

DISS. ETH NO. 20749

NON-CONCAVE UTILITY MAXIMIZATION:
OPTIMAL INVESTMENT, STABILITY AND APPLICATIONS

A dissertation submitted to
ETH ZURICH

for the degree of
Doctor of Sciences

presented by
CHRISTIAN ROCHUS AUGUST REICHLIN
MSc ETH Mathematics
born April 25, 1984
citizen of Schwyz, Switzerland

accepted on the recommendation of

| | |
|----------------------------|-------------|
| Prof. Dr. Martin Schweizer | examiner |
| Prof. Dr. Rose-Anne Dana | co-examiner |
| Prof. Dr. Thorsten Hens | co-examiner |
| Prof. Dr. Miklós Rásonyi | co-examiner |

2012

Abstract

In this thesis, we study optimal investment problems for agents with a non-concave utility function and distorted beliefs. The main motivation for such optimization problems comes from behavioural finance and is to understand the implications of several psychological effects on portfolio optimization and the ensuing consequences for the financial markets. Our results can also be used for other applications in finance, such as manager compensation and portfolio delegation. We present three types of results. We first analyze optimal investment problems for an (exogenously) fixed financial market. For various levels of generality, we provide conditions such that the optimization problem is well posed and admits a solution. We derive fundamental economic properties of the maximizer and analyze the sensitivity of the indirect utility with respect to the initial capital. For a broad class of models, we solve the optimal investment problem explicitly. Secondly, we perform a stability analysis for optimal investment problems. This allows us to analyze the impact of drift misspecification or changing time horizon on the optimal investment. We also use these results to provide numerical procedures for the explicit computation of the optimal investment. Finally, we study implications of behavioural effects on a financial market equilibrium where prices are determined endogenously by (behavioural) demand and supply. We mainly focus on the empirically well-studied relationship between the pricing density and the aggregate endowment and compare the results with empirical estimates from the S&P 500. Throughout, we also provide several examples illustrating our results and exhibiting a number of new and unexpected phenomena.

Kurzfassung

Diese Dissertation beschäftigt sich mit Entscheidungsproblemen für Agenten mit nicht-konkaven Nutzenfunktionen und nicht-linearen (verzerrten) Erwartungen. Diese Optimierungsprobleme sind motiviert durch die Erkenntnisse der Verhaltensökonomie und dienen dazu, Auswirkungen von verschiedenen psychologischen Effekten auf die Portfoliooptimierung sowie die Finanzmärkte zu untersuchen. Die Resultate können auch auf andere ökonomische Fragestellungen wie zum Beispiel die Analyse von erfolgsbezogenen Anreizsystemen angewendet werden. Die Resultate können in drei Klassen eingeteilt werden. Erstens betrachten wir nicht-konkave Nutzenmaximierungsprobleme für einen (exogen) fixierten Finanzmarkt. Wir geben Bedingungen für die Existenz einer Lösung und analysieren die Eigenschaften des Maximierers. Weiter untersuchen wir den indirekten Nutzen als Funktion des Startkapitals. Für verschiedene Modellklassen kann das Optimierungsproblem explizit gelöst werden. Der zweite Teil umfasst Resultate über die Stabilität der oben beschriebenen Optimierungsproblemen. Damit kann zum Beispiel der Einfluss von fehlerhafter Modellspezifikation auf die Entscheidungen untersucht werden. Die Stabilitätsresultate können auch für die numerische Berechnung der optimalen Strategie verwendet werden. Schliesslich untersuchen wir die Konsequenzen von risikosuchendem Verhalten und verzerrten Erwartungen für ein Finanzmarktgleichgewicht, in welchem die Preise endogen durch Angebot und Nachfrage bestimmt werden. Wir untersuchen hauptsächlich die Beziehung zwischen der Preisdichte und dem gesamtwirtschaftlichen Vermögen und vergleichen diese mit den empirischen Schätzungen für den S&P 500. Die verschiedenen Resultate dieser Arbeit werden durch konkrete Beispiele illustriert.

Acknowledgements

First and foremost, I wish to express my deep and sincere gratitude to my supervisors Martin Schweizer and Thorsten Hens. I have benefited from Martin Schweizer in excellent lectures, seminars, and many helpful discussions. His advice, constructive criticism, and suggestions have improved this thesis substantially. I am grateful to Thorsten Hens for his continuing encouragement, valuable research collaboration and generous financial support. His profound intuition about behavioural finance was a great inspiration. Moreover, I wish to thank Rose-Anne Dana and Miklós Rásonyi for readily accepting to act as co-examiners. During my time as a doctoral student, I greatly benefited from numerous discussions with Christoph Czichowsky, Martin Herdegen, Marcel Nutz and Urs Schweri. In addition, I am grateful for the comments of Erik Balder, Laurence Carassus, Walter Farkas, Jens Jackwerth, Kasper Larsen, Johannes Muhle-Karbe, Marc-Oliver Rieger, Stephan Sturm, Josef Teichmann and Alexandre Ziegler. I wish to thank Marc-Oliver Rieger for being the driving force in creating my PhD-position between ETH and UZH. Jens Jackwerth graciously provided the density estimates used for the calibrations in [58] and Chapter V. Wolfgang Härdle, Thorsten Hens, Michael Magill and Martine Quinzii kindly invited me to Berlin, Stanford and Los Angeles; I enjoyed the friendly and inspiring atmosphere at these institutions. Special thanks go to my friends and colleagues at ETH, UZH and outside academia for the time they shared with me. The final thanks go to my family, particularly to my parents, for all their support and encouragement over the years. This thesis is dedicated to them.

Financial support by the National Centre of Competence in Research “Financial Valuation and Risk Management” (NCCR FINRISK) for parts of this thesis is gratefully acknowledged.

Contents

| | | |
|------------|---|------------|
| I | Introduction | 1 |
| I.1 | Existing literature | 3 |
| I.2 | Overview of the thesis | 4 |
| II | Demand problem with a given pricing measure | 9 |
| II.1 | Introduction | 9 |
| II.2 | Problem formulation and intuition | 11 |
| II.3 | Existence and properties of a maximizer | 16 |
| II.4 | The value function | 22 |
| II.5 | Results for an atomless model | 25 |
| II.6 | Conclusion | 34 |
| II.7 | Appendix A: Facts from convex analysis | 35 |
| II.8 | Appendix B: Non-smooth utility maximization | 38 |
| III | Stability and asymptotics along a sequence of models | 43 |
| III.1 | Introduction | 43 |
| III.2 | Problem formulation and main results | 45 |
| III.3 | Applications | 52 |
| III.4 | Stability of the demand problem for RDEU | 62 |
| III.5 | Stability of the demand problem for ENCU | 68 |
| III.6 | Conclusion | 80 |
| III.7 | Appendix | 80 |
| IV | Examples in incomplete markets | 85 |
| IV.1 | Introduction | 85 |
| IV.2 | Problem formulation and existence | 87 |
| IV.3 | Models on finite probability spaces | 93 |
| IV.4 | A sufficient condition for $u(x, U) = u(x, U_c)$ | 97 |
| IV.5 | Examples | 105 |
| IV.6 | Conclusion | 116 |
| V | Implications for a financial market equilibrium | 117 |
| V.1 | Introduction | 117 |
| V.2 | Related literature | 119 |

x Contents

| | | |
|------|--|-----|
| V.3 | Setup | 121 |
| V.4 | The pricing kernel puzzle | 124 |
| V.5 | Boundary behaviour of the pricing kernel | 127 |
| V.6 | Non-concave utility | 129 |
| V.7 | Incorrect beliefs | 132 |
| V.8 | Incomplete markets | 139 |
| V.9 | Conclusion | 143 |
| V.10 | Appendix: Proofs | 143 |

Chapter I

Introduction

How do we make decisions under risk? This is a central and fundamental question in financial economics. Expected concave utility reigned for several decades as one dominant theory for decision making under uncertainty. This theory assumes that agents are risk-averse and that they evaluate possible outcomes linearly with respect to a prior subjective probability. However, there is considerable empirical evidence that agents tend to switch between risk-averse and risk-seeking behaviour depending on the context, and that they overweight extreme events having small probabilities. The impact of these behavioural effects and the resulting theory is exemplified by the award of the Nobel Prize in economics to Daniel Kahneman in 2002 for “having integrated insights from psychological research into economic science, especially concerning human judgment and decision-making under uncertainty”.¹

In the literature, there are several different approaches to model alternative decision criteria accounting for the behavioural effects observed in psychological research. In this thesis, we focus on preference functionals that assign to a positive random payoff f the value

$$V(f) := \int U(f) d(T \circ P) = \int_0^\infty T(P[U(f) > x]) dx \quad (1.1)$$

for a *non-concave* and *non-smooth* utility function U on \mathbb{R}_+ and a strictly increasing function $T: [0, 1] \rightarrow [0, 1]$ representing the *probability distortion* of beliefs. The non-concave utility function allows us to generate risk-seeking behaviour; particular forms of the distortion T lead to overweighting of extreme events. From a theoretical point of view, (1.1) arises naturally as a representation for preference functionals satisfying a certain comonotonicity condition (see, for instance, Schmeidler [100]). In applications, they also serve as the main building block for several behavioural theories such as *rank-dependent expected utility* (RDU or RDEU; see Quiggin [89]) or *cumulative prospect theory* (CPT; see Tversky and Kahneman [105]). If we set $T(p) = p$, then (1.1) covers the classical expected utility functional as well.

¹Excerpt from the official press release of *The Royal Swedish Academy of Sciences*.

Having specified the decision criteria, one important task is to understand how agents described by these preferences behave. This gives an idea whether one can model the decisions observed in empirical experiments. A first step in this direction is to study the *demand problem*

$$\text{maximize } V(f) \text{ over all } f \in L_+^0 \text{ satisfying } E_Q[f] \leq x \quad (1.2)$$

for some probability measure $Q \approx P$. We refer to Q as the *pricing measure*. One possible interpretation of (1.2) is as follows: There are two periods. At the first date 0, there is uncertainty about which state of the world will occur at the second date T . An agent having initial capital x and preferences over date T outcomes, represented by the preference functional V , buys at date 0 for the price $E_Q[f]$ a nonnegative contingent claim f for date T . When choosing f , the agent maximizes the value $V(f)$ of future outcomes subject to the constraint that his expenditure should not be larger than x . In Chapters II and III, we study the demand problem (1.2) in detail.

While problem (1.2) is formulated for an abstract setting, the second step is to study the problem in a more concrete situation, such as a financial market. In this way, one can understand the consequences for portfolio selection of behavioural effects as well as of the associated theories. To formalize this, we assume that a risk-free bank account yielding zero interest and risky assets with price process S are available on the financial market, and that the agent is “small” in the sense that his actions do not influence the financial market. The agent starts at time 0 with initial capital x and trades in a self-financing way along a strategy ϑ . We denote the resulting final position at time T by $X_T^{x,\vartheta}$. The *portfolio selection problem* is then to

$$\text{maximize } V(X_T^{x,\vartheta}) \text{ over all admissible trading strategies } \vartheta. \quad (1.3)$$

In the special case that the financial market is complete, the two problems (1.2) and (1.3) are directly related; but otherwise they differ. In Chapter IV, we examine the portfolio optimization problem (1.3) in the general case.

In order to also examine the consequences for financial markets of behavioural effects as well as of the associated theories, one finally studies the implications of behavioural portfolio selection for a financial market equilibrium where prices are determined by demand and supply. The prices (and the price functional) are therefore not exogenously given as in the problems (1.2) and (1.3), but are determined endogenously from the interactions of the agents. This allows us to understand how behavioural effects aggregate from an individual level to the market level. In Chapter V, we take a step in this direction.

Problems (1.2) and (1.3) and the study of the resulting equilibrium are (or lead to) economically meaningful and mathematically challenging stochastic optimization problems. The overall goal of this thesis is to improve the understanding of these non-standard optimization problems and their resulting economic implications.

