,

$$\mathrm{E}\left(H(T)M^{\gamma-1}X^{\hat{\pi}}(T)\right) = \mathrm{E}\left(M^{\gamma-1}\,\mathrm{E}\left(H(T)X^{\hat{\pi}}(T)\,\middle|\,\mathcal{F}^{W_m}(T)\right)\right) = x_0\mathrm{E}\left(M^{\gamma-1}\right),$$

and it is easily seen that $\mathrm{E}\left(M^{\gamma-1}\right) < \infty$. As for $X^{\pi}H$, one can similarly show that conditionally, given $\mathcal{F}^{W_m}(T)$, $X^{\pi}H$ is an $\mathbb{F}^{(W_r,W_s)'}$-supermartingale (almost surely), so we get

$$\mathrm{E}\left(H(T)M^{\gamma-1}X^{\pi}(T)\right) = \mathrm{E}\left(M^{\gamma-1}\,\mathrm{E}\left(H(T)X^{\pi}(T)\,|\,\mathcal{F}^{W_m}(T)\right)\right) \le x_0\mathrm{E}\left(M^{\gamma-1}\right).$$

Thus we can conclude that

$$\mathrm{E}\left(U(X^{\hat{\pi}}(T))\right) \ge \mathrm{E}\left(U(X^{\pi}(T))\right).$$

Since $\pi \in \mathcal{A}'$ it follows that $\hat{\pi} \in \mathcal{A}'$, and optimality of $\hat{\pi}$ follows from arbitrariness of π. □

Thus, in terms of the relative risk loadings, the optimal portfolio process is exactly as in the complete-market case (although, of course, the market price of interest rate risk is no longer constant, so the risk loadings have to be "updated" at the beginning of each sub-period). However, one should be careful as regards the interpretation of this result. Since the model can be viewed as being "piece-wise" complete, one might think that it would be optimal to invest in each sub-period in accordance with the optimal complete-market strategy (in respect of the time horizons of the sub-periods). This is true in terms of the relative risk loadings, but it does not necessarily mean that the optimal amounts invested in each of the three assets should be the same as in the piece-wise complete markets: If the interest rate derivative in a given sub-period, say (T_{i-1}, T_i), is a T_i-bond, then the volatility of the derivative (bond) in that period is not given by $\sigma_r B(t,T)$, but by $\sigma_r B(t,T_i)$. Thus, the optimal amount (or proportion of wealth) invested in the bond does not equal the amount (or proportion of wealth) that would be invested in the bond if only that sub-period were considered.

Remark 5.5.3 The result (as well as its proof) is inspired by (but does not follow from) similar results in Karatzas et al. (1991) and in Karatzas and Xue (1991).

D. The incomplete-market case, general problem.
The general problem is more difficult to handle because of the terminal wealth constraint. We shall not pursue it in detail; we content ourselves with a brief treatment of the terminal wealth constraint.

The contingent claim given by the constant $\bar{x}$ is not attainable. To see this we note that if $\pi \in \mathcal{A}$ were a hedging portfolio for $\bar{x}$, i.e.,

$$X^{\pi}(T) = \bar{x}, \text{ a.s.},$$

then, during the last sub-period, i.e., for $t \in (T_n, T]$, $X^\pi(t)/\bar{x}$ would necessarily have to equal the well-known value of a T-bond in the (complete) market of the last sub-period. But the value at time T_n of a T-bond certainly depends on $\lambda_r(T_n)$, which is affected by unobservable (and thus unhedgeable) randomness. Since the wealth process of any portfolio is continuous (in particular at T_n), it is not possible to hit the correct time T_n-price with certainty (or almost surely).

However, it may be possible to find super-replicating strategies, i.e., portfolio processes $\pi \in \mathcal{A}$ such that

$$X^\pi(T) \geq \bar{x}, \text{ a.s.},$$

which would in particular imply $\mathcal{A}' \neq \emptyset$.

In fact, we shall argue briefly (and somewhat heuristically) that this is the case if λ_r is bounded from below (or, equivalently, if the jumps of λ_r at $T_1, \ldots, T_n$ are all bounded from below), which is a relatively mild condition from a practical point of view (we believe that it can be weakened, but we shall not pursue this). Thus, assume that the the jumps of λ_r at $T_1, \ldots, T_n$ are bounded from below. Then there is an upper limit for the conditional price at time T_n of a T-bond, given $r(T_n)$. Moreover, as a function of $r(T_n)$ this upper limit has the form

$$C_n e^{-B(T_n,T)r(T_n)},$$

where C_n is an $\mathcal{F}^{(S_0,S_1,S_2)'}(T_{n-1})$-measurable positive random variable. This upper limit is attainable in the complete market of the sub-period $(T_{n-1}, T_n]$, and there is, in turn, an upper limit for its conditional price at time T_{n-1}, given $r(T_{n-1})$, which, in turn, is attainable in the complete market of the sub-period $(T_{n-2}, T_{n-1}]$. By reusing this argument it can be seen that a super-replicating strategy exists.

Although we have not provided a solution of the general problem, the fact that super-replicating strategies may exist is interesting in its own right, because it shows that it may be possible to obey (non-trivial) terminal wealth constraints even when T-bonds are initially unavailable.

5.6 An extended Cox-Ingersoll-Ross term structure

A. Model specification.
In this section we consider the term structure model of Cox et al. (1985). The short rate of interest is given by the dynamics

$$dr(t) = \kappa(\theta - r(t))\,dt - \sigma_r\sqrt{r(t)}\,dW_r(t),\ r(0) > 0, \tag{6.1}$$

where $\kappa, \theta, \sigma_r > 0$ are constants satisfying $\kappa\theta \geq \sigma_r^2/2$ so that $r(t) > 0,\ \forall t \in [0, T]$, almost surely.

We assume in this section that the market price of interest rate risk, $\lambda_r(\cdot)$, has the form

$$\lambda_r(t) = \xi(t)\sqrt{r(t)},\ \forall t \in [0, T], \tag{6.2}$$

where $\xi(\cdot)$ is as in Assumption 5.2.1 (which is then clearly satisfied). Moreover, we assume that the volatility of the interest rate derivative is given by

$$\sigma_{1r}(t) = \sigma_r B(t,T)\sqrt{r(t)},\ t \in [0,T], \tag{6.3}$$

for some deterministic (but possibly time-dependent) function $B(\cdot,T) : [0,T] \to [0,\infty)$, and, finally, that σ_{2r} has the form

$$\sigma_{2r}(t) = \sigma_{2r}^0\sqrt{r(t)},\ t \in [0,T],$$

for some constant $\sigma_{2r}^0 \in \mathbb{R}$.

B. The complete-market case.
In this paragraph we consider the reduced complete-market case of Section 5.4, where W_m is removed from the model. This means that $\xi(\cdot)$ is constant throughout $[0,T]$ in this paragraph. We may and shall assume that the interest rate derivative is a T-bond. The volatility of the T-bond is then given by (6.3), with

$$B(t,T) = \frac{2(e^{\delta(T-t)} - 1)}{\delta - (\kappa - \sigma_r\xi) + e^{\delta(T-t)}(\delta + \kappa - \sigma_r\xi)},\ t \in [0,T],$$

where $\delta = \sqrt{(\kappa - \sigma_r\xi)^2 + 2\sigma_r^2}$.

The unconstrained problem has been studied by Deelstra et al. (2000), who provided explicit solutions under certain restrictions on the investor's risk aversion. The constrained problem given by the utility function (3.1) has been solved explicitly by Deelstra et al. (2003) (under similar restrictions).

For $t \in [0,T]$ we have

$$Z_r(t) = \exp\left(-\int_0^t \xi\sqrt{r(t)}\,dW_r(t) - \frac{1}{2}\int_0^t \xi^2 r(t)\,dt\right). \tag{6.4}$$

The processes Z_r and Z_s are here independent, and Z_s is clearly a true martingale. The following proposition shows that this goes for Z_r as well, so Z_rZ_s is also a martingale. In particular, there is an equivalent martingale measure (with Radon-Nikodym derivative given by $Z_r(T)Z_s(T)$), and the model precludes arbitrage opportunities.

Proposition 5.6.1 *The process Z_r is a true martingale.*

Remark 5.6.2 It is by no means trivial from (6.4) that Z_r is a true martingale, and in much of the existing literature on the Cox-Ingersoll-Ross model this issue is actually not addressed at all. Some authors simply assume that Z_r is a martingale, either implicitly or explicitly, or state this without proof or reference to a known proof, while others work directly with a *risk neutral measure* (and assume the short rate dynamics (6.1) under this measure) and therefore avoid having to address the issue because no particular assumptions about the market price of interest rate risk (such as (6.2)) are imposed. However, a proof that Z_r is a martingale (in a slightly more general model) can be found in Shirakawa (2002). We provide a somewhat different and simpler proof.

Proof. It is well known (see, e.g. Pitman and Yor (1982) or Rogers (1995)) that the distribution of the process $r = (r(t))_{t\in[0,T]}$ can be characterized as follows: Let $\tau : [0,T] \to [0,\infty)$ be the strictly increasing function defined by

$$\tau(t) = \frac{\sigma_r^2(e^{\kappa t} - 1)}{4\kappa},\ 0 \le t \le T,$$

and let $Y^{(\delta)} = (Y^{(\delta)}(u))_{u\in[0,\tau(T)]}$ be a squared Bessel process with dimension $\delta = 4\kappa\theta/\sigma^2 \ge 2$, defined on an auxiliary probability space $(\Omega^\circ, \mathcal{F}^\circ, \mathrm{P}^\circ)$, and starting at $r(0)$. This process satisfies the SDE

$$Y^{(\delta)}(0) = r(0), \tag{6.5}$$

$$dY^{(\delta)}(u) = \delta\, dt + 2\sqrt{Y^{(\delta)}(u)}\, dW^\circ(u), \tag{6.6}$$

where $W^\circ = (W^\circ(u))_{u\in[0,\tau(T)]}$ is a Brownian motion on $(\Omega^\circ, \mathcal{F}^\circ, \mathrm{P}^\circ)$. Then

$$r \stackrel{d}{=} \widetilde{Y},$$

where $\widetilde{Y} = (\widetilde{Y}(t))_{t\in[0,T]}$ is the process given by

$$\widetilde{Y}(t) = e^{-\kappa t} Y^{(\delta)}(\tau(t)),\ 0 \le t \le T.$$

This can also be verified directly by using the Itô and time-change formulae (see, e.g., Karatzas and Shreve (1991)) to show that $\widetilde{Y}$ satisfies the same SDE as r (recall (6.1)), with $(\Omega^\circ, \mathcal{F}^\circ, \mathrm{P}^\circ)$ as the underlying probability space, of course.

Now, let $\delta^\circ = \min\{2n \in \mathbb{N} : 2n \ge \delta\}$ be the smallest even integer greater than δ. From a comparison theorem for solutions of SDE's (Karatzas and Shreve (1991), Theorem 5.2.18) we almost surely have

$$Y^{(\delta)}(u) \le Y^{(\delta^\circ)}(u),\ \forall 0 \le u \le \tau(T),$$

where, of course, $Y^{(\delta^\circ)}$ denotes the solution to the SDE (6.5)-(6.6) with δ replaced by δ°. Moreover, we have

$$Y^{(\delta^\circ)}(\cdot) \stackrel{d}{=} r(0) + \|B(\cdot)\|^2,$$

where $B = (B_1(t), \ldots, B_{\delta^\circ}(t))'_{t\in[0,\tau(T)]}$ is a δ°-dimensional standard Brownian motion (defined on $(\Omega^\circ, \mathcal{F}^\circ, \mathrm{P}^\circ)$).

Let $0 \le t_1 \le t_2 \le T$, and put $u_i = \tau(t_i)$, $i = 1, 2$. We have

$$\begin{aligned}
\mathrm{E}\left(\exp\left(\frac{1}{2}\int_{t_1}^{t_2} \xi_r^2 r(s)\,ds\right)\right) &= \mathrm{E}^{\mathrm{P}^\circ}\left(\exp\left(\frac{\xi_r^2}{2}\int_{t_1}^{t_2} \widetilde{Y}(s)\,ds\right)\right) \\
&= \mathrm{E}^{\mathrm{P}^\circ}\left(\exp\left(\frac{2\xi_r^2}{\sigma_r^2}\int_{u_1}^{u_2} \frac{Y^{(\delta)}(s)}{(1+4\kappa s/\sigma_r^2)^2}\,ds\right)\right) \\
&\le \mathrm{E}^{\mathrm{P}^\circ}\left(\exp\left(\frac{2\xi_r^2}{\sigma_r^2}\int_{u_1}^{u_2} Y^{(\delta)}(s)\,ds\right)\right) \\
&\le \mathrm{E}^{\mathrm{P}^\circ}\left(\exp\left(\frac{2\xi_r^2}{\sigma_r^2}\int_{u_1}^{u_2} Y^{(\delta^\circ)}(s)\,ds\right)\right) \\
&= \exp\left(\frac{2\xi_r^2(u_2-u_1)x}{\sigma_r^2}\right) C(u_1, u_2),
\end{aligned}$$

where

$$C(u_1, u_2) = \mathrm{E}^{\mathrm{P}^\circ}\left(\exp\left(\frac{2\xi_r^2}{\sigma_r^2}\int_{u_1}^{u_2} \|(B(s)\|^2\,ds\right)\right).$$