I.1 Existing literature

There is a vast literature on the empirical evidence as well as the axiomatic foundations for preferences described by V . We refer to Wakker [106] for an overview on the former and to Quiggin [89] for an overview on the latter. We focus here on the literature about optimization problems involving V and its consequences for financial market equilibria. Some results in the literature are formulated for related, but slightly different, preference functionals. Our literature review also includes these results; most of them also hold (with some minor modifications) for our setting.

Results on the demand problem (1.2). There are two main approaches to deal with the demand problem (1.2). One of them is to study (1.2) in the special case without distortion (i.e., $T(p) = p$) and to use the idea of Friedman and Savage [45] and Aumann and Perles [4] to study the non-concave utility maximization problem by reducing it to the classical concave case. Bailey et al. [5], Hartley and Farrell [53] and Rieger [94] use this idea to provide some basic intuition about non-concavities in the utility function. Berkelaar et al. [13] and Carassus and Pham [29] examine the portfolio selection problem (1.3) for a complete financial market in continuous time, rewrite it in the form (1.2) and solve it by reducing it to a concave one. The same technique appears also in a slightly different context: Carpenter [32] and Larsen [77] apply it to study the impact of option compensation and portfolio delegation on portfolio selection problems, and Basak and Makarov [9] analyze performance based salary systems within a bank. In all three articles, the key optimization problem is of the form (1.2) in the special case without distortion, and the problem is solved by reducing it to a concave one. In general, this *concavification* idea is very useful – *if* it can be applied. It then easily gives the existence and several properties of a solution, all in a simple way via classical concave utility maximization. In Section II.5, we present an easily verifiable structural assumption on the model under which this *concavification* approach can indeed be applied.

The second main approach to (1.2) (which works also for the case with distortion) is to exploit the law-invariance of the functional V . This allows one to split the problem into two subproblems: One step is to optimize the functional over all payoff distributions, and the second step is to choose the cheapest final position with a given distribution. Jin and Zhou [64] use this idea, but their assumptions on the coordination between the model and the preference parameters are restrictive and somewhat ad hoc. Carlier and Dana [31] then study the problem more systematically and use a weak convergence approach via Helly's selection principle to show the existence of a maximizer. While these results still rely on some extra assumptions on the underlying model, we prove the existence of a maximizer for a general model (under some very mild technical growth conditions imposed on U and T).

Results on the portfolio selection problem (1.3). In the finance literature, one finds several studies on the one-period model with emphasis on qualitative properties and empirical experiments; see, for instance, Benartzi and Thaler [11], Shefrin and Statman [102], Levy and Levy [80], Gomes [50] and De Giorgi and Hens [34]. A more rigorous mathematical analysis of the problem has started only recently. Bernard and Ghossoub [15], He and Zhou [57] and Pirvu and Schulze [84] study the one-period setting more thoroughly. Carassus and Rásonyi [30] examine a related problem in a multi-period setting; similarly to mixed strategies in game theory, they allow the trading strategies to depend on an external source of randomness. Bichuch and Sturm [17] study the case without distortion (i.e., $T(p) = p$) and give ad hoc assumptions under which the problem (1.3) can be solved by reducing it to a suitable concave utility maximization problem. Rásonyi and Rodrigues [90] use the law-invariance of the functional V and extend the approach of Carlier and Dana [31] to a class of incomplete markets. In Chapter IV, we follow a slightly more abstract approach and give sufficient conditions for the existence of a maximizer in terms of a closedness assumption (under weak convergence) on the set of final position generated by allowed trading.

Results on the implications for the financial market equilibrium. All existing results in the literature on this topic are based on a one-period setting. Shefrin [101] explains intuitively the effect on the equilibrium of non-concave utility functions and distortions. Polkovnichenko and Zhao [86] and Dierkes [41] consider a representative agent economy and study empirically whether a representative agent with a concave utility function and distorted beliefs can explain the observed asset prices. Xi [109] and De Giorgi and Post [35] analyze a complete market economy with finitely many agents; Xi [109] uses the concavification idea to establish the existence of an equilibrium, and De Giorgi and Post [35] derive sufficient conditions on the preferences such that the resulting *pricing density* (or *pricing kernel*) dQ/dP in equilibrium is a decreasing function of the aggregate endowment. De Giorgi et al. [36] study more thoroughly whether there exists an equilibrium; in particular, they show that an equilibrium exists if there is a continuum of agents. In Chapter V, we consider a setup similar to the one in De Giorgi and Post [35] and study whether behavioural effects can lead to a (partially) increasing relation between the pricing density dQ/dP and the aggregate endowment.

I.2 Overview of the thesis

The results are divided into four chapters, three of which correspond essentially to the articles [91], [92] and [58]. To ensure that the chapters can be read independently from each other, we have deliberately allowed some duplication of terms and ideas.

Non-concave utility maximization with a given pricing measure.

Chapter II essentially consists of the article [91]. Here we consider the demand problem (1.2) in the special case without distortion (i.e., $T(p) = p$). The preference functional V therefore reduces to an expected non-concave utility functional $V(f) = E[U(f)]$. We first show necessary and sufficient conditions for the existence of a maximizer by using Fatou’s lemma in several dimensions. In particular, this requires an adaptation to the non-concave case of the notion of asymptotic elasticity. This can be defined via the conjugate of the non-concave utility function. We then present a detailed analysis of how the optimal expected non-concave utility (indirect utility), denoted by $u(x, U)$, depends on the initial capital x . This is a comparative static analysis to investigate the sensitivity of the indirect utility with respect to the key exogenous variable “initial capital” x . The optimal expected non-concave utility is, as a function of x , again increasing and non-concave. We show that its concave envelope coincides with the optimal expected utility $u(x, U_c)$, where U_c is the concave envelope U_c of U . This explains the general relation between the non-concave problem $u(x, U)$ and the “concavified” problem $u(x, U_c)$. While $u(x, U)$ and $u(x, U_c)$ may differ in general, we are then interested in a characterization of models for which $u(x, U)$ and $u(x, U_c)$ coincide. For this, we specialize the setup to the case that the underlying probability space is atomless. Using rearrangement techniques, we show that the two value functions then coincide, and we give a characterization of the maximizer. We also show that this atomlessness assumption is essentially necessary to have $u(x, U) = u(x, U_c)$ for all non-concave utility functions U and for each pricing measure Q . We use our results to discuss explicit examples exhibiting a number of new and unexpected phenomena.

Stability of the demand problem. The analysis in Chapter II and all results in the literature are for a fixed underlying model. Since one is never exactly sure of the accuracy of a proposed model, it is important to know whether the behavioural predictions generated by a model change drastically if one slightly perturbs the model. The main purpose of Chapter III (which corresponds to the article [92]) is to study this issue in detail. Formally, we consider a sequence of models, each represented by some probability space $(\Omega^n, \mathcal{F}^n, P^n)$ and some pricing measure Q^n , and we assume that this sequence converges weakly in a suitable sense to a limit model $(\Omega^0, \mathcal{F}^0, P^0, Q^0)$. For each model, we consider the demand problem

$$v^n(x) := \sup\{V_n(f) \mid f \in L_+^0(\Omega^n, \mathcal{F}^n, P^n), E_{Q^n}[f] \leq x\}, \quad (2.1)$$

where the functional V_n is defined by (1.1). We are interested in the asymptotics of the value $v^n(x)$ and its maximizer $f^n = \arg \max V_n(f)$, and we want to compare them with the analogous quantities in the limit model.

In *concave* utility maximization, the (essentially) sufficient condition for such stability results is the weak convergence of the *pricing densities*

dQ^n/dP^n to dQ^0/dP^0 (see, for instance, He [55] and Prigent [88]). However, in our non-concave setting, we present an example of a sequence of financial markets for which (dQ^n/dP^n) converges weakly to dQ^0/dP^0 , but where the limit $\lim_{n \rightarrow \infty} v^n(x)$ and $v^0(x)$, as well as the corresponding final positions, differ substantially. We discuss these new effects in detail and give sufficient conditions to prevent such unpleasant phenomena.

In order to illustrate the main results, we provide several applications. First, we consider a sequence of binomial models approximating the Black–Scholes model; this is the typical example for the transition from discrete- to continuous-time models and allows one to determine numerically the optimal expected utility for the (computationally difficult) continuous-time models via the one for (computationally tractable) discrete-time models. We also apply our results to study the stability with respect to perturbations of a model’s parameters such as drift, volatility and time horizon.

Some examples in incomplete markets. In Chapter IV, we study the portfolio selection problem (1.3). We first give sufficient conditions for the existence of a maximizer in terms of a closedness assumption (under weak convergence) on the set of final positions generated by allowed trading. This enables us to tackle the problem in a systematic and unified way, which allows us to explain the results on the existence of a maximizer obtained previously in more specific frameworks. We also verify the closedness assumption in some models that are not covered by the existing literature thus far. While we formally only prove the existence result for the special case without distortion, we explain how the same assumption can be used to prove the existence of a maximizer for the functional defined in (1.1) with distortion. For the second part, we then restrict ourselves to the case without distortion, and we study the properties of the maximizer and the optimal expected utility more thoroughly. We start with models on a finite probability space in order to bring out the intuition and structure, eliminating the need for technical complexities. We show that the optimal final position satisfies, as in the classical concave case, the first order condition for optimality in the sense that the marginal utility of the optimal position defines (up to a constant) a local martingale measure. But we also illustrate with a counterexample that the classical interpretation of the resulting martingale measure as *least favourable completion* does not carry over to the non-concave case.

While the optimal expected utility is in general a non-concave function of the initial capital, we next present sufficient conditions for a general model such that the optimal expected non-concave utility is concave in the initial capital. These conditions involve the *utility price* introduced in Jouini and Kallal [66] and can be seen as a natural generalization of the results from Chapter II. To round off the chapter, we illustrate the use of our results with a number of explicit examples.

Equilibrium. In Chapter V, we finally study some implications of behavioural effects on the financial market equilibrium. This chapter is based on the article [58]. The main focus lies on the relation between the pricing density dQ/dP and the aggregate endowment in the economy. In a complete market equilibrium with risk-averse agents, these two quantities are anti-comonotonic. However, this decreasing relation is often seen or claimed to be violated in empirical studies; this observation is called the *pricing kernel puzzle*. We study whether behavioural effects can explain this observation. We consider a one-period setting and show that allowing the agents to be (partially) risk-seeking (in the sense that U is (partially) convex) can lead to a non-decreasing relation between the pricing density dQ/dP and the aggregate endowment. But we also argue that this is a rather pathological phenomenon. We also analyze the effect of the distortion on that relation. For non-concave utility functions as well as for distortions, we compare the resulting pricing densities with the empirical estimates from the S&P 500.

Chapter II

Demand problem with a given pricing measure

In this chapter, which corresponds to [91], we study the demand problem (I.1.2) for a general model with fixed pricing density. We show the existence of a maximizer, discuss its properties and analyze the optimal expected non-concave utility (indirect utility).

II.1 Introduction

For an increasing and upper-semicontinuous function U on \mathbb{R}_+ satisfying a mild growth condition, we study in this chapter the demand problem

$$u(x, U) := \sup\{E[U(f)] \mid f \in C(x)\}, \quad (1.1)$$

where $C(x) := \{f \in L_+^0 \mid E_Q[f] \leq x\}$ for a (pricing) measure $Q \approx P$. As outlined in Chapter I, one can interpret problem (1.1) as abstract demand problem or as non-concave utility maximization problem in a financial market. For the latter interpretation, the function U can be seen as a (non-concave) utility function describing the preferences of an agent in a financial market who is dynamically trading in the underlying discounted assets S . The random variable f can be seen as the final position $x + \int_0^T \vartheta dS$ resulting from a strategy ϑ with initial capital x . In complete markets where there is a unique equivalent martingale measure Q for S , the elements in $C(x)$ are those non-negative positions that can be associated to a trading strategy starting from an initial capital less than or equal to x . Then (1.1) is an abstract version of the utility maximization problem of choosing a trading strategy such that the terminal value f^* of the corresponding wealth process maximizes $E[U(f)]$ over all final outcomes of competing strategies. The idea for this reduction from a dynamic to a static problem can be traced back to Pliska [85]. The key advantage of the static formulation is that it describes in a simple and transparent way the fundamental economic problem.

Moreover, it allows us to consider continuous- and discrete-time models in a unified way.

In the classical case where U is concave (which means that the agent is risk-averse) and smooth (e.g. in C^1), problem (1.1) and its solution are well known and we do not try to survey it here; see Biagini [16] for an attempt in that direction. For a non-concave utility function, problem (1.1) is more involved. There is a broad class of models in which the non-concave problem has been studied by reducing it to the classical concave case; see for instance Aumann and Perles [4], Carpenter [32], Berkelaar et al. [13], Larsen [77], Carassus and Pham [29], Rieger [94], Basak and Makarov [9] and Bichuch and Sturm [17]. At the other end of the scale, there are results on the existence of a solution in a number of (incomplete) discrete-time settings where one does not necessarily have a fixed pricing density, but the structure of the setup allows one to optimize directly over the set of strategies; see Benartzi and Thaler [11], Bernard and Ghossoub [15], He and Zhou [57] and Carassus and Rásonyi [30]. These two approaches provide fundamental results for a particular type of setting, but their economic implications are surprisingly varying.¹

The goal of this article is therefore to analyze (1.1) in a unified setting with one pricing density. From an economic point of view, this allows us to understand the behavioural predictions of non-concave parts in the utility function independently of the model-specific assumptions. In this way, we can explain and generalize in a systematic way the results obtained previously by other authors via ad hoc methods in specific frameworks. We also study an example in continuous time with jumps in the price process where the problem cannot be reduced to a concave one. We therefore not only unify but also extend the previous results in the literature. In contrast to the existing literature, we also present a detailed analysis of the value $u(x, U)$ as a function of the initial capital x . This is a comparative static analysis to investigate the sensitivity of the optimal value with respect to the key exogenous variable “initial capital”.

We present three main results. We first show necessary and sufficient conditions for the existence of a maximizer for $u(x, U)$, by using *Fatou’s lemma in several dimensions* due to Balder [6]; we also describe several fundamental properties of the maximizers. In particular, this requires the notion of asymptotic elasticity adapted to the non-concave case. We then study the value function $u(x, U)$. This is again an increasing and non-concave function, and its concave envelope is shown to coincide with $u(x, U_c)$, which explains

¹One example is the relation between the optimal final position and the pricing density. While they are anti-comonotonic in the typical complete models in continuous time with continuous paths (Theorem B.1 in Jin and Zhou [64]), they are not necessarily anti-comonotonic in complete models in discrete time. Non-concave utility functions might thus explain the *pricing kernel puzzle* in discrete time, but cannot do so in continuous time (see Section V.6 for a detailed discussion).

the general relation between the non-concave problem $u(x, U)$ and the concavified problem $u(x, U_c)$. In the third part, we specialize the setup to the case that the underlying probability space is atomless. Using rearrangement techniques, we show that $u(x, U)$ and $u(x, U_c)$ then coincide and we give a characterization of the maximizer for $u(x, U)$. We use our results to discuss explicit examples exhibiting a number of new and unexpected phenomena.

Besides the articles already mentioned above, there is another important branch of the literature that deals with non-concave problems. For more general preferences than expected utility, Jin and Zhou [64], Carlier and Dana [31] and He and Zhou [56] have developed (under the assumption that the distribution of dQ/dP is continuous) an approach via quantiles to obtain the existence (and in some cases also the structure) of a maximizer. This approach provides interesting new results for a specific setting, but the techniques do not apply to our general unified setup.

This chapter is structured as follows. We start in Section II.2 with a precise definition of the non-concave utility function and its concave envelope. We then formulate the optimization problem and give an illustrative example which provides some basic intuition. In Section II.3, we prove the existence and several properties of a maximizer. The value function is analyzed in Section II.4. Section II.5 contains a detailed analysis of the case that the underlying probability space is atomless. Finally, the appendices contain a number of well-known results from convex analysis and non-smooth utility maximization which are used in the body of the text.

II.2 Problem formulation and intuition

This section introduces the non-concave utility function, describes the framework and formulates the optimization problem we are interested in. For a probability space (Ω, \mathcal{F}, P) , let $L^0(\Omega, \mathcal{F}, P)$ (and $L^1(\Omega, \mathcal{F}, P)$) be the space of (equivalence classes of) \mathcal{F} -measurable (and integrable) random variables. The space $L^0_+(\Omega, \mathcal{F}, P)$ (and $L^1_+(\Omega, \mathcal{F}, P)$) consists of all non-negative elements of $L^0(\Omega, \mathcal{F}, P)$ (and $L^1(\Omega, \mathcal{F}, P)$). Sometimes, we drop the dependence on the probability space if it is clear from the context. For a random variable f , we use $f^\pm = \max(\pm f, 0)$ to denote the positive and negative parts of f .

Definition 2.1. A *non-concave utility* is a function $U : (0, \infty) \rightarrow \mathbb{R}$ with $U(\infty) > 0$, which is non-constant, increasing, upper-semicontinuous and satisfies the growth condition

$$\lim_{x \rightarrow \infty} \frac{U(x)}{x} = 0. \quad (2.1)$$

We only consider non-concave utility functions defined on the positive axis. To avoid any ambiguity, we set $U(x) = -\infty$ for $x < 0$ and define $U(0) := \lim_{x \searrow 0} U(x)$ and $U(\infty) := \lim_{x \nearrow \infty} U(x)$. Note that we do not

assume that U is concave, continuous or strictly increasing. In particular, this also allows us to analyze goal-reaching problems initiated by Kulldorff [76]. In the concave case, condition (2.1) is equivalent to the Inada condition at ∞ that $U'(\infty) = 0$. The assumption $U(\infty) > 0$ is technical but completely harmless, because adding a constant to U does not change the preferences described by U .

Definition 2.2. The *concave envelope* of U is the smallest concave function $U_c : \mathbb{R} \rightarrow \mathbb{R} \cup \{-\infty\}$ such that $U_c(x) \geq U(x)$ holds for all $x \in \mathbb{R}$.

In order to formulate the optimization problem, we fix a probability space (Ω, \mathcal{F}, P) and consider a probability measure Q equivalent to P with density $\varphi := dQ/dP$. We refer to Q as *pricing measure* and to φ as *pricing density* (or *pricing kernel*). The set $C(x)$ is defined by

$$C(x) := \{f \in L_+^0 \mid E_Q[f] \leq x\}.$$

In this chapter, we study the problem

$$u(x, U) := \sup\{E[U(f)] \mid f \in C(x)\}, \quad (2.2)$$

where we define $E[U(f)] := -\infty$ if $U(f)^- \notin L^1$. To exclude the trivial case we impose

Assumption 2.3. $u(x, U) < \infty$ for some $x > 0$.

An element $f \in C(x)$ is *optimal* if $E[U(f)] = u(x, U)$. By a *maximizer* for $u(x, U)$, we mean an optimal element for the optimization problem (2.2). The function $u(\cdot, U)$ is called the *value function* of the problem (2.2).

Remark 2.4. A financial interpretation is as follows. We fix a time horizon $T \in (0, \infty)$ and let $t=0$ be the initial time. Elements in $L^0(\Omega, \mathcal{F}, P)$ describe random payoffs at time T . The price of a discounted payoff f at time 0 is $E_Q[f]$. Problem (2.2) can then be seen as the (non-concave) utility maximization problem faced by an agent with initial capital less than or equal to x and preferences described by U , of choosing a final payoff f that maximizes the expected (non-concave) utility among all those payoffs whose price $E_Q[f]$ is at most x .

To explain the motivation for this static formulation, we recall the classical formulation of the utility maximization problem: Given a financial market, find a strategy ϑ whose resulting gains from trade $G_T(\vartheta)$ maximize $E[U(x + G_T(\vartheta))]$ over all admissible ϑ . Instead of maximizing over strategies, it is also possible to maximize over attainable payoffs, i.e., final positions that can be generated by suitable self-financing strategies starting at $t = 0$ from initial capital x . In the special case of a complete market, there is a unique equivalent martingale measure Q and the set of attainable (non-negative) payoffs is given by $C(x)$. In the general case, the set of attainable payoffs

is more complicated; however, in *concave* utility maximization, one central property is that the solution to the original problem can be obtained by working with only one well-chosen dual object. See Kramkov and Schachermayer [73] for details. \diamond

Remark 2.5. In behavioural finance, payoffs are often evaluated with respect to a (possibly stochastic) reference point R , which means that the agent evaluates $U(f - R)$ rather than $U(f)$. If the reference point is deterministic, this can be embedded in our analysis by defining a new utility $U_1(x) := U(x - R)$ and slightly modifying the arguments to account for the new domain of the utility. If the reference point is stochastic, one has to think carefully about the correct notion of allowed payoffs. Jin et al. [63] allow all payoffs f with $E_Q[f] \leq x$ and $f - R \geq L$ for a fixed constant L . One can then solve the problem by maximizing first $E[U(g)]$ subject to the constraint $E_Q[g] \leq x - E_Q[R]$ to get a maximizer g^* and then choose $f^* := g^* + R$. In the setting of a complete financial market, in particular, this has a clear economic interpretation: the agent uses the amount $E_Q[R]$ to replicate the stochastic reference point R and invests the remaining part $x - E_Q[R]$ as if he had a deterministic reference point $E_Q[R]$. \diamond

Remark 2.6. While several applications in finance use a non-concave utility defined on \mathbb{R}_+ , there are also interesting examples where the non-concave utility is defined on \mathbb{R} . This case is surprisingly different from our analysis. A first problem is that the concave envelope (which plays a key role in our analysis) might not be defined. The problem then might easily be *ill-posed*; see Theorem 3.2 of Jin and Zhou [64]. If, however, the concave envelope U_c is well defined, then several arguments of our analysis carry over to the non-concave utility functions defined on \mathbb{R} ; see Remark 5.4 below. \diamond

II.2.1 Examples and intuition

We present here three representative models to illustrate how the pricing density and the underlying probability space look in explicit settings.

Example 2.7 (Black–Scholes model). We fix some time horizon $T \in (0, \infty)$, a probability space (Ω, \mathcal{F}, P) on which there is a standard Brownian motion $W = (W_t)_{t \geq 0}$ and a (discounted) market consisting of a savings account $B \equiv 1$, assumed to be constant, and one stock S described by

$$dS_t = S_t(\mu_t dt + \sigma_t dW_t), \quad S_0 = s_0 > 0, \sigma > 0,$$

in the augmented filtration generated by W . The drift μ and the volatility σ are assumed to be adapted and the market price of risk $\lambda := \mu/\sigma$ is assumed to satisfy $\int_0^T \lambda_s^2 ds < \infty$ P -a.s. The unique martingale measure, if it exists, is defined by $\varphi := dQ/dP = \exp\left(-\int_0^T \lambda_s dW_s - \frac{1}{2} \int_0^T \lambda_s^2 ds\right)$.

Example 2.8 (Geometric Poisson process). We consider some time horizon $T \in (0, \infty)$ and a probability space (Ω, \mathcal{F}, P) on which there is a Poisson process $N = (N_t)_{t \geq 0}$ with constant intensity $\lambda > 0$. Let $(\mathcal{F}_t)_{0 \leq t \leq T}$ be the augmented filtration generated by N . We consider a (discounted) market consisting of a savings account $B \equiv 1$ and one risky stock S described by

$$dS_t = \alpha S_t dt + \sigma S_{t-} d\tilde{N}_t, \quad S_0 = s_0 > 0, \sigma > -1, \sigma \neq 0, \alpha/\sigma < \lambda,$$

where $\tilde{N}_t := N_t - \lambda t$ is the compensated Poisson process. The unique martingale measure is defined by $\varphi := dQ/dP = e^{(\lambda - \tilde{\lambda})T} \left(\frac{\tilde{\lambda}}{\lambda}\right)^{N_T}$ for $\tilde{\lambda} := \lambda - \alpha/\sigma$.