Thus, if $C(u_1, u_2) < \infty$, then $\mathrm{E}\left(\exp\left(\frac{1}{2}\int_{t_1}^{t_2} \xi_r^2 r(s)\,ds\right)\right) < \infty$. Since $B_1, \ldots, B_{\delta^\circ}$ are independent, the validity of the condition $C(u_1, u_2) < \infty$ does not depend on δ°, so we may assume $\delta^\circ = 2$. Furthermore, using the exponential series expansion it is clear that $C(u_1, u_2) < \infty$ if and only if

$$\sum_{n=2}^{\infty} \frac{1}{n!}\left(\frac{2\xi_r^2}{\sigma_r^2}\right)^n \mathrm{E}^{\mathrm{P}^\circ}\left(\left(\int_{u_1}^{u_2} \|B(s)\|^2\,ds\right)^n\right) < \infty.$$

For $n \ge 2$, the Hölder and Jensen inequalities yield

$$\begin{aligned}
\left(\int_{u_1}^{u_2} \|B(s)\|^2\,ds\right)^n &\le (u_2-u_1)^{n/2}\left(\int_{u_1}^{u_2} \|B(s)\|^4\,ds\right)^{n/2} \\
&\le (u_2-u_1)^{n-1}\int_{u_1}^{u_2} \|B(s)\|^{2n}\,ds,
\end{aligned}$$

and thus

$$\begin{aligned}
\mathrm{E}^{\mathrm{P}^\circ}\left(\left(\int_{u_1}^{u_2} \|B(s)\|^2\,ds\right)^n\right) &\le (u_2-u_1)^{n-1}\int_{u_1}^{u_2} \mathrm{E}^{\mathrm{P}^\circ}\left(\|B(s)\|^{2n}\right)\,ds \\
&= (u_2-u_1)^{n-1}\int_{u_1}^{u_2} (2s)^n n!\,ds \\
&\le 2^n(u_2-u_1)^n u_2^n n!,
\end{aligned}$$

where we have used the fact that $\|B(s)\|^2$ has a χ^2-distribution with $\delta^\circ = 2$ degrees of freedom and scale parameter $2s$. Hence, we have $C(u_1, u_2) < \infty$ if

$$\sum_{n=2}^{\infty}\left(\frac{4\xi_r^2}{\sigma_r^2}\right)^n (u_2-u_1)^n u_2^n < \infty,$$

i.e., if

$$(u_2 - u_1)u_2 < \frac{\sigma_r^2}{4\xi_r}.$$

Thus, if we let $0 = t_0 < t_1 < \ldots < t_m = T$ be a finite partition of $[0, T]$ (not to be confused with the partition given by $T_1, \ldots, T_n$) such that

$$(\tau(t_i) - \tau(t_{i-1}))\tau(t_i) < \frac{\sigma_r^2}{4\xi_r},\ \forall i = 1, \ldots, m,$$

then

$$\mathrm{E}\left(\exp\left(\frac{1}{2}\int_{t_{i-1}}^{t_i} \xi_r^2 r(s)\, ds\right)\right) < \infty,\ \forall i = 1, \ldots, m,$$

which is sufficient to ensure that Z_r is a martingale (Karatzas and Shreve (1991), Corollary 3.5.14). $\square$

We note that since $Z_r Z_s$ is a martingale and r is positive it is clear that Assumption 5.4.1 is met. As for condition (4.5), let us first note that for $\gamma \leq 0$ it is immediately verified that the condition holds by use of Assumption 5.4.1 and Jensen's inequality. For $\gamma \in (0, 1)$ the situation is much more complicated, but we shall here try to provide conditions for (4.5) when $\gamma \in (0, 1)$. The reader with no interest in this can jump directly to the end of Remark 5.6.6 below.

Let us begin with the observation that

$$\mathrm{E}\left(H(T)^{\gamma/(\gamma-1)}\right) = e^a\, \mathrm{E}\left(e^{br(T)+c\int_0^T r(t)\, dt}\right),$$

where

$$\begin{aligned}
a &= -\frac{\gamma}{\gamma-1}\left(T\left[\frac{-\lambda_s^2}{2(\gamma-1)} + \frac{\xi\kappa\theta}{\sigma_r}\right] + \frac{\xi r(0)}{\sigma_r}\right), \\
b &= \frac{\gamma}{\gamma-1}\frac{\xi}{\sigma_r}, && (6.7) \\
c &= -\frac{\gamma}{\gamma-1}\left(1 + \frac{\xi^2}{2} - \frac{\xi\kappa}{\sigma_r}\right), && (6.8)
\end{aligned}$$

which can be verified by straightforward calculations, using (6.1) in integral form and the independence between W_r and W_s. Now, from Lamberton and Lapeyre (1996), Prop. 6.2.5, we have, for any $\beta, \eta \leq 0$,

$$\mathrm{E}\left(e^{\beta r(t)+\eta\int_0^t r(u)\, du}\right) = \exp\left(\kappa\theta\phi_{\beta,\eta}(t) + r(0)\psi_{\beta,\eta}(t)\right),\ t \geq 0, \tag{6.9}$$

where

$$\begin{aligned}
\phi_{\beta,\eta}(t) &= \frac{2}{\sigma_r^2}\log\left(\frac{2\rho e^{(\rho+\kappa)t/2}}{-\sigma_r^2\beta(e^{\rho t}-1) + \rho - \kappa + e^{\rho t}(\rho+\kappa)}\right), \\
\psi_{\beta,\eta}(t) &= \frac{\beta(\rho + \kappa + e^{\rho t}(\rho - \kappa)) + 2\eta(e^{\rho t}-1)}{-\sigma_r^2\beta(e^{\rho t}-1) + \rho - \kappa + e^{\rho t}(\rho+\kappa)},
\end{aligned}$$

with $\rho = \sqrt{\kappa^2 - 2\sigma_r^2\eta}$. The following lemma extends (6.9) in the case $\eta = 0$.

Lemma 5.6.3 *For any $t > 0$ we have*

$$\begin{aligned} \mathrm{E}\left(e^{\beta r(t)}\right) &= \exp\left(\kappa\theta\phi_{\beta,0}(t) + r(0)\psi_{\beta,0}(t)\right) \\ &= e^{\beta r(0)e^{-\kappa t}/(1-2\beta L(t))}(1-2\beta L(t))^{-2\kappa\theta/\sigma^2} \end{aligned}$$

for $\beta < 1/(2L(t))$, where $L(t) = \sigma_r^2(1-e^{-\kappa t})/(4\kappa)$. If $\beta \geq 1/(2L(t))$, then $\mathrm{E}\left(e^{\beta r(t)}\right) = \infty$.

Proof. It follows from (6.9) with $\eta = 0$ (see, e.g. Lamberton and Lapeyre (1996), Ch. 6) that

$$r(t) \stackrel{D}{=} L(t)Y(t), \ \forall t > 0,$$

where $Y(t)$ has a non-central chi-squared distribution with $4\kappa\theta$ degrees of freedom and (non-centrality) parameter $\zeta(t) = r(0)e^{-\kappa t}/L(t) = 4r(0)\kappa/(\sigma^2(e^{\kappa t}-1))$. One easily finds

$$\mathrm{E}\left(e^{\mu Y(t)}\right) = \begin{cases} e^{\mu\zeta(t)/(1-2\mu)}(1-2\mu)^{-2\kappa\theta/\sigma^2}, & \mu < 1/2, \\ \infty, & \mu \geq 1/2, \end{cases}$$

and the assertion follows immediately. □

The following lemma gives a sufficient condition for integrability of $e^{\eta\int_0^t r(u)\,du}$, which is valid for some $\eta > 0$.

Lemma 5.6.4 *Let $t > 0$. If*

$$\eta \leq \frac{\kappa^2}{2\sigma^2}, \tag{6.10}$$

or if

$$\eta < \frac{\kappa}{\sigma^2(1-e^{-\kappa t})t} = \frac{1}{2L(t)t}, \tag{6.11}$$

with $L(\cdot)$ defined as in Lemma 5.6.3, then $\mathrm{E}\left(e^{\eta\int_0^t r(u)\,du}\right) < \infty$. On the other hand, if

$$\eta \geq \frac{2\kappa^3 t e^{\kappa t}}{\sigma^2(2e^{\kappa t} - (1+\kappa t)^2 - 1)},$$

then $\mathrm{E}\left(e^{\eta\int_0^t r(u)\,du}\right) = \infty$.

Proof. Straightforward calculations show, by use of (6.1) in integral form and the independence between W_r and W_s, that

$$\mathrm{E}\left(Z(t)\right) = e^{k(t)}\mathrm{E}\left(e^{\xi r(t)/\sigma_r + (\xi\kappa/\sigma_r - \xi^2/2)\int_0^t r(u)\,du}\right), \ t \geq 0,$$

where $k(t) = -\xi(\kappa\theta t + r(0))/\sigma_r$. It is easily seen that $\xi\kappa/\sigma_r - \xi^2/2 \leq \kappa^2/(2\sigma^2)$ (with equality if and only if $\xi = \kappa/\sigma$). Therefore, since Z is a martingale, we have in particular

$$\mathrm{E}\left(\exp\left(\frac{\kappa^2}{2\sigma^2}\int_0^t r(u)\,du\right)\right) < \infty,$$

and thus $\mathrm{E}\left(e^{\eta\int_0^t r(u)\,du}\right) < \infty$ for any $\eta \leq \frac{\kappa^2}{2\sigma^2}$.

Now, let $0 < \eta < 1/(2L(t)t)$. We first note that for $n \geq 2$ the Hölder and Jensen inequalities imply

$$\left[\int_0^t r(u)\,du\right]^n \leq t^{n/2}\left[\int_0^t r^2(u)\,du\right]^{n/2} \leq t^{n-1}\int_0^t r^n(u)\,du.$$

Therefore,

$$\begin{aligned}
\mathrm{E}\left(e^{\eta\int_0^t r(u)\,du}\right) &= \sum_{n=0}^{\infty}\frac{\eta^n}{n!}\mathrm{E}\left(\left[\int_0^t r(u)\,du\right]^n\right)\\
&\leq 1+\eta\mathrm{E}\left(\int_0^t r(u)\,du\right)+\frac{1}{t}\sum_{n=2}^{\infty}\frac{(\eta t)^n}{n!}\mathrm{E}\left(\int_0^t r^n(u)\,du\right)\\
&\leq 1+\eta\mathrm{E}\left(\int_0^t r(u)\,du\right)+\frac{1}{t}\int_0^t \mathrm{E}\left(e^{\eta t r(u)}\right)\,du < \infty,
\end{aligned}$$

since $\eta t < 1/(2L(t)) \leq 1/(2L(u))$, $\forall u \in [0,t]$. The second assertion is taken from Korn and Kraft (2004). □

Lemma 5.6.5 *Condition (4.5) is satisfied for any $\gamma \leq 0$. For $\gamma \in (0,1)$, (4.5) is satisfied if*

(i) $\xi > 0$ and $\kappa \geq \sigma_r/\xi + \xi\sigma_r/2$, or

(ii) $\xi > 0$, $\kappa < \sigma_r/\xi + \xi\sigma_r/2$, and $\mathrm{E}\left(e^{c\int_0^T r(t)\,dt}\right) < \infty$, or

(iii) $\xi = 0$ and $\mathrm{E}\left(e^{c\int_0^T r(t)\,dt}\right) < \infty$, or

(iv) $\xi < 0$ and $\mathrm{E}\left(e^{br(T)+c\int_0^T r(t)\,dt}\right) < \infty$,

where b and c are given by (6.7) and (6.8).

Proof. For $\gamma \leq 0$ the validity of (4.5) follows directly from Jensen's inequality. For $\gamma \in (0,1)$ each of the conditions (i)-(iv) is easily seen to be sufficient: Under condition (i) we have $b, c \leq 0$. Under conditions (ii) and (iii) we have $b \leq 0$. Under condition (iv) we have $b, c > 0$, but the second part of the condition is equivalent to (4.5). □

Lemmas 5.6.3, 5.6.4, and 5.6.5 can be applied to provide a few simple conditions for the validity of (4.5). An exact necessary and sufficient condition under which it holds is not easy to identify and, to the author's knowledge, nowhere to find in the literature.

Remark 5.6.6 Deelstra et al. (2000) claim that (6.9) holds for $\beta = b$ and $\eta = c$ if only $c \le \kappa^2/(2\sigma_r^2)$, but they do not prove it (and no proof of this is given in the cited references either). They argue that since ρ is well-defined (in $\mathbb{R}$) if only $c \le \kappa^2/(2\sigma_r^2)$, then (6.9) is also valid in this case (presumably with the limiting expressions for $\phi_{b,c}$ and $\psi_{b,c}$ obtained for $c \nearrow \kappa^2/(2\sigma^2)$, which are in fact finite, inserted for $c = \kappa^2/(2\sigma^2)$). However, this need not be the case. Indeed, choosing t, c, and b, such that

$$\frac{\kappa^2}{2\sigma^2} < c < \frac{2\kappa}{\sigma^2(1-e^{-\kappa t})t},$$

and $b \le 0$ (which is clearly possible), we have $\mathrm{E}\left(e^{c\int_0^t r(u)\,du}\right) < \infty$ by Lemma 5.6.4 and thus $\mathrm{E}\left(e^{br(t)+c\int_0^t r(u)\,du}\right) < \infty$, so (6.9) cannot be valid in this case. This suggests (but does not prove) that it is not valid in general for $0 < c \le \kappa^2/(2\sigma^2)$ either.

Now, it is fairly easy to show that (4.7) is satisfied for $\gamma = 0$. Thus, in the unconstrained case the optimal investment strategy, $\hat{\pi}$, is, (at least) for $\gamma \le 0$, given explicitly in terms of the corresponding risk loadings process $\hat{h}$ by

$$\hat{h}_r(t) = \frac{\xi\sqrt{r(t)} - \gamma\sigma_r\sqrt{r(t)}\,k(t)}{1-\gamma} = \frac{\lambda_r(t) - \gamma\sigma_r\sqrt{r(t)}\,k(t)}{1-\gamma}, \tag{6.12}$$

$$\hat{h}_s(t) = \frac{\lambda_s}{1-\gamma}, \tag{6.13}$$

for $t \in [0,T]$, where

$$k(t) = \frac{(e^{\delta'(T-t)}-1)(2+\xi^2/(1-\gamma))}{\delta' - \kappa - \xi\sigma_r\gamma/(1-\gamma) + e^{\delta'(T-t)}(\delta' + \kappa + \xi\sigma_r\gamma/(1-\gamma))},\ 0 \le t \le T,$$

with $\delta' = \sqrt{\kappa^2 + 2\sigma_r^2 c}$ (and c given by (6.8)), see e.g. Deelstra et al. (2000).