Example 2.9 (Complete model in discrete time). The classical example in discrete time is the binomial model consisting of a savings account $B \equiv 1$, assumed to be constant, and one stock S described by

$$\frac{S_k^1}{S_{k-1}^1} = Y_k = \begin{cases} 1 + u & \text{with probability } p \\ 1 + d & \text{with probability } 1 - p \end{cases}$$

for $k = 1, \dots, T$ in the filtration generated by Y . For $u > 0 > d > -1$, this model is complete; the underlying probability space consists of finitely many atoms.

The utility functions studied in this chapter are non-concave. In order to provide some intuition for these non-classical utility functions, we start with a motivating example in which we compare Example 2.7 for $\mu = 0$ and $T = 1$ and Example 2.9 for $u = -d < 1$ and $T = 1$. In both cases, it follows that $\varphi = 1$; in the first case, the underlying probability space is atomless; in the second case, the underlying probability space consists of two atoms. Example 2.10 shows that the two optimization problems are fundamentally different even though the pricing densities are equal: the underlying probability space crucially affects the optimization problem $u(x, U)$. Example 2.10 is of course pathological, but it is nevertheless the simplest possible setting demonstrating both intuition and structure and abstracting them from technical complexities.

Example 2.10. We consider the non-concave utility function U defined by

$$U(x) := \begin{cases} 0, & x < 1, \\ 1, & x \geq 1. \end{cases}$$

This function is increasing and upper-semicontinuous; the concave envelope is given by $U_c(x) = x$ on $(0, 1)$ and $U_c(x) = 1$ for $x \geq 1$; the optimization problem is therefore to maximize $P[f \geq 1]$. Moreover, we set $\varphi \equiv 1$.

For the (concave) utility maximization problem $u(x_0, U_c)$, the random variable $f \equiv x_0$ is optimal (as can be seen by Jensen's inequality, because $\varphi \equiv 1$) and gives $u(x_0, U_c) = U_c(x_0) = x_0$. The maximizer for $u(x_0, U_c)$ and

the value function $u(x_0, U_c)$ are independent of the underlying probabilistic structure.

For $u(x_0, U)$ and $x_0 \in (0, 1)$, the payoff $f \equiv x_0$ is not optimal; it is better to have 1 with probability x_0 while having 0 with probability $1 - x_0$. If the probability space is atomless, one can generate any probability distribution; this allows us to choose $A \in \mathcal{F}$ with $P[A] = x_0$. The element $f = 1_A$ gives the expected utility $U(0)P[A^c] + U(1)P[A] = x_0$ which is equal to $u(x_0, U_c)$. This shows that f is also a maximizer for $u(x_0, U)$, since $U_c \geq U$.

To illustrate the other extreme case, take $\Omega := \{\omega_1, \omega_2\}$, $\mathcal{F} := 2^\Omega$ and $P[\{\omega_1\}] := P[\{\omega_2\}] := \frac{1}{2}$. In this model, a set A with $P[A] = x \in (0, 1) \setminus \{\frac{1}{2}\}$ cannot be generated since the elements in $L_+^0(\Omega, \mathcal{F})$ take at most two values, each with probability $1/2$, and so $u(x, U) = 1_{\{x \in [1/2, 1)\}}1/2 + 1_{\{x \geq 1\}}$.

The important insight from this example is that the underlying probability space crucially affects the optimization problem $u(x, U)$. In the atomic case, the value function $u(x, U)$ is non-concave and (for some x) strictly below $u(x, U_c)$. In the atomless case, the value function $u(x, U)$ is concave and equal to $u(x, U_c)$. Note that this has nothing to do with whether or not the pricing density φ has a continuous distribution.

II.2.2 Some tools from convex analysis

Example 2.10 already highlights the importance of the concave envelope U_c . This section summarizes several results which will be used throughout the chapter. The proofs can be found in Appendix II.7.

Lemma 2.11. *The concave envelope U_c of U is finite and continuous on $(0, \infty)$. The set $\{U < U_c\} := \{x \in \mathbb{R}_+ \mid U(x) < U_c(x)\}$ is open and can be written as a countable union of finite disjoint open intervals. Moreover, U_c is locally affine on the set $\{U < U_c\}$, in the sense that it is affine on each of the above intervals.*

A key tool to study the relation between U and U_c is the conjugate of U defined by

$$J(y) := \sup_{x>0} \{U(x) - xy\}.$$

Because of the non-concavity of U , the conjugate J is no longer smooth; we therefore work with the subdifferential which is denoted by ∂J for the convex function J and by ∂U_c for the concave function U_c (for precise definitions, see Appendix II.7). The right- and left-hand derivatives of J are denoted by J'_+ and J'_- . The next lemma summarizes several properties, whose proof is deferred to Appendix II.7.

Lemma 2.12. *The function J is convex, decreasing, finite on $(0, \infty)$ and satisfies $J(x) = \infty$ for $x < 0$. The non-concave utility U and its concave envelope U_c have the same conjugate. Moreover, it holds that*

$$U_c(x) - xy = J(y) \iff x \in -\partial J(y) \iff y \in \partial U_c(x). \quad (2.3)$$

In classical concave utility maximization, the asymptotic elasticity (AE) of the utility function is important. In particular, many results impose an upper bound on $AE(U)$. Deelstra et al. [37] showed that this condition may also be formulated for the asymptotic elasticity of the conjugate. This turns out to be useful for extending asymptotic elasticity to the non-concave case. We define

$$AE_0(J) := \limsup_{y \rightarrow 0} \sup_{q \in \partial J(y)} \frac{|q| y}{J(y)}.$$

The condition $AE_0(J) < \infty$ has useful implications, which are summarized in the next lemma. The equivalence is proved in Lemma 4.1 of Deelstra et al. [37]. The proof of the second part is presented in Appendix II.7.

Lemma 2.13. *The asymptotic elasticity condition $AE_0(J) < \infty$ is equivalent to the existence of two constants $\gamma > 0$ and $y_0 > 0$ such that*

$$J(\mu y) \leq \mu^{-\gamma} J(y) \quad \text{for all } \mu \in (0, 1] \text{ and } y \in (0, y_0].$$

Moreover, if $AE_0(J) < \infty$ is satisfied, then there are $x_0 > 0$ and $k > 0$ such that $0 \leq U_c(x) \leq kU(x)$ on (x_0, ∞) .

In most applications, the non-concave utility U is strictly concave and differentiable on (x, ∞) for x large enough. For such utilities, the required asymptotic elasticity condition could also be formulated in the classical way in terms of $AE(U)$ or $AE(U_c)$, and the classical interpretation of the condition via risk aversion then still applies.

II.3 Existence and properties of a maximizer

The goal of this section is to prove the existence of a maximizer for $u(x, U)$. We also discuss several properties such as uniqueness, first order condition for optimality, and the relation between the maximizer and the pricing density.

Recall that the only assumption we made so far is Assumption 2.3 that $u(x, U) < \infty$ for some $x > 0$. Even in the case of concave utilities, this is not sufficient to guarantee the existence of a maximizer; see Examples 2.3 and 3.1 in Jin et al. [62]. We introduce a stronger assumption, see Remark 3.3 below.

Assumption 3.1. $E[J(\lambda\varphi)] < \infty$ for all $\lambda > 0$.

In the case that $U = U_c$, sufficient conditions for this assumption are $u(x, U_c) < \infty$ for some $x > 0$ and $AE_0(J) < \infty$. The next lemma extends this to the non-concave case.

Lemma 3.2. $AE_0(J) < \infty$ and Assumption 2.3 imply Assumption 3.1.

Proof. It is known from *concave* utility maximization that $AE_0(J) < \infty$ and $u(x, U_c) < \infty$ for some $x > 0$ imply Assumption 3.1 (see for instance Lemma 5.4 of Westray and Zheng [107]). It is therefore sufficient to show that $AE_0(J) < \infty$ and Assumption 2.3 imply $u(x, U_c) < \infty$ for some $x > 0$. For this, let $x_0 > 0$ and k be as given in the second part of Lemma 2.13 so that

$$0 \leq U_c(x) \leq kU(x) \quad \text{on } (x_0, \infty). \quad (3.1)$$

If U is positive on $(0, \infty)$, then fix some $f \in C(x)$ and apply (3.1) on the set $\{f > x_0\}$. This gives $E[U_c(f)] \leq U_c(x_0) + kE[U(f)]$ and taking the supremum over all $f \in C(x)$ implies $u(x, U_c) \leq U_c(x_0) + ku(x, U)$. So if $u(x, U)$ is finite for some $x > 0$, then $u(x, U_c)$ is also finite.

If $U(0) < 0$, we choose ϵ small enough such that $x - \epsilon > 0$, fix $f \in C(x - \epsilon)$ and apply the above argument to $f_\epsilon := f + \epsilon$ and $U(x) - U(\epsilon)$. This gives

$$\begin{aligned} E[U_c(f)] &\leq E[U_c(f_\epsilon)] \\ &\leq U_c(x_0) + kE[(U(f_\epsilon) - U(\epsilon))1_{\{f_\epsilon \geq x_0\}}] + kE[U(\epsilon)1_{\{f_\epsilon \geq x_0\}}] \\ &\leq U_c(x_0) + kE[U(f_\epsilon) - U(\epsilon)] + kE[U(\epsilon)1_{\{f_\epsilon \geq x_0\}}] \\ &\leq U_c(x_0) + ku(x, U) + k|U(\epsilon)|, \end{aligned}$$

where $f_\epsilon \in C(x)$ is used in the last step. Taking the supremum over all $f \in C(x - \epsilon)$ then gives $u(x - \epsilon, U_c) \leq ku(x, U) + k$ for some constant k and the result follows. \square

Remark 3.3. Conversely, Assumption 3.1 also implies Assumption 2.3 since $E[U(f)] \leq E[U_c(f)] \leq E[J(\lambda\varphi)] + \lambda x_0$ for $f \in C(x_0)$ and $\lambda > 0$. \diamond

We now turn to the main result of this section. In Section IV.2.2 and IV.5, we prove a similar result for the preference functional with distortion. We will discuss there also the relation to the present theorem.

Theorem 3.4. *Suppose that Assumption 3.1 is satisfied. For all $x_0 \in (0, \infty)$, there exists some $\hat{f} \in C(x_0)$ such that $u(x_0, U) = E[U(\hat{f})]$.*

Because of the non-concavity of U , proving the existence in the usual way via a dual problem is not possible. We work directly with a maximizing sequence (f^n) and try to get a suitable subsequence. Due to the lack of concavity of U , Komlós-type arguments (see, for instance, Lemma 5.5 in Westray and Zheng [107]) also do not work. Instead, we consider the sequence $(-U(f^n), \varphi f^n)$ and look, in the spirit of Fatou's lemma, for some two-dimensional "lim inf" (\hat{g}_1, \hat{g}_2) . The random variable $\hat{f} = \hat{g}_2/\varphi$ is then a natural candidate for a maximizer. This idea and the kind of Fatou-type result in several dimensions needed to get a two-dimensional "lim inf" already exist in the literature. They were originally developed to prove the existence of equilibria for economies with a measure space of economic agents.

Balder and Pistorius [7] also apply them to demand problems. The condition needed for this approach is a lower closure result for the sequence $(-U(f^n), \varphi f^n)$. This boils down to showing, as in the concave case, that $(U(f^n)^+)$ is uniformly integrable. For later applications, we prove a more general statement.