Note the similarity with the optimal relative risk loadings, (5.3)-(5.4), in the unconstrained complete-market Vasicek case. In particular, the relative stock risk loading, $\hat{h}_s$, is the same in the two models.

As for the general constrained cases (where $\bar{x} > 0$) we refer to the remarks at the end of Paragraph 5.5.B, which are also valid here.

C. The incomplete-market case, unconstrained problem.
We now turn our attention to the incomplete market of Paragraph A, and we first consider the unconstrained problem. For simplicity, we assume $\gamma < 0$ (the case $\gamma = 0$ is the case of logarithmic utility, and it is well known that in this case

the optimal portfolio is given by the generalized forms of the optimal relative risk loadings of the complete-market case, see e.g. Karatzas et al. (1991)).

It turns out that the result obtained in the extended Vasicek model does not carry over to the extended Cox-Ingersoll-Ross model. In other words, the generalized forms of the optimal relative risk loadings of the complete-market case are (in general) suboptimal in the incomplete-market case.

We shall now construct a method for obtaining an optimal portfolio process, based on ideas from discrete-stage dynamic programming. Thus, we construct optimal sub-portfolios (i.e., portfolios for each sub-period) such that the portfolio obtained by concatenating them is optimal. With a slight abuse of notation, the sets $\mathcal{A}$ and $\mathcal{A}'$ will be used to characterize sub-portfolios in an obvious way.

Let us begin by considering the situation at time T_n, i.e., at the beginning of the last sub-period. Mathematically, we work under the conditional probability, given the σ-algebra $\mathcal{F}^{(S_0,S_1,S_2)'}(T_n)$. To ease the notation we denote $\mathcal{F}^{(S_0,S_1,S_2)'}(T_i)$ by $\mathcal{F}_i^S$, $i = 1, \ldots, n$, in this paragraph. As viewed from time T_n, the "remaining market" of the sub-period $[T_n, T]$ is complete, and we can therefore use the general methodology and results of Section 5.4 to determine the optimal strategy in $[T_n, T]$. The state price density process of this sub-period, which we denote by $H(T_n, \cdot) = (H(T_n, \cdot))_{t \in [T_n, T]}$, is given by

$$H(T_n, t) = \frac{H(t)}{H(T_n)},\ T_n \le t \le T,$$

where

$$H(t) = S_0^{-1}(t) Z_r(t) Z_s(t),\ 0 \le t \le T.$$

We have $H(T_n, T_n) \equiv 1$, and

$$\begin{aligned} H(T_n, t) \quad &= \quad e^{-\int_{T_n}^t r(s)\,ds} \exp\left(-\int_{T_n}^t \xi(T_n)\sqrt{r(s)}\,dW_r(s) - \frac{1}{2}\int_{T_n}^t \xi^2(T_n) r(s)\,ds \right) \\ &\qquad \times \exp\left(-\lambda_s(W_s(t) - W_s(T_n)) - \frac{\lambda_s^2}{2}(t - T_n) \right),\ T_n < t \le T. \end{aligned}$$

From the analysis of Paragraph 5.4.B it follows that there exists an optimal portfolio $\hat{\pi} \in \mathcal{A}'$ (for the sub-period), and the corresponding optimal terminal wealth is given by

$$X^{\hat{\pi}}(T) = (H(T_n, T)\mathcal{Y}(T_n, X(T_n)))^{1/(\gamma-1)}, \tag{6.14}$$

where $X(T_n)$ is, of course, the wealth at time T_n, and $\mathcal{Y}(T_n, \cdot) : (0, \infty) \to (0, \infty)$ is given by

$$\mathcal{Y}(T_n, x) = x^{\gamma-1}\left[\mathrm{E}\left(H(T_n, T)^{\gamma/(\gamma-1)} \,\middle|\, \mathcal{F}_n^S \right) \right]^{1-\gamma},\ 0 < x < \infty.$$

The optimal portfolio in $[T_n, T]$ is of course given by the relative risk loadings (6.12)-(6.13), with $\xi(T_n)$ in the place of ξ.

Now, inserting the expression for $\mathcal{Y}(T_n, \cdot)$ in (6.14) and taking the conditional expected utility yields

$$\mathrm{E}\left(U(X^{\hat{\pi}}(T)) \,\middle|\, \mathcal{F}_n^S\right) = \left[\mathrm{E}\left(H(T_n,T)^{\gamma/(\gamma-1)} \,\middle|\, \mathcal{F}_n^S\right)\right]^{1-\gamma} X(T_n)^\gamma/\gamma. \tag{6.15}$$

The value function at time T_n therefore has the form

$$V(T_n,x) = Y(T_n)\, x^\gamma/\gamma, \;\; x > 0,$$

where

$$Y(T_n) = \left[\mathrm{E}\left(H(T_n,T)^{\gamma/(\gamma-1)} \,\middle|\, \mathcal{F}_n^S\right)\right]^{1-\gamma}.$$

Now, consider the situation at time T_{n-1}. By the dynamic programming principle we have

$$\begin{aligned} V(T_{n-1},x) &= \sup_{\pi\in\mathcal{A}'} \mathrm{E}\left(V(T_n, X^\pi(T_n)) \,|\, \mathcal{F}_{n-1}^S\right) \\ &= \sup_{\pi\in\mathcal{A}'} \mathrm{E}\left(Y(T_n)(X^\pi(T_n))^\gamma/\gamma \,|\, \mathcal{F}_{n-1}^S\right), \end{aligned}$$

where the supremum is taken over all admissible strategies for the sub-period $[T_{n-1}, T_n]$, financed by x. A portfolio process for which the supremum is achieved is optimal. We shall now show how the methodology of Section 5.4, slightly generalized, can be used to obtain an optimal portfolio process.

First, we note, as above, that, when viewed from time T_{n-1}, the market of the sub-period $[T_{n-1}, T_n]$ is complete (as long as only $\mathcal{F}^S(T_n-)$-measurable contingent claims are allowed), with state price density process $H(T_{n-1},\cdot)$ given by

$$H(T_{n-1},t) = \frac{H(t)}{H(T_{n-1})}, \;\; T_{n-1} \le t \le T,$$

where H is as above.

Now, if we set

$$Y(T_n-) = \mathrm{E}\left(Y(T_n) \,|\, \mathcal{F}^S(T_n-)\right),$$

then it suffices to find an optimal portfolio process $\hat{\pi} \in \mathcal{A}'$ for the sub-problem of maximizing the (conditional) expected value of $Y(T_n-)(X^\pi(T_n))^\gamma/\gamma$ (rather than $Y(T_n)(X^\pi(T_n))^\gamma/\gamma$); in other words, if

$$\mathrm{E}\left(Y(T_n-)(X^{\hat{\pi}}(T_n))^\gamma/\gamma \,\middle|\, \mathcal{F}_{n-1}^S\right) \ge \mathrm{E}\left(Y(T_n-)(X^\pi(T_n))^\gamma/\gamma \,|\, \mathcal{F}_{n-1}^S\right), \; \forall \pi \in \mathcal{A}',$$

then also

$$\mathrm{E}\left(Y(T_n)(X^{\hat{\pi}}(T_n))^\gamma/\gamma \,\middle|\, \mathcal{F}_{n-1}^S\right) \ge \mathrm{E}\left(Y(T_n)(X^\pi(T_n))^\gamma/\gamma \,|\, \mathcal{F}_{n-1}^S\right), \; \forall \pi \in \mathcal{A}'.$$

This follows from the identity

$$\begin{aligned} \mathrm{E}\left(Y(T_n)(X^\pi(T_n))^\gamma/\gamma \,|\, \mathcal{F}_{n-1}^S\right) &= \mathrm{E}\left(\mathrm{E}\left(Y(T_n)(X^\pi(T_n))^\gamma/\gamma \,|\, \mathcal{F}^S(T_n-)\right) \,|\, \mathcal{F}_{n-1}^S\right) \\ &= \mathrm{E}\left(Y(T_n-)(X^\pi(T_n))^\gamma/\gamma \,|\, \mathcal{F}_{n-1}^S\right), \end{aligned}$$

which holds for every $\pi \in \mathcal{A}'$.

Note that $Y(T_n)$ is in general not $\mathcal{F}^S(T_n-)$-measurable, because $\xi(T_n)$ is $\mathcal{F}_n^S$-measurable but in general not $\mathcal{F}^S(T_n-)$-measurable. In contrast we have that $Y(T_n-)$ is $\mathcal{F}^S(T_n-)$-measurable by construction.

We thus consider the sub-problem of maximizing the (conditional) expected value of $Y(T_n-)(X^\pi(T_n))^\gamma/\gamma$, given $\mathcal{F}_{n-1}^S$. This corresponds to the complete-market problem of Section 5.4 with the utility function $U(\cdot)$ replaced by $\Gamma U(\cdot)$ for some strictly positive $\mathcal{F}^S(T)$-measurable random variable Γ. The generalization of the methodology needed to allow for such a factor is straightforward (note, however, that it is crucial that it is $Y(T_n-)$, and not $Y(T_n)$, that enters as a factor in the maximization problem in order for the generalization to work).

Therefore, there is an optimal portfolio process $\hat{\pi} \in \mathcal{A}'$, and after some tedious calculations it can be shown that the corresponding optimal time T_n wealth has the form

$$X^{\hat{\pi}}(T_n) = X^{\hat{\pi}}(T_{n-1})L(T_n),$$

where $L(T_n)$ is a certain strictly positive $\mathcal{F}^S(T_n-)$-measurable random variable. Thus, the value function at time T_{n-1} has the form

$$V(T_{n-1}, x) = Y(T_{n-1})\, x^\gamma/\gamma,\ x > 0,$$

where

$$Y(T_{n-1}) = \mathrm{E}\left(L(T_n)^\gamma \,|\, \mathcal{F}_{n-1}^S \right).$$

The value function at time T_{n-1} thus has the same form as at time T_n, so by reusing this argument for each sub-period one can obtain an optimal portfolio process.

We have thus outlined a method by which an optimal portfolio process can be obtained (in principle, at least). In order to obtain the optimal portfolio explicitly one would have to solve each of the sub-period problems recursively. We shall not pursue this here (we could, without much difficulty, have gone a bit further so as to characterize the optimal relative risk loadings in an abstract form in terms of certain martingale representations, but this would not in itself be of much benefit). It is clear, though, that the optimal portfolio in general will depend on how the parameter process ξ is specified.

D. The incomplete-market case, general problem.
As in the Vasicek case (see Paragraph 5.5.D), the general problem is more difficult to handle because of the terminal wealth constraint, and apart from the following simple observations we do not treat it here.

Once again, the contingent claim given by the constant $\bar{x}$ is unattainable. However, super-replicating strategies may exist in general in this case, i.e., without particular assumptions about λ_r: If $\bar{x} \leq x$, then the portfolio process given by $\pi_1 \equiv \pi_2 \equiv 0$ is super-replicating and belongs to $\mathcal{A}'$, since $r(t) > 0$, $\forall t \in [0, T]$, a.s.

5.7 Conclusion

We have shown that the optimal portfolio of an investor with a CRRA (constant relative risk aversion) utility function in a fairly general extended and incomplete Vasicek market (with partial information) is given by the same relative risk loadings as in the complete-market case (except for the fact that the market price of interest risk parameter has to be updated in the incomplete-market case).

This result does not carry over to the extended Cox-Ingersoll-Ross market. There, the investor takes the future uncertainty into account in a somewhat more complex manner.

We have not obtained explicit results in the general cases with terminal wealth constraints, so this is a subject for future research. However, we have demonstrated that super-replicating strategies may exist, so that it is possible to obey a long-term guarantee in a market without long-term bonds.

Appendix A

Some Utility Theory

Motivated by the fact that in most (if not all) financial optimization problems, the considered objective plays a crucial role for the solution (see the discussion in Section 2.10), we offer in this appendix an account of some relevant aspects of utility theory, which forms the basis for the objectives considered in this thesis (as well as in most of the financial literature). The main purpose is to bring forward the logical foundations for using the so-called *expected utility maxim* (explained below). Apart from this we broach a few other important aspects and provide some comments on various issues that are often neglected in the literature.

The appendix may enlighten the reader with an actuarial or mathematical background, who may be only vaguely familiar with utility theory. However, for the reader who is well informed on the topic (such as an economist), most or all of the material here may be well known. We make absolutely no attempt to cover all relevant aspects; nor do we mention all the relevant literature (and we apologize for any obviously relevant references that have been omitted).

A.1 The expected utility maxim: Basics

Utility theory has been a topic of much debate in the economic literature. Basically, it forms a part of the general theory of (rational) decision making (*decision theory*), but in order to ease the presentation we shall not discuss the topic from the most general perspective.

We consider, for simplicity, an (economic) agent (e.g., an individual or an insurance company), who faces a set of different alternatives, each of which has a specified, but possibly random, effect on the wealth of the individual. An obvious example is that of an investor who invests in a financial market and can choose between different investment strategies. Other examples are those of an individual, who is to decide whether or not he wants to buy a lottery ticket, or whether or not he wants to buy insurance coverage against some risk.