Proposition 3.5. *Suppose that Assumption 3.1 is satisfied. Consider a sequence $x_n \rightarrow x > 0$ and a sequence (f^n) with $f^n \in C(x_n)$. There is some $\hat{f} \in C(x)$ such that $\limsup_{n \rightarrow \infty} E[U(f^n)] \leq E[U(\hat{f})]$.*

Proof. 1) In the first part, we show that the family $(U(f^n)^+)_{n \in \mathbb{N}}$ is uniformly integrable. If U is bounded from above, this is clear. Hence, we may assume $U(\infty) = \infty$. The sequence (x_n) is bounded by x_0 , say, and it follows that

$$\begin{aligned} E[U(f^n)^+ 1_{\{U(f^n)^+ > \alpha\}}] &\leq E[(J(\lambda\varphi)^+ + f^n\varphi\lambda) 1_{\{U(f^n)^+ > \alpha\}}] \\ &\leq E[J(\lambda\varphi)^+ 1_{\{U(f^n)^+ > \alpha\}}] + \lambda x_0. \end{aligned}$$

Therefore, it is sufficient to show that for any $\lambda > 0$, we have

$$\lim_{\alpha \rightarrow \infty} \sup_n E[J(\lambda\varphi)^+ 1_{\{U(f^n)^+ > \alpha\}}] = 0.$$

Since the single random variable $J(\lambda\varphi)^+ \in L^1$ is trivially uniformly integrable, we only need to show that $\sup_n P[U(f^n)^+ > \alpha] \rightarrow 0$ for $\alpha \rightarrow \infty$. For this, fix a sequence $\alpha_i \rightarrow \infty$ and let $\tilde{x} := \inf\{x > 0 \mid U(x) > 0\}$ denote the first point where the utility becomes positive. By the definition of \tilde{x} and x_0 , we have $U(f^n)^+ \leq U(\tilde{x} + f^n)$ and $f^n + \tilde{x} \in C(x_0 + \tilde{x})$. And since we have by Remark 3.3 that $u(x, U_c) < \infty$ for all $x > 0$, we get

$$\sup_n P[U(f^n)^+ > \alpha_i] \leq \sup_n \frac{E[U(f^n)^+]}{\alpha_i} \leq \frac{u(x_0 + \tilde{x}, U)}{\alpha_i} \rightarrow 0,$$

which completes the proof of the first part.

2) By passing to a subsequence that realizes the lim sup, we can assume that the sequence $(E[U(f^n)])$ converges, and we denote the limit by γ_1 . In the same way, by passing to a further subsequence (again, relabelled as f^n) that realizes $\limsup E[\varphi f^n]$, we can also assume that $(E[\varphi f^n])$ converges to some $\gamma_2 \leq x$. Define (g^n) by $g^n := (-U(f^n), \varphi f^n)$ and denote the i -th component of g^n by g_i^n . The sequence $(E[U(f^n)])$ converges to γ_1 , hence

$$\lim_{n \rightarrow \infty} E[g_1^n] = - \lim_{n \rightarrow \infty} E[U(f^n)] = -\gamma_1,$$

and $\lim_{n \rightarrow \infty} E[g_2^n] = \gamma_2$. Since $g_2^n \geq 0$ P -a.s., the family $(\max(0, -g_2^n))$ is trivially uniformly integrable. Moreover, we have shown in 1) that the sequence $(\max(0, -g_1^n))$ is uniformly integrable. By Corollary 3.9 of Balder [6], it follows that there exists $\hat{g} \in L^1(\Omega; \mathbb{R}^2)$ such that

$$E[\hat{g}_1] \leq -\gamma_1 \quad \text{and} \quad E[\hat{g}_2] \leq \gamma_2, \quad (3.2)$$

and for a.e. $\omega \in \Omega$, there exists a subsequence $n_k(\omega)$ such that

$$\lim_{k \rightarrow \infty} g_1^{n_k(\omega)}(\omega) = \hat{g}_1(\omega) \quad \text{and} \quad \lim_{k \rightarrow \infty} g_2^{n_k(\omega)}(\omega) = \hat{g}_2(\omega). \quad (3.3)$$

Since $\varphi > 0$ P -a.s., we can define $\hat{f}(\omega) := \hat{g}_2(\omega)/\varphi(\omega)$. This gives $\hat{f}\varphi = \hat{g}_2$ and together with (3.2), it follows that $E[\varphi\hat{f}] = E[\hat{g}_2] \leq \gamma_2 \leq x$, which means that $\hat{f} \in C(x)$. Moreover, it follows from (3.3) that

$$\lim_{k \rightarrow \infty} f^{n_k(\omega)}(\omega) = \lim_{k \rightarrow \infty} \frac{g_2^{n_k(\omega)}(\omega)}{\varphi(\omega)} = \frac{\hat{g}_2(\omega)}{\varphi(\omega)} = \hat{f}(\omega).$$

Together with upper-semicontinuity of U and (3.3) we obtain

$$U(\hat{f}(\omega)) \geq \lim_{k \rightarrow \infty} U(f^{n_k(\omega)}(\omega)) = - \lim_{k \rightarrow \infty} g_1^{n_k(\omega)}(\omega) = -\hat{g}_1(\omega).$$

Taking expectations and using (3.2) gives $E[U(\hat{f})] = E[-\hat{g}_1] \geq \gamma_1$. \square

The proof of Theorem 3.4 is now a direct application of Proposition 3.5.

Proof of Theorem 3.4. Consider a maximizing sequence (f^n) in $C(x_0)$ which also satisfies $E[\varphi f^n] = x_0$ for each $n \in \mathbb{N}$. Proposition 3.5 gives some $\hat{f} \in C(x_0)$ such that $E[U(\hat{f})] \geq \lim_n E[U(f^n)] = u(x_0, U)$, which shows that \hat{f} is a maximizer. \square

Remark 3.6. Balder and Pistorius [7] provide an existence result for a multi-good consumption problem with a non-concave utility function on \mathbb{R}_+^m . They impose on the utility function a growth condition that also involves the pricing density. By exploiting convex duality, our Assumption 3.1 relaxes their conditions and highlights that the classical assumptions via asymptotic elasticity are sufficient for the existence of a maximizer also in the non-concave case with one pricing measure. \diamond

Having clarified the existence of a solution to $u(x, U)$, it is natural to ask about uniqueness. Here the answer is negative: *A maximizer for $u(x, U)$ is not necessarily unique.* Similarly to the case of concave utilities which are not strictly concave, we can manipulate the solution on those parts of $\text{dom}(U)$ where U_c is locally affine. This is illustrated in the next example.

Example 3.7. Take a sufficiently rich model, $\varphi \equiv 1$ and a utility function satisfying $\{U < U_c\} = (a, b) \cup (b, c)$. Jensen's inequality shows that $f \equiv b$ is a maximizer for $u(b, U)$. But on the other hand, for a set $A \in \mathcal{F}$ satisfying $P[A] = (c-b)/(c-a)$ (which exists by the richness of the model), it follows that $a1_A + c1_{A^c}$ is in $C(b)$ and satisfies $E[U(a1_A + c1_{A^c})] = U(b)$, which means that $f' = a1_A + c1_{A^c}$ is also a maximizer. Lemma 5.9 below describes model classes where the solution for $u(x, U)$ is unique.

For the economic interpretation, the mere existence result is not very satisfying. In the sequel, our goal is therefore to describe the properties of maximizers in more detail. We start with the first order condition for optimality that the marginal utility $U'(f^*)$ of a maximizer f^* is proportional to the pricing density. In the classical case when U is concave, this property of the optimizer is a by-product of the convex duality arguments. But in the non-concave case, convex duality cannot be applied. Instead, we use a standard marginal variation argument. In order to avoid dealing with generalized derivatives for non-concave functions, we impose here a slightly stronger assumption on U . Corollary 5.5 below gives a more general statement about the first order condition for atomless models.

Lemma 3.8. *Let U be continuously differentiable satisfying the Inada condition $\lim_{x \nearrow \infty} U'(x) = 0$ at ∞^2 and let f^* be a maximizer for $u(x, U)$. Then there is $\lambda > 0$ such that f^* satisfies $U'(f^*) = \lambda\varphi$ on $\{f^* > 0\}$.*

Proof. Fix $\epsilon > 0$, define $A_\epsilon := \{f^* > \epsilon\}$ and consider

$$\max E[U(f)1_{A_\epsilon}] \quad \text{subject to } f \in L^0_+, E[f\varphi 1_{A_\epsilon}] \leq E[f^*\varphi 1_{A_\epsilon}]. \quad (3.4)$$

If there is some \hat{f} in (3.4) with $E[U(\hat{f})1_{A_\epsilon}] > E[U(f^*)1_{A_\epsilon}]$, the candidate $f' := \hat{f}1_{A_\epsilon} + f^*1_{A_\epsilon^c}$ is feasible for the problem $u(x, U)$ and satisfies $E[U(f')] > E[U(f^*)]$ which contradicts the optimality of f^* . Hence f^* also solves (3.4). Now fix some uniformly bounded $f \in L^0$ and define $f_\lambda := f^* + \lambda(f - c)1_{A_\epsilon}$ for $c = E[\varphi f 1_{A_\epsilon}]/E[\varphi 1_{A_\epsilon}]$. First, note that

$$E[\varphi f_\lambda 1_{A_\epsilon}] = E[\varphi f^* 1_{A_\epsilon}] + \lambda(E[\varphi f 1_{A_\epsilon}] - cE[\varphi 1_{A_\epsilon}]) = E[\varphi f^* 1_{A_\epsilon}]$$

holds for every λ . Moreover, since f is uniformly bounded, it follows that $f_\lambda \geq 0$ on A_ϵ for λ small. Thus f_λ is a feasible candidate for (3.4) and this yields

$$0 \geq \limsup_{\lambda \rightarrow 0} \frac{E[(U(f_\lambda) - U(f^*))1_{A_\epsilon}]}{\lambda}.$$

Since U is continuously differentiable and satisfies the Inada condition at ∞ , the derivative U' is bounded on (ϵ, ∞) , and using the mean value theorem, we find that $(U(f_\lambda) - U(f^*))/\lambda$ is bounded by a constant on A_ϵ . Hence we may interchange limit and expectation to obtain $0 \geq E[U'(f^*)(f - c)1_{A_\epsilon}]$. Replacing f by $-f$ shows that the expectation must vanish. Using the notation $\gamma := E[U'(f^*)1_{A_\epsilon}]/E[\varphi 1_{A_\epsilon}]$, we thus see that $E[(U'(f^*) - \varphi\gamma)f 1_{A_\epsilon}] = 0$ holds for all $f \in L^\infty$. This implies $U'(f^*) = \varphi\gamma$ on A_ϵ . The same approach for $\tilde{\epsilon} \in (0, \epsilon)$ gives $U'(f^*) = \varphi\tilde{\gamma}$ on $A_{\tilde{\epsilon}}$ for some constant $\tilde{\gamma}$. Since $A_\epsilon \subset A_{\tilde{\epsilon}}$, we have that $U'(f^*) = \varphi\gamma = \varphi\tilde{\gamma}$ on A_ϵ and we infer $\gamma = \tilde{\gamma}$. This can be done for any $\tilde{\epsilon} > 0$ and we obtain $U'(f^*) = \varphi\gamma$ on $\bigcup_{\epsilon > 0} A_\epsilon = \{f^* > 0\}$, which proves the assertion. \square

²One sufficient condition for the Inada condition is that U satisfies (2.1) and $\{U < U_c\} = \bigcup_{i=1}^n (a_i, b_i)$ for some finite n .

In financial economics, there is a broad strand of literature (see for instance Jackwerth [59] and Beare [10]) that analyzes the pricing density in financial markets. From an equilibrium perspective, this boils down to an analysis of the relation between the pricing density dQ/dP and the optimal final position f^* of a single agent. In the classical case where U is concave, the marginal utility U' is decreasing and the first order condition $U'(f^*) = \lambda dQ/dP$ hence gives a decreasing relation between f^* and dQ/dP . This means that f^* and dQ/dP are *anti-comonotonic* (see Definition 4.82 in Föllmer and Schied [43]). Dybvig [42] considers a finite-dimensional setup and shows, by using rearrangement techniques, that the decreasing relation between f^* and dQ/dP also holds for non-concave utility functions if all states $\omega \in \Omega$ have the same probability. The economic intuition for this result is that it is (cost-)efficient to have a higher payoff in those states where the pricing density is low. Jin and Zhou [64] and Carlier and Dana [31] show that the decreasing relation between the optimal final position and dQ/dP also holds for more general preferences if the pricing density has a continuous distribution. However, this result cannot be generalized to an arbitrary combination of pricing density and underlying probability space as we illustrate in the next example.

Example 3.9. Take $\Omega = \{\omega_1, \omega_2\}$, $\mathcal{F} = 2^\Omega$, $P[\{\omega_1\}] = 2/3$, define the pricing measure by $Q[\{\omega_1\}] = 3/4$ and consider the non-concave utility function

$$U(x) := \begin{cases} (x-1)^{1/3}, & x \geq 1, \\ -(1-x)^{2/3}, & x < 1. \end{cases}$$

The problem $u(1, U)$ can be solved explicitly for this function U (see Theorem 3.1 in Bernard and Ghossoub [15]), and it turns out that f^* defined by $f^*(\omega_1) := 10/9$ and $f^*(\omega_2) := 2/3$ is the optimal final position for $u(1, U)$. It follows that f^* and φ are not anti-comonotonic.

Remark 3.10. 1) By extending, for the case of expected non-concave utility, the results of Jin and Zhou [64] and Carlier and Dana [31], it is shown below in Corollary 5.6 that φ and the maximizer f^* for $u(x, U)$ are anti-comonotonic if the probability space is atomless.

2) At first glance, Example 3.9 seems to contradict the main result of Rieger [93] who claims to generalize the above result in Dybvig [42] to a general pricing density. However, a closer inspection of the setup and the results in Rieger [93] shows that the optimization problem (Definitions 2.10 and 2.11 and Theorem 2.12 there) is not of the same classical form as our problem (2.2). Rieger [93] considers probability distributions where the pricing density φ and the final payoff f may have any joint distribution (in contrast to our setup, where only those joint distributions are allowed which can be actually supported by the underlying probability space), and the constraint is formulated in terms of the covariance between the payoff and the pricing

density. In the context of Example 3.9, this allows one to choose a payoff f via the joint distribution

$$P[\varphi = 3/4, f = 10/9] = P[\varphi = 9/8, f = 10/9] = P[\varphi = 9/8, f = 2/3] = 1/3,$$

which is not possible in our framework.

One can think of the approach in Rieger [93] as allowing for *randomized* payoffs. In the above example, instead of a fixed payoff $f(\omega_2)$, one is allowed to choose a lottery with outcomes 10/9 and 2/3 with probability 0.5 each. For mathematical purposes, having randomized payoffs is very useful since it allows one to prove anti-comonotonicity of f^* and φ . But if the payoffs must be generated by trading in a specified financial market, the underlying probabilistic structure matters, and allowing randomized payoffs does not match up well with the given financial problem.

3) Instead of the (primal) problem to maximize the expected (non-concave) utility, one can also look at a dual problem of minimizing over payoffs f the cost $E_Q[f]$ for a given value of $E[U(f)]$. Since the expected non-concave utility depends only on the distribution of f , this problem is closely related to finding the cheapest price to generate a given distribution. In Dybvig [42], this is called the *distributional price* of f or of its distribution, and a payoff is called *cost-efficient* if any other payoff that generates the same distribution costs at least as much. The ideas of Dybvig [42] have recently been revisited, formalized and extended in Bernard and Boyle [14]. One central result there is to show that a payoff is cost-efficient if the payoff and the pricing density are anti-comonotonic. Example 3.9 above shows that the converse direction is not true in general: A pricing density φ and a cost-efficient payoff are not necessarily anti-comonotonic. \diamond

We conclude this section with one clarification. Throughout, we have focused on the problem $u(x, U)$ for a generic $x > 0$. For a particular x , the arguments might simplify drastically because one might happen to have more structure. As an example, fix $\lambda > 0$ and $A \in \mathcal{F}$ and consider $x := E[\varphi(-J'_-(\lambda\varphi)1_A - J'_+(\lambda\varphi)1_{A^c})]$. It then follows from classical convex duality arguments that $f^* := -J'_-(\lambda\varphi)1_A - J'_+(\lambda\varphi)1_{A^c}$ is a maximizer for $u(x, U_c)$, and Lemma 7.2 shows that $P[f^* \in \{U < U_c\}] = 0$. This implies that f^* is a maximizer for $u(x, U)$. In this way, the existence and several properties of the maximizer follow here directly from the *concave* utility maximization. But for arbitrary $x > 0$, this simplification is not possible.

II.4 The value function

In Section II.3, we have considered $u(x, U)$ for a fixed $x > 0$. This section analyzes how the optimal value $u(x, U)$ depends on x . Not surprisingly, the value function $u(x, U)$ is again increasing, upper-semicontinuous and non-concave (upper-semicontinuity is proved below in Proposition 4.2). So the

value function itself has similar properties as the utility U and, as in the case of U , we are interested in the concave envelope of $u(x, U)$. The first important result is that the concave envelope of $u(x, U)$ is (essentially) given by $u(x, U_c)$. This gives a general relation between $u(x, U)$ and $u(x, U_c)$.

Theorem 4.1. *Take any non-concave utility function U with conjugate J . If Assumption 3.1 is satisfied, then the value function $u(x, U_c)$ and the concave envelope of $u(x, U)$ coincide on $(0, \infty)$ and we have*

$$j(\lambda) := E[J(\lambda\varphi)] = \sup_{x>0} \{u(x, U) - x\lambda\}, \quad \lambda > 0. \quad (4.1)$$

The relation (4.1) states that j is the conjugate of the value function $u(x, U)$. In the *concave* case, the corresponding result from the literature is more general since j there is defined as an infimum (over pricing measures) in a dual problem; see for instance Lemma 3.4 of Kramkov and Schachermayer [73]. However, a closer inspection of that proof there shows that concavity is not needed for our present case with a single pricing measure. More precisely, Kramkov and Schachermayer [73] (or Westray and Zheng [107] for the nonsmooth case) determine the conjugate of the value function in a setting with infinitely many pricing measures. Then concavity of the utility function U is (essentially) used to apply the minimax theorem in order to exchange the supremum over (bounded) payoffs and the infimum over pricing densities. In the present case with one single pricing density, this step is not necessary and one can proceed without concavity. For completeness, we carry out the details.

Proof of Theorem 4.1. 1) The main part of the proof is to show (4.1). Once we have that, we note from Lemma 2.12 that U and U_c have the same conjugate J so that applying (4.1) for U and U_c implies that j is the conjugate of both $u(x, U)$ and $u(x, U_c)$. By the classical concave utility theory and the assumption that $u(x_0, U_c) < \infty$, the function $x \mapsto u(x, U_c)$ is finite and concave on $(0, \infty)$, hence continuous. Applying part (iii) of Lemma 7.1 to $f_1(x) := -u(-x, U_c)$ then gives that $f_1^{**} = f_1$ on $(-\infty, 0)$. On the other hand, $f_2(x) := -u(-x, U)$ is proper on $(-\infty, 0)$ because $U(x) \leq u(x, U) \leq u(x, U_c) < \infty$ for all $x > 0$, and we have argued above that $f_2^* = f_1^*$ due to (4.1) so that $f_2^{**} = f_1^{**} = f_1$ on $(-\infty, 0)$. Moreover, the convex envelope \bar{f}_2 of f_2 is also proper on $(-\infty, 0)$ since $\bar{f}_2 \geq f_1$, and \bar{f}_2 is continuous on $(-\infty, 0)$ since it is convex and finite like f_1 and U . So applying part (iii) of Lemma 7.1 now to f_2 (which is lower-semicontinuous due to Proposition 4.2 below) finally gives $\bar{f}_2 = f_2^{**} = f_1$ on $(-\infty, 0)$, as claimed.

2) To show (4.1), we start with the claim that

$$\lim_{n \rightarrow \infty} \sup_{f \in C_n} E[U(f) - f\lambda\varphi] = \sup_{x>0} \{u(x, U) - x\lambda\}, \quad (4.2)$$

where $C_n := \{f \in L_+^0 \mid 0 \leq f \leq n\}$ is the ball of radius n in the positive orthant of L^0 . For “ \leq ”, note that the left-hand side of (4.2) is an increasing limit in n ; so we have to show that for each n and each $f \in C_n$,

$$E[U(f) - f\lambda\varphi] \leq \sup_{x>0} \{u(x, U) - x\lambda\}.$$

To do that, fix f and define $x^* := E[\varphi f]$. For $x^* = 0$, we get $E[f\varphi] = 0$, hence $f \equiv 0$, and so $E[U(f) - f\lambda\varphi] = U(0) \leq u(x, U) - x\lambda + x\lambda$ for any $x > 0$. This gives the above inequality, and so we consider the case where $x^* > 0$. By the definition of x^* , we have $f \in C(x^*)$ and it follows that

$$E[U(f) - f\lambda\varphi] \leq u(x^*, U) - x^*\lambda \leq \sup_{x>0} \{u(x, U) - x\lambda\}.$$

This proves (4.2).

3) To get (4.1) from (4.2), we now want to interchange supremum and expectation on the left-hand side of (4.2) and then let $n \rightarrow \infty$. For each n , a measurable selection argument (see Theorem 18.19 in Aliprantis and Border [3]) shows that we can choose a measurable selector $x^*(\omega) \in L_+^0$ such that

$$\sup_{0 \leq x \leq n} \{U(x) - x\lambda\varphi(\omega)\} = U(x^*(\omega)) - x^*(\omega)\lambda\varphi(\omega).$$

With $J_n(\lambda) := \sup_{0 \leq x \leq n} \{U(x) - x\lambda\} \geq U(n) - n\lambda$, it thus follows that we have indeed

$$\sup_{f \in C_n} E[U(f) - f\lambda\varphi] = E[J_n(\lambda\varphi)]. \quad (4.3)$$

But J_n is increasing in n and dominated by J , so in view of (4.2) and (4.3), we have to show for (4.1) that $\lim_{n \rightarrow \infty} E[J_n(\lambda\varphi)] \geq E[J(\lambda\varphi)]$. Because of $J_n^\pm \rightarrow J^\pm$ and Fatou’s lemma, it is sufficient to show that $(J_n(\lambda\varphi)^-)_{n \in \mathbb{N}}$ is uniformly integrable. Since J_n is increasing in n , J_n^- is decreasing in n and $J_1(\lambda\varphi)^- \leq |U(1)| + \lambda\varphi$ is an integrable upper bound for $J_n(\lambda\varphi)^-$, $n \in \mathbb{N}$. \square

Example 2.10 shows that the value function $u(x, U)$ is not continuous in general. If, however, the non-concave utility function is continuous, then also the value function $u(x, U)$ is continuous.

Proposition 4.2. *Let Assumption 3.1 be satisfied. Then the value function $u(x, U)$ is upper-semicontinuous. If U is, in addition, continuous, then the value function $u(x, U)$ is continuous on $(0, \infty)$.*

Proof. For upper-semicontinuity, consider a sequence $x_n \searrow x \in (0, \infty)$ and, using Theorem 3.4, denote the maximizer for $u(x_n, U)$ by f^n . Proposition 3.5 gives $\hat{f} \in C(x)$ satisfying

$$\limsup_{n \rightarrow \infty} u(x_n, U) = \limsup_{n \rightarrow \infty} E[U(f^n)] \leq E[U(\hat{f})] \leq u(x, U),$$

which completes the proof of upper-semicontinuity for $x \in (0, \infty)$. For $x = 0$, the upper-semicontinuity of $u(x, U)$ follows from the one of $u(x, U_c)$ since $U_c(0) = U(0) \leq u(x, U) \leq u(x, U_c)$ for all $x > 0$. The latter is proved (in greater generality) in Theorem 6 of Siorpaes [103].