To formalize the problem, we denote the agent's wealth (prior to making his decision) by W, and we denote by $\mathbb{P}$ the set of alternatives. Each alternative is formalized by a probability distribution P on $(\mathbb{R}, \mathcal{B}(\mathbb{R}))$, which is the distribution of the agent's resulting (random) wealth corresponding to the alternative in question. It is assumed that the probability distribution P corresponding to each alternative is known.

In the abovementioned investment example each $\mathrm{P} \in \mathbb{P}$ could denote the distribution of the wealth at some fixed time horizon T corresponding to an admissible investment strategy. In the insurance example there could be only two alternatives, i.e., $\mathbb{P} = \{\mathrm{P}_1, \mathrm{P}_2\}$, with P_1 given by $\mathrm{P}_1(\{W - \pi\}) = 1$, corresponding to the alternative where the insurance was purchased at the premium π, yielding the wealth $W - \pi$ with certainty, and P_2 given by the distribution of $W - Y$, with Y representing the total loss incurred, corresponding to the alternative without insurance.

Of course, the assumed probability distributions of the various alternatives play an important role. It should be mentioned that they need not be objective (universally agreed-upon); rather, they represent the agent's own probability assessments. In either case the are taken to be given, so it is assumed that, e.g., any necessary statistical estimations have been carried out.

Remark A.1.1 Witout loss of generality, one could work with random variables (defined on some underlying probability space) instead of probability distributions. However, since the distribution of each alternative is all that matters for the agent, we work with distributions, as in most of the literature on utility theory.

Now, the *expected utility maxim* or *hypothesis* basically states that the agent should choose the alternative with the greatest expected utility according to some (measurable) *utility function* $u : \mathbb{R} \to [-\infty, \infty)$ assigning a numerical value of utility to each possible outcome under consideration (the assumed range of u, $[-\infty, \infty)$, will be discussed below). In other words, the agent should act so as to maximize

$$\mathrm{E}^{\mathrm{P}}(u(X)) = \int_{\mathbb{R}} u(x)\, d\mathrm{P}(x) \tag{1.1}$$

over all $\mathrm{P} \in \mathbb{P}$. Here, X is simply the random variable given as the identity function on $(\mathbb{R}, \mathcal{B}(\mathbb{R}))$. It is tacitly assumed that the expectation is well defined for all $\mathrm{P} \in \mathbb{P}$; we discuss this further below.

We shall discuss the expected utility maxim at some length in Section A.4 below. Here we just provide a few comments. The agent's utility function represents his general preferences concerning his wealth and the risks associated with a set of alternatives. Nothing more, nothing less. One should be careful not to overinterpret the utility function and take it to be, e.g., a measure of "happiness", which is somewhat meaningless.

A.2 Definition and basic properties of utility functions

In this section we shall, as is the usual approach to optimization problems, take the expected utility maxim for granted. We provide a fairly general definition of utility functions and discuss a few basic properties.

Definition A.2.1 A utility function is a nondecreasing and concave function $u : \mathbb{R} \to [-\infty, \infty)$ with a nonempty *proper* domain, i.e.,

$$\mathrm{dom}(u) := \{x \in \mathbb{R} : u(x) > -\infty\} \neq \emptyset.$$

The condition that a utility function must be nondecreasing is due to the obvious assumption that the agent prefers “more money” to “less money”. The concavity condition is less obvious (and has indeed been debated in the literature, see, e.g. Machina (1982) and the references therein); an interpretation is that the wealthier the agent is, the smaller is his *marginal* utility, which is fairly plausible. It also implies that the agent is *risk averse*, i.e., prefers “less risk” to “more risk”, everything else being equal. It should be noted that these properties are not implied by the expected utility maxim, though. Thus, from a mathematical point of view, nothing prevents the use of a “pseudo-utility” function without the properties of Definition A.2.1, but this may lead to very counterintuitive results. Apart from a few more comments below this will not be discussed further.

For the remaning part of this section, let u be a given utility function. We impose in this section the technical assumption that

$$\int_{\{x\in\mathbb{R}:u(x)\geq 0\}} u(x)\, d\mathrm{P}(x) < \infty, \ \forall \mathrm{P} \in \mathbb{P}. \tag{2.1}$$

Then $\int_{\mathbb{R}} u(x)\, d\mathrm{P}(x)$ is well defined as a value in $[-\infty, \infty)$.

We set $x_0 = \inf(\mathrm{dom}(u))$ (with $\inf \mathbb{R} = -\infty$). Then x_0 constitutes a lower bound on feasible values of x in the sense that

$$u(x) = -\infty, \ \forall x < x_0. \tag{2.2}$$

This means that, with the terminology of Section A.1, any alternative P for which there is even the slightest chance of ending up with a terminal wealth below x_0, i.e., $\mathrm{P}((-\infty, x_0)) > 0$, has an expected utility of $-\infty$ and is therefore at least as “bad” as any other alternative.

If $x_0 = -\infty$, then $\mathrm{dom}(u) = \mathbb{R}$. Otherwise, we have

$$\mathrm{dom}(u) = \begin{cases} (x_0, \infty) & \text{if } u(x_0) = -\infty, \\ [x_0, \infty) & \text{if } u(x_0) > -\infty, \end{cases}$$

By concavity, u is continuous on $(\mathrm{dom}(u))^\circ$, the interior of $\mathrm{dom}(u)$.

We have included the element $-\infty$ in the assumed range of u. This allows for a natural definition of u on the entire set $\mathbb{R}$ of real numbers by (2.2), also if u is "originally" defined only on $\mathrm{dom}(u)$. It is also often quite convenient, in particular in an optimal investment problem with an explicit constraint on the terminal wealth that it must exceed som fixed level x_0 almost surely; this can be turned into an implicit constraint simply by setting $u(x) = -\infty$ for $x < x_0$.

If $\tilde{u}$ is another utility function, then u and $\tilde{u}$ are said to be *equivalent* if they yield the same preference system according to the expected utility maxim (we give precise meaning to the notion of a "preference system" in Section A.4 below). This is the case if and, in general, only if they are (positive) affine transformations of each other, which means that there exist $a > 0$, $b \in \mathbb{R}$ such that $\tilde{u} = au + b$. In particular this also shows that the particular values of a utility function are completely irrelevant; only the differences between the values at different points matter.

If u is strictly increasing and belongs to $C^2(\mathrm{dom}(u)^\circ)$, we define the *(absolute) risk aversion* function $a : \mathrm{dom}(u)^\circ \to \mathbb{R}$ by

$$a(x) = -\frac{u''(x)}{u'(x)},\ x \in \mathrm{dom}(u)^\circ.$$

The absolute risk aversion function is motivated and studied in Pratt (1964) (which, incidentally, is largely based on "actuarial" arguments). This function is positive and measures the degree of risk aversion (locally at x). If u is only piece-wise twice continuously differentiable we can of course define a in a piece-wise fashion.

The definition of a does not rely on the concavity of a utility function, and one could therefore define a risk aversion function $\tilde{a}$ corresponding to any increasing "pseudo-utility" function $\tilde{u}$. An agent with a preference structure given by $\tilde{u}$ would then be called risk averse, risk loving or risk neutral (locally at x), according to whether $\tilde{a}(x)$ were strictly positive, strictly negative or 0, respectively. The requirement that a utility function u must be concave is equivalent to requiring that the corresponding risk aversion function must be positive.

We note that the risk aversion functions corresponding to equivalent utility functions are identical.

A *certainty equivalent* of an alternative P is an alternative $\widetilde{\mathrm{P}} \in \mathbb{P}$ for which the outcome is known with certainty, i.e., $\widetilde{\mathrm{P}}$ is degenerate with $\widetilde{\mathrm{P}}(\{x_0\}) = 1$ for some $x_0 \in [-\infty, \infty)$, and

$$\int_{\mathbb{R}} u(x)\, d\widetilde{\mathrm{P}}(x) = u(x_0) = \int_{\mathbb{R}} u(x)\, d\mathrm{P}(x),$$

meaning that the agent is indifferent as to a choice between P and $\widetilde{\mathrm{P}}$. Note that a certainty equivalent need not exist, and need not be unique if it exists.

A.3 HARA utility functions

In this section we consider a widely used parameterized class of utility functions known as HARA (*hyperbolic absolute risk aversion*) utility functions. This class consists of utility functions in $C^2(\mathrm{dom}(u)^\circ)$ for which the corresponding risk aversion function a is strictly positive and hyperbolic (on its domain of defintion), i.e., takes the form

$$a(x) = \frac{1}{cx+d}, \; cx+d>0. \tag{3.1}$$

If $d = 0$ (and $c > 0$), one speaks of *constant relative risk aversion* (CRRA) or *isoelastic* utility, because in this case $a(x)x$ is constant on $(0,\infty)$ (see Pratt (1964) for at motivation of the notion of relative risk aversion). If $c = 0$ we speak of *constant absolute risk aversion* (CARA) utility, because in this case a is constant (the abbreviation "CARA" does not appear to be standard in the literature, though). From (3.1) it can be deduced that, up to positive affine transformations, HARA utility functions have either of the forms

$$u(x) = \begin{cases} -e^{-x/d}, & \text{if } c=0, d>0, \; (x\in\mathbb{R}), \\ \log(x+d), & \text{if } c=1, d\in\mathbb{R}, \; (x>-d), \\ c\frac{(x+d/c)^{1-\frac{1}{c}}-1}{1-\frac{1}{c}}, & \text{if } c\in\mathbb{R}\setminus\{0,1\}, d\in\mathbb{R}, \; (cx+d>0), \end{cases}$$

with $u(x) = -\infty$ for $cx + d < 0$ and $u(-d/c) = \lim_{x\searrow -d/c} u(x)$ in the case $c < 0$; $u(x) = u(-c/d) = -c^2/(c-1)$ for $cx+d < 0$ in the case $c > 0$.

For $c \in (0,\infty)$, $d \in \mathbb{R}$, we have $\mathrm{dom}(u) = (-d/c,\infty)$ if $0 < c \le 1$ and $\mathrm{dom}(u) = [-d/c,\infty)$ if $c > 1$, and the risk aversion function a is defined (in both cases) on $(-d/c,\infty)$. The CRRA utility functions correspond to the case $d = 0$, and they are often referred to as *power* utility functions (if $c \neq 1$) or *logarithmic* utility functions (if $c = 1$). The CRRA *coefficient* is given by $1/c$.

The CARA utility functions correspond to the case $c = 0$, $d > 0$, and they are usually referred to as *exponential* utility functions. In this case $\mathrm{dom}(u) = \mathbb{R}$, and the risk aversion function a is defined on all of $\mathbb{R}$.

The *quadratic* utility functions correspond to the case $c = -1$, $d \in \mathbb{R}$. The utility functions corresponding to the case $c \in (-\infty,0)\setminus\{-1\}$ are rarely used. In general, for $c < 0$, we have $\mathrm{dom}(u) = \mathbb{R}$, and the risk aversion function a is only defined on $(-\infty,-d/c)$. Note that u is constant on $[-d/c,\infty)$, which means that $-d/c$ represents a *satiation level*, i.e., a level where the maximum utility level is attained. In particular, for $c = -1$, u is only quadratic on $(-\infty, d/c]$.

Finally we note that the expression for $c \in \mathbb{R}\setminus\{0,1\}$ can obviously be simplified, but it is convenient in the stated form because

$$\lim_{c\to 1} c\frac{(x+d/c)^{1-\frac{1}{c}}-1}{1-\frac{1}{c}} = \log(x+d).$$

A.4 The expected utility maxim: Discussion

The expected utility maxim is widely accepted by most economists (but not all, as discussed below) and widely used in portfolio optimization problems as well in other areas of finance. As stated in Section A.1 it is by no means objectively obvious that it is a "correct" optimization objective, though. In particular, as already mentioned, it does not in itself imply rational behaviour.

However, the expected utility maxim enjoys substantial support from the fact that it can, at least under certain regularity conditions, be deduced from a fairly simple set of axioms concerning the agent's preferences among the alternatives in $\mathbb{P}$, as we shall now explain. The classical axiomatic approach considered here appears in several versions (which seem to differ only in minor respects, though) in the economic literature, see e.g. von Neumann and Morgenstern (1947), Samuelson (1952), Herstein and Milnor (1953), Savage (1954), Markowitz (1959), and their references. We adopt the version of Herstein and Milnor (1953).

To state the axioms we first assume that $\mathbb{P}$ is closed with respect to "finite mixtures", i.e., for any $\mathrm{P}_1, \mathrm{P}_2 \in \mathbb{P}$ and any $\alpha \in [0,1]$, the mixture distribution $\alpha \mathrm{P}_1 + (1-\alpha)\mathrm{P}_2$ is also in $\mathbb{P}$. We introduce an *order relation* $\succeq$ on $\mathbb{P}$ that represents the agent's preferences. Thus, for $\mathrm{P}_1, \mathrm{P}_2 \in \mathbb{P}$, $\mathrm{P}_1 \succeq \mathrm{P}_2$ means that the agent considers P_1 at least as "good" as P_2. Note that, by the definition of an order relation, we have $\mathrm{P} \succeq \mathrm{P}$, $\forall \mathrm{P} \in \mathbb{P}$ (reflectivity), and if $\mathrm{P}_1 \succeq \mathrm{P}_2$ and $\mathrm{P}_2 \succeq \mathrm{P}_3$, then $\mathrm{P}_1 \succeq \mathrm{P}_3$, $\forall \mathrm{P}_1, \mathrm{P}_2, \mathrm{P}_3 \in \mathbb{P}$ (transitivity). The axioms can now be stated as follows:

I The set $(\mathbb{P}, \succeq)$ is totally ordered, i.e., for any two alternatives $\mathrm{P}_1, \mathrm{P}_2 \in \mathbb{P}$ we have either $\mathrm{P}_1 \succeq \mathrm{P}_2$, $\mathrm{P}_2 \succeq \mathrm{P}_1$, or both $\mathrm{P}_1 \succeq \mathrm{P}_2$ and $\mathrm{P}_2 \succeq \mathrm{P}_1$. In the latter case we write $\mathrm{P}_1 \sim \mathrm{P}_2$.

II For any $\mathrm{P}_1, \mathrm{P}_2, \mathrm{P}_3 \in \mathbb{P}$, the sets $\{\alpha \in [0,1] : \alpha \mathrm{P}_1 + (1-\alpha)\mathrm{P}_2 \succeq \mathrm{P}_3\}$ and $\{\alpha \in [0,1] : \mathrm{P}_3 \succeq \alpha \mathrm{P}_1 + (1-\alpha)\mathrm{P}_2\}$ are closed.

III For any $\mathrm{P}_1, \mathrm{P}_2, \mathrm{P}_3 \in \mathbb{P}$ such that $\mathrm{P}_1 \sim \mathrm{P}_2$ we have $\frac{1}{2}\mathrm{P}_1 + \frac{1}{2}\mathrm{P}_3 \sim \frac{1}{2}\mathrm{P}_2 + \frac{1}{2}\mathrm{P}_3$.

Some comments are in order. Axiom I may be difficult to accept at first glance because of the seemingly implied practical problem of actually having the agent specify his preferences among all alternatives. However, it can be argued that this is not a serious problem: The important thing is the assumption that the agent *would* be able to state his preferences among any two (well-behaved) alternatives *if* he were confronted with them, which is fairly plausible. It would indeed constitute quite a problem if a preference system was to be fully determined in such a way in practice, but this is somewhat irrelevant.

As we shall see below, a much more serious problem is the assumption that the agent has preferences concerning *any* pair of alternatives, due to which the axiom is called the *completeness axiom*. It may very well be the case that there exist alternatives, which are not well behaved (in a sense that we specify below) and thus cannot be compared to others.

On a more technical note, $(\mathbb{P}, \succeq)$ is not necessarily well ordered, i.e., if $\mathrm{P}_1 \sim \mathrm{P}_2$, then this does not imply that $\mathrm{P}_1 = \mathrm{P}_2$. However, if each $\mathrm{P} \in \mathbb{P}$ is identified with its equivalence class $\{\mathrm{P}' \in \mathbb{P} : \mathrm{P} \sim \mathrm{P}'\}$, then the set is well ordered.

Axiom II is equivalent to assuming that for any $\mathrm{P}_1, \mathrm{P}_2, \mathrm{P}_3 \in \mathbb{P}$ and any convergent sequence $(\alpha_n)_{n \geq 1} \in [0,1]^{\mathbb{N}}$ we have that if $\alpha_n \mathrm{P}_1 + (1-\alpha_n)\mathrm{P}_2 \succeq \mathrm{P}_3$, $\forall n \geq 1$, then $\alpha P_1 + (1-\alpha)P_2 \succeq P_3$, where, of course, $\alpha = \lim_{n\to\infty} \alpha_n$. Thus, it basically states that the agent's preference ordering is continuous with respect to the mixing distribution, and it is known as the *continuity axiom.* It might seem a bit unmotivated, but it is necessary for the result that is reported below.

Axiom III can be interpreted as a statement that if the agent is indifferent as to the choice between P_1 and P_2, then he is also indifferent as to the choice between the alternative consisting in a fifty-fifty chance of either P_1 or P_3 and the alternative consisting in a fifty-fifty chance of either P_2 or P_3, regardless of his preferences regarding P_3. It is known as the *independence axiom*, because the choice between the two fifty-fifty alternatives is independent of P_3.

These axioms are discussed further below. Before we continue with the derivation of the expected utility maxim we interpose the remark that the axioms are necessarily obeyed if the expected utility maxim is taken as given (and the (assumed) utility function is integrable with respect to every $P \in \mathbb{P}$). That is, if $\succeq$ is defined by

$$P_1 \succeq P_2 \Leftrightarrow \int_{\mathbb{R}} u(x)\, dP_1 \geq \int_{\mathbb{R}} u(x)\, dP_2,$$

then the axioms are satisfied.

Now, Herstein and Milnor (1953) show that the axioms imply the existence of a functional $U : \mathbb{P} \to \mathbb{R}$ such that for any $P_1, P_2 \in \mathbb{P}$ and $\alpha \in [0,1]$ we have

$$P_1 \succeq P_2 \Leftrightarrow U(P_1) \geq U(P_2), \tag{4.1}$$

and

$$U(\alpha P_1 + (1-\alpha)P_2) = \alpha U(P_1) + (1-\alpha)U(P_2). \tag{4.2}$$

Moreover, U is unique up to a positive affine transformation, i.e., if $\widetilde{U} : \mathbb{P} \to \mathbb{R}$ is another functional with these properties, then there exist $a > 0$, $b \in \mathbb{R}$, such that

$$\widetilde{U} \equiv aU + b.$$

The functional U thus represents the preference ordering numerically. We shall show below that, under certain regularity assumptions, U can be (uniquely) represented by a (measurable) function $u : \mathbb{R} \to \mathbb{R}$ in the sense that

$$U(P) = E^P\left(u(X)\right) = \int_{\mathbb{R}} u(x)\, dP(x),\ \forall P \in \mathbb{P}, \tag{4.3}$$

(where it is implicitly understood that u is integrable with respect to every $P \in \mathbb{P}$). In combination with (4.1) this leads us to the conclusion that the axioms imply the validity of the expected utility maxim (under the regularity assumptions), cf. (1.1).

Now, for each $P \in \mathbb{P}$ we let $F_P : \mathbb{R} \to \mathbb{R}$ denote the corresponding distribution function. We can and shall consider U as a functional on the set of distribution functions

$$\mathbb{F} := \{F_P : P \in \mathbb{P}\}.$$

We need the following lemma.

Lemma A.4.1 *Let $U : \mathbb{F} \to \mathbb{R}$ be a functional satisfying (4.1) and (4.2) with F_{P_i} in the place of P_i, $i = 1, 2$. If U is bounded, then there exists a unique bounded finitely additive set function $\nu : \mathcal{B}(\mathbb{R}) \to \mathbb{R}$ such that*

$$U(F) = \int_{\mathbb{R}} F(x)\, d\nu(x),\ \forall F \in \mathbb{F}. \tag{4.4}$$

Remark A.4.2 Since ν is not necessarily a measure on $\mathcal{B}(\mathbb{R})$, the integral on the right-hand side of (4.4) is not defined in the usual sense. We refer to Dunford and Schwartz (1958) for the definition of the integral and for other "simple" properties of the integral and bounded finitely additive set functions used below.

Proof. Consider the set

$$\mathbb{F}^* := \{\beta_1 \mathrm{F}_{\mathrm{P}_1} - \beta_2 \mathrm{F}_{\mathrm{P}_2} : \mathrm{P}_1, \mathrm{P}_2 \in \mathbb{P},\ \beta_1, \beta_2 \geq 0\}.$$

Clearly, $\mathbb{F} \subseteq \mathbb{F}^*$, and we can obviously extend U to a linear functional $U^* : \mathbb{F}^* \to \mathbb{R}$ such that $U^*|_{\mathbb{F}} = U$ by setting

$$U^*(F) = \beta_1 U(F_{\mathrm{P}_1}) - \beta_2 U(F_{\mathrm{P}_2}),\ F \in \mathbb{F}^*, \tag{4.5}$$

where $F = \beta_1 F_{\mathrm{P}_1} - \beta_2 F_{\mathrm{P}_2}$ is the minimal decomposition of F in this form. It is straightforward to verify that $\mathbb{F}^*$ is a linear space, and we can equip it with the supremum norm, which is given by

$$\|F\| = \max(\beta_1, \beta_2),\ F \in \mathbb{F}^*.$$

It is then easily seen that

$$\sup_{F \in \mathbb{F}^* : \|F\| \leq 1} |U^*(F)| = \sup_{F \in \mathbb{F} : \|F\| \leq 1} |U(F)| = \sup_{F \in \mathbb{F}} |U(F)|,$$

and since U is bounded (by assumption) this shows that U^* is bounded as a linear functional. Since $\mathbb{F}$ is a subspace of the linear space $B(\mathbb{R}, \mathcal{B}(\mathbb{R}))$ of measurable, bounded functions on $(\mathbb{R}, \mathcal{B}(\mathbb{R}))$ with values in $\mathbb{R}$, it follows from the Hahn-Banach theorem (Dunford and Schwartz (1958), Theorem II.3.10) that U^* can be (further) extended to a bounded linear functional $U^{**} : B(\mathbb{R}, \mathcal{B}(\mathbb{R})) \to \mathbb{R}$ such that $U^{**}|_{\mathbb{F}^*} = U^*$.

Now, according to Dunford and Schwartz (1958), Theorem IV.5.1, there is a unique bounded finitely additive set function $\nu : \mathcal{B}(\mathbb{R}) \to \mathbb{R}$ satisfying (4.4) for every $F \in B(\mathbb{R}, \mathcal{B}(\mathbb{R}))$. Since $\mathbb{F} \subseteq B(\mathbb{R}, \mathcal{B}(\mathbb{R}))$, the assertion follows. □

The set function ν has bounded variation, and it can therefore be decomposed as

$$\nu = \nu^+ - \nu^-,$$

where $\nu^+, \nu^- : \mathcal{B}(\mathbb{R}) \to [0, \infty)$ are bounded finitely additive nonnegative set functions. The corresponding functions $G^+, G^- : \mathbb{R} \to [0, \infty)$, defined by

$$G^{\pm}(x) = \nu^{\pm}((-\infty, x]),\ x \in \mathbb{R},$$

are increasing and therefore have left-hand limits at each $x \in \mathbb{R}$, denoted by $G^{\pm}(x-)$. Note that $G^{\pm}$ is right-continuous if and only if $\nu^{\pm}$ is countably additive.

Now, if all distibutions in $\mathbb{P}$ have finite support (a much milder assumption from a practical point of view than from a theoretical point of view), then each $F \in \mathbb{F}$ is *simple*, and it is easily shown that

$$U(F) = a + \int_{\mathbb{R}} (G^{-}(x-) - G^{+}(x-))\, dF(x),\ \forall F \in \mathbb{F},$$

where $a = G^{+}(\infty) - G^{-}(\infty) = \nu(\mathbb{R})$. This identity also holds in general if ν is countably additive, which can be verified by use of the product rule. Thus, under either of these conditions we have

$$U(F) = \int_{\mathbb{R}} u(x)\, dF(x),\ \forall F \in \mathbb{F},$$

i.e., the expected utility maxim is valid; we have the representation (4.3) with

$$u(x) = a + G^{-}(x-) - G^{+}(x-),\ x \in \mathbb{R}.$$

In general, however, ν is only finitely additive, and when $\mathbb{P}$ contains distributions that do not have finite support, a representation of U of the form (4.3) may not exist. Moreover, we had to impose the somewhat strict assumption in Lemma A.4.1 that U was bounded. These imperfections indicate that the axiomatic approach (at least the one taken here) is not flawless. However, Fishburn (1976) and Wakker (1993) have studied the case where U is not bounded (but still finite) and shown that the expected utility maxim remains valid under certain additional conditions (or axioms). We shall not go into more detail on the issue of axiomatic deduction of the expected utility maxim, though, as it is considered to be beyond the scope and purpose of this appendix.

We shall, however, comment on the issue of boundedness under the assumption that the expected utility maxim is taken as given. In this case one must be careful when working with unbounded utility functions such as the HARA ones, which are all unbounded. If a utility function is unbounded from below (resp. above), then an “infinitely bad” (resp. “infinitely good”) probability distribution P with $\int_{\mathbb{R}} u(x)\, d\mathrm{P}(x) = -\infty$ (resp. $\int_{\mathbb{R}} u(x)\, d\mathrm{P}(x) = \infty$) will always exist. Naturally, all “infinitely bad” distributions should be excluded from the set of alternatives $\mathbb{P}$, but in many cases this must be done explicitly, i.e., they need not be excluded automatically by the problem setup. Similarly, “infinitely good” distributions should be excluded; a problem setup that allows of “infinitely good” alternatives cannot be meaningful (Samuelson (1977) discusses this, as well as other aspects). Again we note that this must be done explicitly in many cases (Korn and Kraft (2004) point out examples from the financial literature where these issues have not been taken care of, and wrong conclusions have been drawn). We also note that if “infinitely good” and “infinitely bad” were not excluded from $\mathbb{P}$, then the expected utility maxim would no longer necessarily obey the axioms, as Axiom II would be violated in general. Finally, distributions for which both the integrals of the positive and negative parts of u are infinite should also be excluded, as they cannot be compared to any other alternative by the expected utility criterion.

The independence axiom (Axiom III) has been the subject of many discussions in the literature (for an early discussion, see Wold et al. (1952), Manne and Charnes (1952), Samuelson (1952), and Malinvaud (1952)). The main (modern) argument against it seems to be that it does not conform with observed behaviour and, in particular, experimental studies, where a large number of individuals have been asked to specify their preferences among certain (hypothetical) alternatives. We refer to Machina (1982) for an elaboration of this and for an axiomatic approach to utility without the independence axiom. However, this argument can (to some extent, at least) be rejected by the presumption that such observations are mainly due to the fact that individuals with a very limited knowledge of probability theory cannot be expected to make rational decisions even in a fairly simple setup, because they do not have the insight to analyze the setup sufficiently. Whatever stand one takes, the axiom can certainly be considered rational, and we shall not discuss it further.

An argument against utility theory, which is sometimes put forth, is that the notion of a utility function is too abstract and theoretical, and that it is very difficult, if not impossible, to state one's preferences in terms of some specific utility function. However, although this may be true in some degree, it hardly constitutes an argument against the expected utility maxim as such, nor against the use of utility theory as a theoretical way to study rational behaviour and determine theoretically optimal strategies.