It remains to show the lower-semicontinuity of $u(x, U)$ if U is continuous. Fix some $f_0 \in C(x)$ with $E[U(f_0)^-] < \infty$ and $E_Q[f_0] = x$ and consider some sequence $\epsilon_i \searrow 0$. We define sequences $f_i := f_0 1_{\{f_0 < \frac{x}{2}\}} + (1 - \epsilon_i) f_0 1_{\{f_0 \geq \frac{x}{2}\}}$ and $\tilde{\epsilon}_i := x - E_Q[f_i]$. Dominated convergence gives $\tilde{\epsilon}_i \searrow 0$, and continuity of U gives $U(f_i) \nearrow U(f_0)$ for $i \rightarrow \infty$. For ϵ_i small enough, we have $f_i > \delta$ on $\{f_0 \geq \frac{x}{2}\}$ for some $\delta > 0$. Thus, we have $U(f_i)^- < U(\delta)^-$ on $\{f_0 \geq \frac{x}{2}\}$, and dominated convergence yields $E[U(f_i)^- 1_{\{f_0 \geq \frac{x}{2}\}}] \rightarrow E[U(f_0)^- 1_{\{f_0 \geq \frac{x}{2}\}}]$. The equality $f_i = f_0$ on $\{f_0 \leq \frac{x}{2}\}$ gives $E[U(f_i)^- 1_{\{f_0 \leq \frac{x}{2}\}}] = E[U(f_0)^- 1_{\{f_0 \leq \frac{x}{2}\}}]$. Together with Fatou's lemma for $U(f_i)^+$, we obtain

$$\liminf_{i \rightarrow \infty} u(x - \tilde{\epsilon}_i, U) \geq \liminf_{i \rightarrow \infty} E[U(f_i)] \geq E[U(f_0)].$$

The result follows by taking the supremum over $f_0 \in C(x)$. \square

II.5 Results for an atomless model

In this section, we specialize our setup to the case when the underlying probability space is atomless. This is equivalent to assuming that the probability space supports a continuous distribution (see Definition A.26 and Proposition A.27 in Föllmer and Schied [43] for a precise definition and equivalent formulations). One special example of this occurs when we have a pricing density with a continuous distribution. This case often appears in the literature; see for instance Larsen [77]. It is shown there in the setting of the Black–Scholes model that $u(x, U_c) = u(x, U)$ and that the non-concave utility maximization problem $u(x, U)$ can be reduced to the concavified problem $u(x, U_c)$. However, the assumption that the distribution of φ is continuous is not very satisfactory from an economic point of view since it does not provide any structural explanation for the results. Also, from a mathematical point of view, the assumption is not very elegant since it might be tedious or even impossible to verify the continuity of the distribution. Carassus and Pham [29] and Bichuch and Sturm [17] use techniques from Malliavin calculus (and additional assumptions) to show the continuity of the distribution of φ in Example 2.7; we also present an example with jumps in the price process where the underlying probability space is atomless and the distribution of the pricing density is not continuous.

The main result of this section is that $u(x, U) = u(x, U_c)$ already holds if the underlying probability space is atomless. This unifies and generalizes the existing results in the literature, and the proof also provides an economically intuitive explanation. Note also that the assumption of having an atomless probability space is easily verifiable in applications.

Theorem 5.1. *Let (Ω, \mathcal{F}, P) be atomless. Then it holds that*

$$u(x, U) = u(x, U_c) \quad \text{for all } x > 0. \quad (5.1)$$

The non-concave problem $u(x, U)$ admits a maximizer if and only if the concavified problem $u(x, U_c)$ admits a maximizer. Every maximizer for the non-concave problem $u(x, U)$ also maximizes the concavified problem $u(x, U_c)$.

While every maximizer for $u(x, U)$ is also a maximizer for $u(x, U_c)$, the opposite conclusion is not true in general.

Example 5.2. *A maximizer for $u(x, U_c)$ is in general not a maximizer for $u(x, U)$. Fix $\varphi \equiv 1$ and $x \in \{U < U_c\}$. Doing nothing, i.e., $f \equiv x$, is optimal for $u(x, U_c)$, but is not optimal for $u(x, U)$. Lemma 5.7 below describes model classes where a solution for $u(x, U_c)$ is also a solution for $u(x, U)$.*

As we see below in Lemma 5.7, Theorem 5.1 is straightforward and (essentially) known in the case in which the distribution of φ is continuous. In general, it is therefore a natural idea to approximate possible mass points in that distribution by continuous distributions. Such an approach is possible, but it requires some new and involved convergence results as well as additional restrictive integrability conditions on φ (see Section III.3.3). We use another approach based on rearrangement techniques that works without additional assumptions on φ . This approach is based on a slightly stronger result.

Proposition 5.3. *Suppose that (Ω, \mathcal{F}, P) is atomless. For any $f \in C(x)$, there exists $f^* \in C(x)$ satisfying $\{f^* \in \{U < U_c\}\} = \emptyset$ and*

$$E[U(f^*)] = E[U_c(f^*)] = E[U_c(f)] \geq E[U(f)], \quad (5.2)$$

where the last inequality is strict if and only if $P[f \in \{U < U_c\}] > 0$.

This proposition is inspired by Rieger [94] who solved the case of a pricing density with a continuous distribution. In addition to giving a rigorous proof, our contribution compared to Rieger [94] is to prove the statement without any assumption on the distribution of φ . This is the crucial step in extending the existing results in the literature. Let us first explain the idea. Fix a non-concave utility U with $\{U < U_c\} = (a, b)$, choose $x \in (a, b)$ and consider the random variable $f \equiv x$. The idea is to use *rearrangement techniques* to construct $f^* \in C(x)$ with the same expected concavified utility and no probability weight in $\{U < U_c\}$. The method to achieve this is to concentrate all the weight on a and b (see Figure II.1), i.e., to choose f^* of the form $f^* = a1_A + b1_{A^c}$. The set $A \in \mathcal{F}$ has to be chosen in such a way that the agent can still afford the claim.

In the special case $\varphi \equiv 1$, a feasible choice for A is any set satisfying $P[A] = \frac{b-x}{b-a}$. For an arbitrary φ , the condition on A is a bit more involved;




Figure II.1: concentrate all the weight in (a, b) on the corners a and b

in particular, we have to put the “expensive” states ($\varphi(\omega)$ high) on a and the “cheap” states ($\varphi(\omega)$ low) on b . For a general non-concave utility U with $\{U < U_c\} = \bigcup_i (a_i, b_i)$ as in Lemma 2.11, every set $\{\omega \in \Omega \mid f \in (a_i, b_i)\}$ is rearranged separately in a similar way by concentrating the weight on a_i and b_i . For $\omega \in \Omega$ with $f(\omega) \in \{U = U_c\}$, the random variable $f^*(\omega) := f(\omega)$ is not changed.

Proof of Proposition 5.3. We split the proof into several steps; the plan is as follows. In the first part, we start with some preliminary remarks and define the sets needed for the construction of f^* . We then construct f^* and show that $f^* \in C(x)$ in the second part. (5.2) is shown in the third part. The final “if and only if” statement is a direct consequence of (5.2).

1) It is shown in Lemma 2.11 that the concave envelope U_c is locally affine on $\{U < U_c\}$ and that $\{U < U_c\} = \bigcup_i (a_i, b_i)$ for some a_i and b_i . We define

$$S_i := \{\omega \in \Omega \mid f(\omega) \in (a_i, b_i)\}$$

and $S := \bigcup_i S_i = \{f \in \{U < U_c\}\}$, which means that S contains all the states where f takes values in the non-concave part of U . For $\omega \in S_i$, the weight $\lambda(\omega)$ is defined by $\lambda(\omega) := (b_i - f(\omega))/(b_i - a_i)$, and since $f(\omega) \in (a_i, b_i)$ by the definition of S_i , it holds that $\lambda(\omega) \in (0, 1)$. An elementary calculation shows that for $\omega \in S_i$, we have (by construction)

$$f(\omega) = \lambda(\omega)a_i + (1 - \lambda(\omega))b_i. \quad (5.3)$$

The idea is now to decompose every set S_i into two parts. Because of the atomless structure of the probability space, there exist (Lemma A.28 in Föllmer and Schied [43]) a random variable \mathcal{U} with a uniform distribution on $(0, 1)$ and an increasing function q_φ such that $\varphi = q_\varphi(\mathcal{U})$ holds P -a.s. (q_φ is a *quantile* function of φ). Since $f_i(s) := P[S_i \cap \{\mathcal{U} < s\}]$ is a continuous function from $[0, 1]$ to $[0, P[S_i]]$, we find s_i such that $f_i(s_i) = E[(1 - \lambda)1_{S_i}]$. We now define the disjoint sets

$$S_{i1} := S_i \cap \{\mathcal{U} \geq s_i\} \text{ and } S_{i2} := S_i \cap \{\mathcal{U} < s_i\}.$$

Note that $S_{i1} \cup S_{i2} = S_i$, which gives $P[S_{i1}] = P[S_i] - P[S_{i2}]$. Moreover, due to the definition of s_i , we can express the probabilities of S_{i2} and S_{i1} in terms of S_i and λ . More precisely, we have

$$P[S_{i2}] = f_i(s_i) = E[(1 - \lambda)1_{S_i}] \text{ and } P[S_{i1}] = E[\lambda 1_{S_i}]. \quad (5.4)$$

Finally, note that $\mathcal{U} < s_i$ on S_{i2} and $\mathcal{U} \geq s_i$ on S_{i1} by definition. Since $\varphi = q_\varphi(\mathcal{U})$ is an increasing function of \mathcal{U} , we have

$$\sup_{\omega \in S_{i2}} \varphi(\omega) \leq \inf_{\omega \in S_{i1}} \varphi(\omega). \quad (5.5)$$

2) The modified random variable f^* is defined by

$$f^*(\omega) := \begin{cases} f(\omega), & \omega \in S^c, \\ a_i, & \omega \in S_{i1}, \\ b_i, & \omega \in S_{i2}. \end{cases}$$

The measurability of f implies that $S_i = \{f \in (a_i, b_i)\} \in \mathcal{F}$, and the measurability of \mathcal{U} implies that $\{\mathcal{U} \geq s_i\} \in \mathcal{F}$; so also the intersection $S_{i1} = S_i \cap \{\mathcal{U} \geq s_i\}$ is in \mathcal{F} . Since $\{f^* \leq x\}$ can be written in terms of f and S_{i1} , we get measurability of f^* . We show below in step 4) that

$$a_i E[\varphi 1_{S_{i1}}] + b_i E[\varphi 1_{S_{i2}}] \leq E[\varphi(\lambda a_i + (1 - \lambda)b_i) 1_{S_i}], \quad (5.6)$$

which means that $E[\varphi f^* 1_{S_i}] \leq E[\varphi f 1_{S_i}]$ holds for every i . Since $f^* = f$ on $S^c = (\bigcup_i S_i)^c$, the assumption $f \in C(x)$ then gives $E[\varphi f^*] \leq E[\varphi f] \leq x$, which means that $f^* \in C(x)$.