Other approaches to decision making under uncertainty do exist, of course. In particular, in the theory of optimal investment, some authors have advocated that in multi-period problems one should employ the strategy that maximizes the expected geometric mean rate of return, at least when the number of periods, say n, is large, and the returns in the periods are i.i.d. The argument is based on the fact that as $n \to \infty$, the geometric mean rate of return obtained with this strategy has an almost sure limit (by the law of large numbers), which is larger than the geometric mean rate of return obtained with any other strategy. However, this only holds in the limit. If the number of periods is finite, which is always the case in practice, the situation changes dramatically. Although the probability, that the abovementioned strategy will lead to a larger geometric mean return, and hence a larger terminal wealth, than any other strategy, is close to 1 for large n, the criterion totally neglects the outcomes in the event that it does not, and they may lead to extremely losses. It should be noted that maximizing the expected mean rate of return corresponds to maximizing the expected utility with the logarithmic utility function, and as such does not violate the expected utility maxim, but this is a consequence of pure mathematics, which should not disturb the discussion of principles. For an interesting an easily accessible discussion of the two approaches we refer to Ophir (1978, 1979), Latané (1978, 1979), and Samuelson (1979).

We end this section with a few additional remarks. To be able to state the axioms we assumed that $\mathbb{P}$ was closed with respect to finite mixtures. As such, this is a rather strong assumption, which is not satisfied by most natural sets of alternatives (of course, finite mixtures can be constructed (or approximated) in practice if it is possible to conduct an independent random experiment, which has a certain outcome with a given probability $\alpha \in (0,1)$, but this is somewhat irrelevant). However, it is not necessary. The crucial assumption is that the order relation is defined for any two distributions in the set of convex linear combinations of distributions in $\mathbb{P}$ (which is the smallest set of distributions containing $\mathbb{P}$ that is closed with respect to finite mixtures). Thus, the set of alternatives need not be closed with respect to finite mixtures; it is sufficient that the agent has preferences as to any pair of (hypothetical) mixture of alternatives.

Finally, the notion of a *certainty equivalent* (cf. Section A.2) does not rest on the validity of the expected utility maxim and can be defined in the general case (we do not elaborate on this, though).

A.5 Proper risk aversion

In this section we briefly discuss a concept called *proper risk aversion* — defined and studied in Pratt and Zeckhauser (1987) — mainly because it has some features of actuarial interest. As in Pratt and Zeckhauser (1987) we work here with random variables rather than probability distributions, as it is more convenient, and we use $\succeq$ as an order relation between random variables in an obvious sense. All expectations appearing below are implicitly assumed to be well defined.

Consider an agent with the utility function u (this implicitly means that he obeys the expected utility maxim). Let W denote the agent's wealth, which may be random, and let X and Y be independent random variables, which can be interpreted as the (random) outcomes of risky prospects. It is assumed that X and Y are independent of W. We have in mind the following particular interpretation: The agent is an insurance company, X and Y represent the net profits corresponding to two particular insurance policies covering independent risks, and W represents the company's wealth (i.e., equity or reserve, depending on the terminology) stemming from all other economic activities of the company (over some specific time period).

The agent's utility function u is said to be *proper* if

$$W \succeq W + X + Y \text{ whenever } W \succeq W + X \text{ and } W \succeq W + Y. \tag{5.1}$$

The interpretation of *properness* is that if the agent prefers not to take either of the risks X or Y on their own (as opposed to doing nothing), then he also prefers not to take both risks. In the insurance interpretation, if the company prefers not to issue either of the policies on their own, then it also prefers not to issue both of them. It is immediately seen that properness implies similar preference relations in the presence of any finite number of independent prospects $X_1, \ldots, X_n$, i.e., if u

is proper, then

$$W \succeq W + \sum_{i=1}^{n} X_i \text{ whenever } W \succeq W + X_i,\ \forall i = 1, \ldots, n. \tag{5.2}$$

Thus, if the insurance company has a proper utility function, then a diversification argument alone is not sufficient to ensure that having a large portfolio of policies increases the company's expected utility. Each individual policy must in itself represent an increase in expected utility, even in the (hypothetical) case where it is the company's only policy.

Moreover, if u is proper, then (5.1) (resp. (5.2)) also hold if X and Y (resp. $X_1, \ldots, X_n$) are positively correlated (but still jointly independent of W). This is particularly relevant for an insurance company, where many policies are positively correlated because they are affected by the same systematic background risks.

It seems fairly reasonable for an individual to have a proper utility function. However, in the case of a typical insurance company it could be argued that it is the effect of diversification (having many similar policies) that makes it desirable for the company to issue policies at the premiums they charge. In other words, it could be argued that if the company had to choose between (A) not having any policies and (B) having only a single policy in its portfolio at a given (competitive) premium, then it might choose (A) (most *individuals* certainly would choose not to take on an insurance obligation for a competitive premium), whereas if it had to choose between (A) and (B') having a large portfolio of similar, independent policies at the same premium rate as in (B) above, then it might prefer (B'). Indeed, this is ideally the purpose of an insurance company: To insure individuals against risks they do not want by pooling a large number of independent risks so as to make it profitable. If this argument is deemed valid, then the insurance company cannot have a proper utility function.

It adds considerable interest to the concept of proper risk aversion, in particular in respect of the above discussion in relation to insurance, that all commonly used parametric utility functions are proper. Indeed, Pratt and Zeckhauser (1987) show that all exponential (CARA) and CRRA utility functions are proper, and it is easily shown that this goes for all quadractic utility functions as well.

This raises some interesting questions. In particular, one might ask: Could it be the case that the utility function of an insurance is not even concave? Trying to provide a good answer to this question is way beyond the scope of this appendix, so we leave it as an open question.

Appendix B

Convex Analysis

In this appendix we give a brief account of some results from convex analysis, which play an important role in mathematical finance, in particular in the so-called martingale approach to optimal investment problems.

B.1 Convex analysis

Convex analysis is a mathematical subject concerned with optimization (i.e., mimimization) of complex functions (or functionals) under constraints. It forms a part of the more general discipline of mathematical analysis known as *nonlinear programming*, which deals with constrained optimization of general (not necessarily convex) function(al)s. The theory of convex analysis is in many ways much nicer than the general theory, and it is sufficient for our purposes. For the results quoted here we refer to Holmes (1975) and Kreyszig (1978).

Let X be a Banach space (i.e., a complete, normed linear space) over the scalar field $\mathbb{R}$. Let X' and X^* denote, respectively, its *algebraic conjugate (dual)* and its *(topological) conjugate (dual)* space, that is, X' (X^*) is the space of linear (linear and continuous) functionals or operators $\phi : X \to \mathbb{R}$. (Note that in general, linearity does not imply continuity, so X^* is a proper subset of X'. In fact, a linear operator is continuous if and only if it is bounded (in the linear operator sense). Furthermore, continuity at a single point is a sufficient (and of course necessary) condition for continuity of a linear operator). Both X' and X^* are linear spaces, and X^* is a Banach space with norm defined by $\|\phi\|_{X^*} = \sup_{x\in X,\, \|x\|_X \le 1} |\phi(x)|$ (this is actually the case even if X is not a Banach space but just a normed linear space).

Now, let $f : A \to \mathbb{R}$ be a convex functional defined on a convex subset $A \subseteq X$, i.e.,

$$f(tx + (1-t)y) \le tf(x) + (1-t)f(y),\ \forall x, y \in A,\ t \in [0, 1].$$

If $A^\circ \neq \emptyset$ and f is locally bounded above on a neighbourhood of a point $x \in A^\circ$, then f is continuous at every point in A°. A *subgradient* of f at a point $x \in A$ is a linear operator $\phi \in X'$ such that $\phi(y - x) = \phi(y) - \phi(x) \le f(y) - f(x)$, $\forall y \in A$.

The set of subgradients of f at $x \in A$ is called the *subdifferential* and is denoted by $\partial f(x)$. Note that a subgradient is a generalization of a derivative or gradient: If $X = \mathbb{R}^n$ and f is differentiable at x, then $\partial f(x) = \{\nabla f(x)\}$. If $\partial f(x) \neq \emptyset$ then f is said to be *subdifferentiable* at x, and $\partial f(x)$ is clearly seen to be convex. For $x \in A$ the set $\partial f(x) \cap X^*$ of continuous subgradients is a convex set, which is non-empty if $x \in A^\circ$ and f is continuous at x.

The *directional derivative* of f at $x_0 \in A^\circ$ in the direction $x \in X$ is

$$f'(x_0; x) = \lim_{t \searrow 0} \frac{f(x_0 + tx) - f(x_0)}{t}.$$

For any x_0 the directional derivative exists for any $x \in X$ and is a sublinear function of x, i.e. $f'(x_0; x + y) \leq f'(x_0; x) + f'(x_0; y),\ x, y \in X$. If $f'(x_0; x) = -f'(x_0; -x)$ for all $x \in X$, then $f'(x_0; \cdot)$ is linear and is called the *gradient* of f at x_0, and one writes $f'(x_0) = \nabla f(x_0)$.

The following result (Holmes (1975), Section 14D) establishes a relationship between directional derivatives and subdifferentials.

Theorem B.1.1 *If f is continuous at $x_0 \in A^\circ$ then for every $x \in X$ we have*

$$\begin{aligned} f'(x_0; x) &= \max\{\psi(x) : \psi \in \partial f(x_0)\}, \\ -f'(x_0; -x) &= \min\{\psi(x) : \psi \in \partial f(x_0)\}. \end{aligned}$$

It is seen that if f has a gradient at x_0 then $\partial f(x_0) = \{\nabla f(x_0)\}$.

Now, consider the convex optimization problem

$$\min_{x \in A} f(x),$$

for which we use the shorthand notation (A, f).

For $x \in A$ we denote by $F(x; A)$ the set of *feasible directions* of A at x as the set of $y \in X$ for which there exists $\delta > 0$ such that $x + ty \in A$ for $0 \leq t < \delta$, intuitively it is the set of directions in which we can move from x without leaving A immediately. Then $F(x; A)$ is a *wedge*, i.e. it is closed under multiplication by non-negative scalars. We let $F(x; A)^*$ denote the continuous dual wedge, that is, the wedge $\{\phi \in X^* \,|\, \phi(y) \geq 0,\ \forall y \in F(x; A)\}$.

We have (Holmes (1975), Section 14E)

Theorem B.1.2 *If f is continuous at $\hat{x} \in A$ then $\hat{x}$ is a solution to (A, f) if and only if*

$$\partial f(\hat{x}) \cap F(\hat{x}; A)^* \neq \emptyset.$$

The continuity assumption is only needed to prove necessity of the condition (the "only if" part).

Now, suppose that $g : X \to \mathbb{R}$ is convex and continuous and that $A = \{x \in X : g(x) \leq 0\}$. Let $\hat{x} \in A$ be a candidate for a solution. If $\hat{x} \in A^\circ$ then $F(\hat{x}; A) = X$

and thus $F(\hat{x};A)^* = \{\mathbf{0}\}$, so $\hat{x}$ is a solution if and only if $\mathbf{0} \in \partial f(\hat{x})$ (this conclusion does not depend on the form of A). If $\hat{x}$ is a boundary point, then $g(\hat{x}) = 0$ by continuity of g. Assume that $\{x \in X : g(x) < 0\} \neq \emptyset$ (*Slater's condition*, which is needed to show necessity, i.e. the "only if" parts in the conclusion that follows below). It can then be shown that

$$F(\hat{x};A)^* = (-\infty, 0]\,\partial g(\hat{x}),$$

i.e. $F(\hat{x};A)^*$ is the set of $\phi \in X^*$ of the form $\phi = -\hat{\lambda}\psi$ for some $\hat{\lambda} \geq 0$ and $\psi \in \partial g(\hat{x})$. Thus, in this case $\hat{x}$ is a solution if and only if there exist $\phi \in \partial f(\hat{x})$, $\psi \in \partial g(\hat{x})$ and a *multiplier* $\hat{\lambda} \geq 0$ such that $\phi + \hat{\lambda}\psi = \mathbf{0}$.

The overall conclusion is that $\hat{x}$ (whether it is a boundary point or not) is a solution to (A, f) if and only if there are subgradients $\phi \in \partial f(\hat{x})$, $\psi \in \partial g(\hat{x})$ and a *multiplier* $\hat{\lambda} \geq 0$ such that $\phi + \hat{\lambda}\psi = \mathbf{0}$ and $\hat{\lambda} g(\hat{x}) = 0$. This can be generalized to the case where $A = \{x \in X : g_i(x) \leq 0,\ i = 1, \ldots, n\}$ for convex, continuous functions $g_i : X \to \mathbb{R}$. Slater's condition (which is still only needed to show the "only if" parts below) is that $\{x \in X : g_i(x) < 0,\ i = 1, \ldots, n\} \neq \emptyset$. Then $\hat{x}$ is a solution to (A, f) if and only if there exist subgradients $\phi \in \partial f(\hat{x})$, $\psi_i \in \partial g_i(\hat{x})$ and *multipliers* $\hat{\lambda}_i \geq 0$ such that $\hat{\lambda}_i g_i(\hat{x}) = 0$, $i = 1, \ldots, n$ and $\phi + \sum_{i=1}^n \hat{\lambda}_i \psi_i = \mathbf{0}$, or, equivalently, if and only if there exist *multipliers* $\hat{\lambda}_i \geq 0$ such that $\hat{\lambda}_i g_i(\hat{x}) = 0$, $i = 1, \ldots, n$, and $(\hat{x}, \hat{\lambda})$ is a *saddle point* for the *Lagrangian* $L : A \times [0, \infty)^n \to \mathbb{R}$ given by

$$L(x, \lambda) = f(x) + \sum_{i=1}^n \lambda_i g_i(x),\ x \in A, \lambda \in [0, \infty)^n,$$

that is, for all $(x, \lambda) \in A \times [0, \infty)^n$,

$$L(\hat{x}, \lambda) \leq L(\hat{x}, \hat{\lambda}) \leq L(x, \hat{\lambda}).$$

B.2 Some remarks on intuition

The results of the previous section are somewhat abstract, and we shall therefore briefly provide some remarks to promote the intuitive understanding of the results. To allow for a geometric illustration we consider the case where $X = \mathbb{R}^2$, and both f and g are differentiable (and convex). Then $\partial f(\hat{x}) = \{\nabla f(x)\}$ and $\partial g(\hat{x}) = \{\nabla g(x)\}$ for every $x \in \mathbb{R}^2$. We assume that $A^\circ \neq \emptyset$, where $A = \{x \in X : g(x) \leq 0\}$ (otherwise the problem is trivial). From the previous section we then have that $\hat{x} \in \mathbb{R}^2$ is a solution to the problem of minimizing f over the set A if and only if $\hat{x} \in A$, and there exists a $\hat{\lambda} \geq 0$ such that $\hat{\lambda} g(\hat{x}) = 0$ and

$$\nabla f(\hat{x}) = -\hat{\lambda} \nabla g(\hat{x}). \tag{2.1}$$

To see this inuitively, let us first note that there are two possibilities: Either $\hat{x}$ is also an *unconstrained* minimum, or it is an *effectively constrained* minimum. It is well known that $\hat{x}$ is an *unconstrained* minimum of f if and only if $\nabla f(\hat{x}) = \mathbf{0}$, where $\mathbf{0}$ is the zero vector in $\mathbb{R}^2$, so in this case the condition is clearly necessary and sufficient (and $\hat{\lambda} = 0$). If $\hat{x}$ is an effectively constrained minimum, then $\nabla f(\hat{x}) \neq \mathbf{0}$, and $\hat{x}$ must be a boundary point of A. Then the geometric interpretation of (2.1) is that the gradients of f and g point in opposite directions (in $\mathbb{R}^2$). The gradients of a differentiable function always points in the direction of its steepest slope from the point in question. In particular, the gradient of g points "out of" A, and it is orthogonal to the boundary of A at $\hat{x}$. Thus, (2.1) says that at the point $\hat{x}$ the gradient of f must be orthogonal to the boundary of A and point "into" A. The condition that it must point into A is clear, as it it simply means that f must be increasing in the direction into A. The condition that it must be orthogonal to the boundary of A is also natural: If the direction of the steepest slope of f were not orthogonal to the boundary of A, then it would be possible to find a point along the slope at which f were smaller, because it would have to be "going downhill" from $\hat{x}$ in one of the directions along the boundary.

Appendix C

Elliptical Distributions

In this appendix we summarize some results about elliptical distributions. All definitions and results that appear below can be found in Fang et al. (1990).

C.1 Definition and properties

Definition C.1.1 A random vector X with values in $\mathbb{R}^n$ is said to have a *spherically symmetric* (or just *spherical*) distribution if

$$X \overset{d}{=} OX$$

for any orthogonal $n \times n$ matrix O.

Theorem C.1.2 *For a random vector X with values in $\mathbb{R}^n$, the following statements are equivalent:*

(i) X has a spherical distribution.

(ii) The characteristic function of X, $\psi : \mathbb{R}^n \to \mathbb{C}$, has the form

$$\psi(t_1, \ldots, t_n) = \phi(t_1^2 + \ldots + t_n^2),\ (t_1, \ldots, t_n) \in \mathbb{R}^n,$$

for some function $\phi : \mathbb{R} \to \mathbb{C}$.

(iii) X has a distributional representation of the form

$$X \overset{d}{=} RU,$$

where R and U are independent random variables with values in $[0, \infty)$ and $\mathbb{R}^n$, respectively, and U is uniformly distributed on the unit sphere $\{x \in \mathbb{R}^n : \|x\| = 1\}$.

The representation $X \overset{d}{=} RU$ is of course unique in the sense that if $X \overset{d}{=} \widetilde{R}\widetilde{U}$ is another similar representation then $R \overset{d}{=} \widetilde{R}$. The function ϕ is called the *characteristic generator* of the spherical distribution.

Definition C.1.3 A random vector X with values in $\mathbb{R}^n$ is said to have an *elliptically symmetric* (or just *elliptical*) distribution with parameters μ, Σ and ϕ, and we write $X \sim EC_n(\mu, \Sigma, \phi)$, if

$$X \stackrel{d}{=} \mu + AY,$$

where $\mu \in \mathbb{R}^n$, A is an $n \times k$ matrix such that $AA' = \Sigma$ and $\operatorname{rank}(\Sigma) = k$, and Y has a spherical distribution in $\mathbb{R}^k$ with characteristic function given by

$$(t_1, \dots, t_k) \mapsto \phi(t_1^2 + \dots + t_k^2),\ (t_1, \dots, t_k) \in \mathbb{R}^k.$$

Note that the Σ, A, ϕ and Y in Definition C.1.3 are not unique. However, we have

Theorem C.1.4 *If $X \sim EC_n(\mu, \Sigma, \phi)$ and $X \sim EC_n(\widetilde{\mu}, \widetilde{\Sigma}, \widetilde{\phi})$, then $\mu = \widetilde{\mu}$. Furthermore, if X is not degenerate, then there exists a constant $c > 0$ such that*

$$\widetilde{\Sigma} = c\Sigma, \qquad \widetilde{\phi}(t) = \phi(t/c),\ t \in \mathbb{R}.$$

In particular, $\operatorname{rank}(\Sigma)$ is unique.

Theorem C.1.5 *If $X \sim EC_n(\mu, \Sigma, \phi)$, B is a $k \times n$ matrix, and $\nu \in \mathbb{R}^k$, then*

$$\nu + BX \sim EC_k(\nu + B\mu, B\Sigma B', \phi).$$

Theorem C.1.6 *Let $X \sim EC_n(\mu, \Sigma, \phi)$ and assume that $\mathrm{E}\left(\|X\|^2\right) < \infty$. Then*

$$\mathrm{E}\,(X) = \mu, \qquad \operatorname{Var}(X) = -2\phi'(0)\Sigma.$$

Bibliography

Aase, K. K. (1984). Optimum portfolio diversification in a general continuous-time model. *Stochastic Processes and their Applications*, 18: 81–98.

Bajeux-Besnainou, I., Jordan, J. V., and Portait, R. (2003). Dynamic asset allocation for stocks, bonds, and cash. *Journal of Business*, 76(2): 263–287.

Bajeux-Besnainou, I. and Portait, R. (1998). Dynamic asset allocation in a mean-variance framework. *Management Science*, 44(11/2): 79–95.

Basak, S. (1995). A general equilibrium model of portfolio insurance. *The Review of Financial Studies*, 8(4): 1059–1090.

Bouiler, J.-F., Huang, S. J., and Taillard, G. (2001). Optimal management under stochastic interest rates: the case of a proctected defined contribution pension fund. *Insurance: Mathematics and Economics*, 28: 173–189.

Brennan, M. J. and Xia, Y. (2000). Stochastic interest rates and the bond-stock mix. *European Finance Review*, 4: 197–210.

Brennan, M. J. and Xia, Y. (2002). Dynamic asset allocation under inflation. *Journal of Finance*, 57(3): 1201–1238.

Cairns, A. (2000). Some notes on the dynamics and optimal control of stochastic pension fund models in continuous time. *ASTIN Bulletin*, 30: 19–55.

Canestrelli, E. and Pontini, S. (2000). Inquiries on the applications of multidimensional stochastic processes to financial investments. *Economics & Complexity*, 2(3): 44–62.

Canner, N., Mankiw, N. G., and Weil, D. N. (1997). An asset allocation puzzle. *American Economic Review*, 87(1): 181–191.

Cass, D. and Stiglitz, J. E. (1970). The structure of investor preferences and asset returns, and separability in portfolio allocation: A contribution to the pure theory of mutual funds. *Journal of Economic Theory*, 2: 122–160.

Chamberlain, G. (1983). A characterization of the distributions that imply mean-variance utility functions. *Journal of Economic Theory*, 29: 185–201.

Chang, S. C., Tzeng, L. Y., and Miao, J. C. Y. (2003). Pension funding incorporating downside risk. *Insurance: Mathematics and Economics*, 32: 217–228.

Cox, J. C. and Huang, C.-F. (1989). Optimal consumption and portfolio policies when asset prices follow a diffusion process. *Journal of Economic Theory*, 49: 33–83.

Cox, J. C. and Huang, C.-F. (1991). A variational problem arising in financial economics. *Journal of Mathematical Economics*, 20: 465–487.

Cox, J. C., Ingersoll, J. E., and Ross, S. A. (1985). A theory of the term structure of interest rates. *Econometrica*, 53(2): 385–408.

Cvitanić, J., Schachermayer, W., and Wang, H. (2001). Utility maximization in incomplete markets with random endowment. *Finance and Stochastics*, 5: 259–272.

Dahl, M. (2005). *On Mortality and Investment Risk in Life Insurance.* Ph.D. thesis, Laboratory of Actuarial Mathematics, University of Copenhagen.

Davis, M. H. A. (1993). *Markov Models and Optimization.* Chapman & Hall, London.

Deelstra, G., Grasselli, M., and Koehl, P.-F. (2000). Optimal investment strategies in a CIR framework. *Journal of Applied Probability*, 37: 936–946.

Deelstra, G., Grasselli, M., and Koehl, P.-F. (2003). Optimal investment strategies in the presence of a minimum guarantee. *Insurance: Mathematics and Economics*, 33: 189–207.

Delbaen, F. and Schachermayer, W. (1994). A general version of the fundamental theorem of asset pricing. *Mathematische Annalen*, 300: 463–520.

Delbaen, F. and Schachermayer, W. (1995a). The existence of absolutely continuous local martingale measures. *Annals of Applied Probability*, 5(4): 926–945.

Delbaen, F. and Schachermayer, W. (1995b). The no-arbitrage property under a change of numéraire. *Stochastics and Stochastics Reports*, 53: 213–226.

Delbaen, F. and Schachermayer, W. (1998). The fundamental theorem of asset pricing for unbounded stochastic processes. *Mathematische Annalen*, 312: 215–250.

Devolder, P., Princep, M. B., and Fabian, I. D. (2003). Stochastic optimal control of annuity contracts. *Insurance: Mathematics and Economics*, 33: 227–238.

Duffie, D. (1996). *Dynamic Asset Pricing Theory.* Princeton University Press, Princeton, NJ.

Dunford, N. and Schwartz, J. T. (1958). *Linear Operators, Part I: General Theory.* Interscience Publishers, New York.

Epstein, L. G. and Zin, S. E. (1989). Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework. *Econometrica*, 57(4): 939–969.

Fang, K. T., Kotz, S., and Ng, K.-W. (1990). *Symmetric Multivariate and Related Distributions.* Chapman and Hall, London.

Fishburn, P. C. (1976). Unbounded utility functions in expected utility theory. *The Quarterly Journal of Economics*, 90(1): 163–168.

Fleming, W. H. and Rishel, R. W. (1975). *Deterministic and Stochastic Optimal Control.* Springer-Verlag, New York Berlin Heidelberg.

Fleming, W. H. and Soner, H. M. (1993). *Controlled Markov Processes and Viscosity Solutions.* Springer-Verlag, New York Berlin Heidelberg.

Gerber, H. U. (1997). *Life Insurance Mathematics.* Third edition, Springer-Verlag, New York Berlin Heidelberg.

Gerber, H. U. and Shiu, E. S. W. (2000). Investing for retirement: Optimal capital growth and dynamic asset allocation. *North American Actuarial Journal*, 4(2): 42–62.

Grossman, S. J. and Zhou, Z. (1996). Equilibrium analysis of portfolio insurance. *Journal of Finance*, 51(4): 1379–1403.

Haberman, S. (1994). Autoregressive rates of return and the variability of pension contributions and fund levels for a defined benefit pension scheme. *Insurance: Mathematics and Economics*, 14: 219–240.

Haberman, S. and Sung, J. H. (1994). Dynamic approaches to pension funding. *Insurance: Mathematics and Economics*, 15: 151–162.

Haberman, S. and Vigna, E. (2002). Optimal investment strategies and risk measures in defined contribution pension schemes. *Insurance: Mathematics and Economics*, 31: 35–69.

Haberman, S. and Wong, L. Y. P. (1997). Moving average rates of return and the variability of pension contributions and fund levels for a defined benefit pension scheme. *Insurance: Mathematics and Economics*, 20: 115–135.

Hakansson, N. H. (1969). Optimal investment and consumption strategies under risk, an uncertain lifetime, and insurance. *International Economic Review*, 10(3): 443–466.

Hakansson, N. H. (1970). Optimal investment and consumption strategies under risk for a class of utility functions. *Econometrica*, 38: 587–607.

Harrison, J. M. and Kreps, D. M. (1979). Martingales and arbitrage in multiperiod security markets. *Journal of Economic Theory*, 20: 381–408.

Harrison, J. M. and Pliska, S. R. (1981). Martingales and stochastic integrals in the theory of continuous trading. *Stochastic Processes and their Applications*, 11: 215–260.

Harrison, J. M. and Pliska, S. R. (1983). A stochastic calculus model of continuous trading: Complete markets. *Stochastic Processes and their Applications*, 15: 313–316.

He, H. and Pearson, N. D. (1991a). Consumption and portfolio policies with incomplete markets and short-sale constraints: The finite-dimensional case. *Mathematical Finance*, 1(3): 1–10.

He, H. and Pearson, N. D. (1991b). Consumption and portfolio policies with incomplete markets and short-sale constraints: The infinite dimensional case. *Journal of Economic Theory*, 54: 259–304.

Herstein, I. N. and Milnor, J. (1953). An axiomatic approach to measurable utility. *Econometrica*, 21(2): 291–297.

Hindy, A. and Huang, C.-F. (1993). Optimal consumption and portfolio rules with durability and local substitution. *Econometrica*, 61(1): 85–121.

Ho, T. and Lee, S. (1986). Term structure movements and pricing interest rate contingent claims. *Journal of Finance*, 41: 1011–1029.

Hoem, J. M. (1969). Markov chain models in life insurance. *Blätter der Deutschen Gesellschaft für Versicherungsmathematik*, 9: 91–107.

Holmes, R. B. (1975). *Geometric Functional Analysis and its Applications.* Springer, New York.

Hull, J. C. and White, A. (1990). Pricing interest-rate derivative securities. *Review of Financial Studies*, 3: 573–592.

Jacod, J. and Shiryaev, A. N. (2003). *Limit Theorems for Stochastic Processes.* Springer.

Jensen, B. A. and Sørensen, C. (2001). Paying for minimum interest guarantees: Who should compensate who? *European Financial Management*, 7(2): 183–211.

Josa-Fombellida, R. and Rincon-Zapatero, J. P. (2001). Minimization of risk in pension funding by means of contributions and portfolio selection. *Insurance: Mathematics and Economics*, 29: 35–45.

Karatzas, I., Lehoczky, J. P., Sethi, S. P., and Shreve, S. E. (1986). Explicit solution of a general investment/consumption problem. *Mathematics of Operations Research*, 11: 261–294.

Karatzas, I., Lehoczky, J. P., and Shreve, S. E. (1987). Optimal portfolio and consumption decisions for a “small investor” on a finite horizon. *SIAM Journal on Control and Optimization*, 25: 1557–1586.

Karatzas, I., Lehoczky, J. P., Shreve, S. E., and Xu, G.-L. (1991). Martingale and duality methods for utility maximization in an incomplete market. *SIAM Journal on Control and Optimization*, 29: 702–730.

Karatzas, I. and Shreve, S. E. (1991). *Brownian Motion and Stochastic Calculus.* Second Edition, Springer-Verlag, New York Berlin Heidelberg.

Karatzas, I. and Shreve, S. E. (1998). *Methods of Mathematical Finance.* Springer-Verlag, New York Berlin Heidelberg.

Karatzas, I. and Xue, X.-X. (1991). A note on utility maximization under partial observations. *Mathematical Finance*, 1(2): 57–70.

Khanna, A. and Kulldorff, M. (1998). A generalization of the mutual fund theorem. *Finance and Stochastics*, 3: 167–185.

Korn, R. (1997). *Optimal Portfolios.* World Scientific, Singapore New Jersey London Hong Kong.

Korn, R. and Kraft, H. (2001). A stochastic control approach to portfolio problems with stochastic interest rates. *SIAM Journal on Control and Optimization*, 40(4): 1250–1269.

Korn, R. and Kraft, H. (2004). On the stability of continuous-time portfolio problems with stochastic opportunity set. *Mathematical Finance*, 14(3): 403–414.

Kramkov, D. and Schachermayer, W. (1999). The asymptotic elasticity of utility functions and optimal investment in incomplete markets. *Annals of Applied Probability*, 9(3): 904–950.

Kreyszig, E. (1978). *Introductory Functional Analysis with Applications.* John Wiley & Sons, New York.

Krylov, N. V. (1980). *Controlled Diffusion Processes.* Springer, New York Berlin Heidelberg.

Lakner, P. and Slud, E. (1991). Optimal consumtion by a bond investor: The case of random interest rate adapted to a point process. *SIAM Journal on Control and Optimization*, 29: 638–655.

Lamberton, D. and Lapeyre, B. (1996). *Introduction to Stochastic Calculus Applied to Finance*. Chapman & Hall, London.

Latané, H. A. (1978). The geometric-mean principle revisited: a reply. *Journal of Banking and Finance*, 2: 395–398.

Latané, H. A. (1979). The geometric mean criterion continued. *Journal of Banking and Finance*, 3: 309–311.

Levental, S. and Skorohod, A. V. (1995). A necessary and sufficient condition for absence of arbitrage with tame portfolios. *Annals of Applied Probability*, 5: 906–925.

Lintner, J. (1965). The valuation of risk assets and the selection of risky investments in stock portfolios and capital budgets. *Review of Economics and Statistics*, 47(1): 13–37.

Machina, M. J. (1982). "Expected utility" analysis without the Independence Axiom. *Econometrica*, 50(2): 277–324.

Malinvaud, E. (1952). Note on von Neumann-Morgenstern's strong independence axiom. *Econometrica*, 20(4): 679. See the editiorial note on page 661.

Manne, A. S. and Charnes, A. (1952). The strong independence assumption — gasoline blends and probability mixtures. *Econometrica*, 20(4): 665–669. See the editiorial note on page 661.

Markowitz, H. M. (1952). Portfolio Selection. *Journal of Finance*, 7: 77–91.

Markowitz, H. M. (1959). *Portfolio Selection: Efficient Diversification of Investments*. John Wiley & Sons, New York.

Merton, R. C. (1969). Lifetime portfolio selection under uncertainty: The continuous-time case. *Review of Economics and Statistics*, 51: 247–257.
Reprinted as Chapter 4 in Merton (1992).

Merton, R. C. (1971). Optimum consumption and portfolio rules in a continuous-time model. *Journal of Economic Theory*, 3(4): 373–413.
Erratum. *Journal of Economic Theory*, 6: 213-214 (1973).
Corrected version reprinted as Chapter 5 in Merton (1992).

Merton, R. C. (1973). An intertemporal capital asset pricing model. *Econometrica*, 41: 867–887.
Extended version reprinted as Chapter 15 in Merton (1992).

Merton, R. C. (1992). *Continuous-Time Finance*. Blackwell Publishers, Cambridge, MA, Oxford, UK.

Møller, T. (2001). Risk-minimizing hedging strategies for insurance payment processes. *Finance and Stochastics*, 5(4): 419–446.

Mossin, J. (1966). Equilibrium in a capital asset market. *Econometrica*, 34(4): 768–783.

Mossin, J. (1968). Optimal multiperiod portfolio policies. *The Journal of Business*, 41: 215–229.

Munk, C. and Sørensen, C. (2004). Optimal consumption and investment strategies with stochastic interest rates. *Journal of Banking and Finance*, 28(8): 1987–2013.

Nielsen, P. H. (2004). Utility maximization and risk minimization in life and pension insurance. Working paper no. 201, Laboratory of Actuarial Mathematics, University of Copenhagen.

Nielsen, P. H. (2005). Optimal bonus strategies in life insurance: The Markov chain interest rate case. *Scandinavian Actuarial Journal*, pages 81–102.

Nielsen, P. H. (2006). Utility maximization and risk minimization in life and pension insurance. *Finance and Stochastics*, 10: 75–97.

Norberg, R. (1991). Reserves in life and pension insurance. *Scandinavian Actuarial Journal*, pages 3–24.

Norberg, R. (1995). A time-continuous Markov chain interest model with applications to insurance. *Insurance: Mathematics and Economics*, 17: 171–180.

Norberg, R. (1999). A theory of bonus in life insurance. *Finance and Stochastics*, 3: 373–390.

Norberg, R. (2000). Basic Life Insurance Mathematics. Lecture notes, Laboratory of Actuarial Mathematics, University of Copenhagen.

Norberg, R. (2001). On bonus and bonus prognoses in life insurance. *Scandinavian Actuarial Journal*, 2: 126–147.

Norberg, R. (2003). The Markov chain market. *ASTIN Bulletin*, 33: 265–287.

Norberg, R. (2005). Anomalous PDEs in Markov chains: Domains of validity and numerical solutions. *Finance and Stochastics*, 9(4): 519–537.

Ophir, T. (1978). The geometric-mean principle revisited. *Journal of Banking and Finance*, 2: 103–107.

Ophir, T. (1979). The geometric-mean principle revisited: a reply to a reply. *Journal of Banking and Finance*, 3: 301–303.

Pitman, J. and Yor, M. (1982). A decomposition of Bessel bridges. *Zeitschrift für Wahrscheinlichkeitsteorie und verwandte Gebiete*, 59: 425–457.

Pliska, S. R. (1986). A stochastic calculus model of continuous trading: Optimal portfolios. *Mathematical Methods of Operations Research*, 11: 371–382.

Pratt, J. W. (1964). Risk aversion in the small and in the large. *Econometrica*, 32: 122–136.
Erratum. *Econometrica*, 44(2): 420 (1976).

Pratt, J. W. and Zeckhauser, R. J. (1987). Proper risk aversion. *Econometrica*, 55(1): 143–154.

Protter, P. (1990). *Stochastic Integration and Differential Equations*. Springer-Verlag, Berlin Heidelberg.

Radner, R. and Shepp, L. (1996). Risk vs. profit potential: A model for corporate strategy. *Journal of Economic Dynamics & Control*, 20: 1373–1393.

Richard, S. F. (1975). Optimal consumption, portfolio and life insurance rules for an uncertain lived individual in a continuous time model. *Journal of Financial Economics*, 2: 187–203.

Rogers, L. C. G. (1995). Which model for term-structure of interest rates should one use? In Davis, M., Duffie, D., Fleming, W., and Shreve, S., editors, *Mathematical Finance, IMA Volume 65*, pages 93–116. Springer, New York.

Rogers, L. C. G. and Williams, D. (1987). *Diffusions, Markov Processes, and Martingales, Vol. 2*. John Wiley & Sons, Chichester New York Brisbane Toronto Singapore.

Samuelson, P. A. (1952). Probability, utility, and the independence axiom. *Econometrica*, 20(4): 670–678. See the editorial note on page 661.

Samuelson, P. A. (1969). Lifetime portfolio selection by dynamic stochastic programming. *The Review of Economics and Statistics*, 51: 239–246.

Samuelson, P. A. (1977). St. Petersburg Paradoxes: Defanged, dissected, and historically described. *Journal of Economic Literature*, 15(1): 24–55.

Samuelson, P. A. (1979). Why we should not make mean log of wealth big though years to act are long. *Journal of Banking and Finance*, 3: 305–307.

Savage, L. J. (1954). *The Foundations of Statistics*. John Wiley & Sons, New York.

Schachermayer, W. (1992). A Hilbert space proof of the fundamental theorem of asset pricing in finite discrete time. *Insurance: Mathematics and Economics*, 11: 249–257.

Schachermayer, W. (2002). Optimal investment in incomplete financial markets. In Geman, H., Madan, D., Pliska, S. R., and Vorst, T., editors, *Mathematical Finance, Bachelier Congress 2000*, pages 427–462. Springer-Verlag, New York Berlin Heidelberg.

Schweizer, M. (2001). A guided tour through quadratic hedging approaches. In Jouini, E., Cvitanic, J., and Musiela, M., editors, *Option Pricing, Interest Rates and Risk Management*, pages 538–574. Cambridge University Press.

Sethi, S. and Taksar, M. (1988). A note on Merton's "Optimum consumption and portfolio rules in a continuous-time model". *Journal of Economic Theory*, 46: 395–401.

Sharpe, W. F. (1964). Capital asset prices: A theory of market equilibrium under conditions of risk. *Journal of Finance*, 19(3): 425–442.

Shirakawa, H. (2002). Squared Bessel processes and their applications to the square root interest rate model. *Asia-Pacific Financial Markets*, 9: 169–190.

Sommer, D. (1997). Pricing and hedging of contingent claims in term structure models with exogenous issuing of new bonds. *European Financial Management*, 3(3): 269–292.

Sørensen, C. (1999). Dynamic asset allocation and fixed income management. *Journal of Financial and Quantitative Analysis*, 34: 513–531.

Steffensen, M. (2001). *On Valuation and Control in Life and Pension Insurance*. Ph.D. thesis, Laboratory of Actuarial Mathematics, University of Copenhagen.

Steffensen, M. (2002). Intervention options in life insurance. *Insurance: Mathematics and Economics*, 31: 71–85.

Steffensen, M. (2004). On Merton's problem for life insurers. *ASTIN Bulletin*, 34(1): 5–25.

Taksar, M. I. (2000). Optimal risk and dividend distribution control models for an insurance company. *Mathematical Methods of Operations Research*, 51: 1–42.

Taylor, G. (2002). Stochastic control of funding systems. *Insurance: Mathematics and Economics*, 30: 323–350.

Tepla, L. (2001). Optimal investment with minimum performance constraints. *Journal of Economic Dynamics & Control*, 25: 1629–1645.

Tobin, J. (1958). Liquidity preference as behavior towards risk. *Review of Economic Studies*, 25(2): 65–86.

Tobin, J. (1965). The theory of portfolio selection. In Hahn, F. H. and Brechling, F. P. R., editors, *The Theory of Interest Rates*, pages 3–51. Macmillan, London.

Vasicek, O. (1977). An equilibrium characterization of the term structure. *Journal of Financial Economics*, 5: 177–188.

Vigna, E. and Haberman, S. (2001). Optimal investment strategy for defined contribution pension schemes. *Insurance: Mathematics and Economics*, 28: 233–262.

von Neumann, J. and Morgenstern, O. (1947). *Theory of Games and Economic Behavior*. Princeton University Press, Princeton, NJ.

Wakker, P. (1993). Unbounded utility for Savage's "Foundations of Statistics" and other models. *Mathematics of Operations Research*, 18(2): 446–485.

Wold, H., Shackle, G. L. S., and Savage, L. J. (1952). Ordinal preference or cardinal utility? *Econometrica*, 20(4): 661–664. See the editiorial note on page 661.

Yong, J. and Zhou, X. Y. (1999). *Stochastic Controls*. Springer-Verlag, New York Berlin Heidelberg.