3) For the first equality in (5.2) that $E[U(f^*)] = E[U_c(f^*)]$, note that $f^* = f$ on S^c by the definition of f^* . On $S = \bigcup_i S_i$, the definition of f^* gives $f^* \in \{a_i, b_i\}$ on S_i . By Lemma 2.11, the constants a_i and b_i are also in $\{U = U_c\}$; hence we arrive at $f^*(\omega) \in \{U = U_c\}$ for all $\omega \in \Omega$. The latter implies $E[U(f^*)] = E[U_c(f^*)]$. For the second equality in (5.2) that $E[U_c(f^*)] = E[U_c(f)]$, we first show that $E[U_c(f) 1_{S_i}] = E[U_c(f^*) 1_{S_i}]$ holds for every i . This needs four ingredients. In the first step, we apply (5.3) to rewrite f in terms of a_i , b_i and λ . Second, we use that the concave envelope U_c is affine on $[a_i, b_i]$ and the fact that $\lambda(\omega) \in (0, 1)$. We then apply (5.4) and finally, we rewrite the resulting convex combination in terms of f^* . Following these steps, we obtain

$$\begin{aligned} E[U_c(f) 1_{S_i}] &= E[U_c(\lambda a_i + (1 - \lambda)b_i) 1_{S_i}] \\ &= U_c(a_i) E[\lambda 1_{S_i}] + U_c(b_i) E[(1 - \lambda) 1_{S_i}] \\ &= U_c(a_i) P[S_{i1}] + U_c(b_i) P[S_{i2}] \\ &= E[U_c(f^*) 1_{S_i}] \end{aligned}$$

for every i . Recall that $S = \bigcup_i S_i$ holds by definition and that $f = f^*$ on the complement S^c by the definition of f^* . Monotone convergence and the above equality show that $U_c(f)^\pm \in L^1$ if and only if $U_c(f^*)^\pm \in L^1$ and that

$$E[U_c(f)1_S] = \sum_i E[U_c(f)1_{S_i}] = \sum_i E[U_c(f^*)1_{S_i}] = E[U_c(f^*)1_S]$$

if $U_c(f)^\pm \in L^1$. In all cases, we obtain $E[U_c(f)] = E[U_c(f^*)]$. The inequality in (5.2) holds because $U_c \geq U$, and the “if and only if” part is clear from the definition of $\{U < U_c\}$.

4) Finally, to show (5.6), recall that $P[S_{i2}] = E[(1 - \lambda)1_{S_i}]$. Subtracting $E[(1 - \lambda)1_{S_{i2}}]$ on both sides gives

$$E[\lambda 1_{S_{i2}}] = E[(1 - \lambda)1_{S_i}] - E[(1 - \lambda)1_{S_{i2}}] = E[(1 - \lambda)1_{S_{i1}}].$$

This equality can be combined with (5.5) to deduce

$$E[\varphi \lambda 1_{S_{i2}}] \leq \sup_{\omega \in S_{i2}} \varphi(\omega) E[\lambda 1_{S_{i2}}] \leq \inf_{\omega \in S_{i1}} \varphi(\omega) E[(1 - \lambda)1_{S_{i1}}] \leq E[\varphi(1 - \lambda)1_{S_{i1}}].$$

Because $a_i \leq b_i$, the product of the positive terms $E[\varphi(1 - \lambda)1_{S_{i1}}] - E[\varphi \lambda 1_{S_{i2}}]$ and $b_i - a_i$ is again positive, and this can be rewritten as

$$a_i E[\varphi(1 - \lambda)1_{S_{i1}}] + b_i E[\varphi \lambda 1_{S_{i2}}] \leq a_i E[\varphi \lambda 1_{S_{i2}}] + b_i E[\varphi(1 - \lambda)1_{S_{i1}}]. \quad (5.7)$$

Finally, rewriting $a_i E[\varphi 1_{S_{i1}}] + b_i E[\varphi 1_{S_{i2}}]$ in terms of λ , applying (5.7) and rewriting the resulting terms again in a compact form yields

$$\begin{aligned} & a_i E[\varphi 1_{S_{i1}}] + b_i E[\varphi 1_{S_{i2}}] \\ &= a_i E[\varphi \lambda 1_{S_{i1}}] + a_i E[\varphi(1 - \lambda)1_{S_{i1}}] + b_i E[\varphi \lambda 1_{S_{i2}}] + b_i E[\varphi(1 - \lambda)1_{S_{i2}}] \\ &\leq a_i E[\varphi \lambda 1_{S_{i1}}] + a_i E[\varphi \lambda 1_{S_{i2}}] + b_i E[\varphi(1 - \lambda)1_{S_{i1}}] + b_i E[\varphi(1 - \lambda)1_{S_{i2}}] \\ &= E[\varphi(\lambda a_i + (1 - \lambda)b_i)1_{S_i}], \end{aligned}$$

which finishes the proof of (5.6). \square

Theorem 5.1 follows now directly from Proposition 5.3, as follows.

Proof of Theorem 5.1. The inequality “ \leq ” for (5.1) follows from $U \leq U_c$. For “ \geq ”, we start with some $f \in C(x)$. Proposition 5.3 gives $f^* \in C(x)$ with $E[U(f^*)] = E[U_c(f)]$. The inequality “ \geq ” follows since $f \in C(x)$ is arbitrary.

For the second part, we first assume that there is a maximizer $f \in C(x)$ for the non-concave problem $u(x, U)$. Then (5.1) gives

$$u(x, U) = E[U(f)] \leq E[U_c(f)] \leq u(x, U_c) = u(x, U),$$

which shows that f also maximizes $u(x, U_c)$. Conversely, assume that f maximizes $u(x, U_c)$. Proposition 5.3 gives a candidate $f^* \in C(x)$ satisfying $E[U(f^*)] = E[U_c(f^*)] = E[U_c(f)]$. We deduce that

$$u(x, U) = u(x, U_c) = E[U_c(f)] = E[U_c(f^*)] = E[U(f^*)],$$

which shows that f^* is a maximizer for $u(x, U)$. \square

Remark 5.4. 1) Looking more closely at the proof of Proposition 5.3 shows that we have not directly used that U is defined on \mathbb{R}_+ . The essential ingredients are that U_c exists and that $\{U < U_c\}$ can be written as a countable (finite) union of finite intervals. The proof is also valid for non-concave utilities defined on \mathbb{R} if these two assumptions are satisfied.

2) Theorem 5.1 shows that an atomless underlying probability space is sufficient for $u(x, U) = u(x, U_c)$. Under the (harmless) Assumption 3.1, an atomless probability space is in fact also necessary for $u(x, U) = u(x, U_c)$ in the following sense: if the underlying probability space contains (at least) one atom $\{\omega_1\}$, then we can choose a particular non-concave utility function, a pricing density φ and an initial capital x such that $u(x, U) < u(x, U_c)$. Indeed, fix $U(x) := 1_{\{x \geq 1\}}$, $x \in (0, P[\{\omega_1\}]\varphi(\omega_1))$ and a pricing density φ in such a way that $\varphi(\omega_1) < \inf_{\omega \in \Omega \setminus \{\omega_1\}} \varphi(\omega)$ holds. The problem $u(x, U)$ admits a maximizer \tilde{f} due to Theorem 3.4. If $u(x, U) = u(x, U_c)$ holds, then \tilde{f} is also a maximizer for $u(x, U_c)$ and Proposition 8.3 gives $\tilde{f} \in -\partial J(\lambda\varphi)$ for some $\lambda > 0$. It follows by part (iv) of Lemma 7.1 that \tilde{f} and φ are anti-comonotonic. But on the other hand, being a maximizer for $u(x, U)$ as well as for $u(x, U_c)$ implies that \tilde{f} satisfies $P[\tilde{f} \in (0, 1)] = 0$. This gives $\tilde{f}(\omega_1) = 0$ since $\tilde{f} \in C(x)$ and $P[\{\omega_1\}]\varphi(\omega_1) > x$. But this implies $(\tilde{f}(\omega) - \tilde{f}(\omega_1))(\varphi(\omega) - \varphi(\omega_1)) \geq 0$ for all $\omega \in \Omega$ where the inequality is strict for some ω . This means that \tilde{f} and φ are not anti-comonotonic giving the required contradiction.

However, for a particular pricing density, it might be possible to obtain $u(x, U) = u(x, U_c)$ for all $x > 0$ even for models where the underlying probability space is not atomless; see Example IV.5.3. \diamond

The first message of Theorem 5.1 is that, in the atomless case, any solution for the non-concave optimization problem also solves the concavified one, which in turn is well understood (see, for instance, Bouchard et al. [22] and Westray and Zheng [107]). In particular, the structure of the solution for the concavified problem is known. In this way, Theorem 5.1 can be used to describe explicitly the structure of all solutions to the non-concave optimization problem. Theorem 5.1 also says that the existence of a maximizer for the concavified optimization problem already guarantees the existence of a maximizer for the non-concave problem. Since the conditions for existence for the concavified problem are slightly weaker than those in Theorem 3.4, Theorem 5.1 can be used to slightly relax the assumptions in the present atomless case, as follows.

Corollary 5.5. *Suppose that (Ω, \mathcal{F}, P) is atomless and Assumption 2.3 is satisfied. Then there exists $\tilde{x} \in (0, \infty]$ such that for $x \in (0, \tilde{x})$, the problem $u(x, U)$ has a maximizer and every maximizer satisfies $f^* \in -\partial J(\lambda^*\varphi)$ for some $\lambda^* > 0$. If Assumption 3.1 is satisfied, then $\tilde{x} = \infty$.*

Proof. The argument is a combination of known results on non-smooth util-

ity maximization, non-smooth versions of results of Jin et al. [62], and Theorem 5.1. For completeness, we summarize the material on non-smooth utility maximization in Appendix II.8.

Proposition 8.1 shows that there is $\tilde{x} \in (0, \infty]$ such that for $x \in (0, \tilde{x})$, the concavified optimization problem $u(x, U_c)$ has a maximizer. By Theorem 5.1, the non-concave problem $u(x, U)$ then also admits a solution f^* which also maximizes $u(x, U_c)$. The particular form of f^* then follows from Proposition 8.1. If Assumption 3.1 is satisfied, the statements follow in much the same way by using Proposition 8.3 instead of Proposition 8.1. \square

For models with an atomless underlying probability space, Corollary 5.5 shows that any maximizer for $u(x, U)$ satisfies $f^* \in -\partial J(\lambda\varphi)$ for some $\lambda \geq 0$. But J is a convex function and the subdifferential is monotone in the sense that $J'_+(z_1) \leq J'_-(x) \leq J'_+(x) \leq J'_-(z_2)$ when $z_1 < x < z_2$ (see part (iv) of Lemma 7.1). This implies that f^* and $\lambda\varphi$ are anti-comonotonic. If U_c is strictly increasing, the first order condition implies $\lambda > 0$ and we obtain:

Corollary 5.6. *Let U_c be strictly increasing, let (Ω, \mathcal{F}, P) be atomless and let f^* be a maximizer for $u(x, U)$ derived in Corollary 5.5. Then f^* and φ are anti-comonotonic.*

In some examples, the distribution of the pricing density φ is continuous. This simplifies several arguments and allows for further results. Larsen [77], Carassus and Pham [29] and Rieger [94] consider a framework with $\{U < U_c\} = \bigcup_{i=1}^N (a_i, b_i)$ for a fixed N and argue that a maximizer for $u(x, U_c)$ also maximizes $u(x, U)$. The arguments of Larsen [77] can easily be adapted to our slightly more general case.

Lemma 5.7. *Let φ have a continuous distribution and let $f^* \in -\partial J(\lambda\varphi)$ be a maximizer f^* for $u(x, U_c)$. Then f^* is also a maximizer for $u(x, U)$.*

Proof. Let f^* be a maximizer for the concavified problem $u(x, U_c)$. If

$$P[f^* \in \{U < U_c\}] = 0, \quad (5.8)$$

it follows that $u(x, U) \leq u(x, U_c) = E[U_c(f^*)] = E[U(f^*)] \leq u(x, U)$ which means that f^* is also a maximizer for $u(x, U)$. To prove (5.8), note that $f^* \in -\partial J(\lambda\varphi)$ for some $\lambda \geq 0$. The case $\lambda = 0$ gives $f^* = -J'_+(0)$ and the result follows from Lemma 7.2. So consider $\lambda > 0$. It is shown in Lemma 2.11 that U_c is locally affine on $\{U < U_c\}$. So if $f^*(\omega) \in \{U < U_c\}$, then $U'_c(f^*(\omega)) = \lambda\varphi(\omega)$ by Lemma 2.12. But if U_c is affine on an interval (a, b) with slope c , then it follows that J is not differentiable in c since $x \in -\partial J(c)$ for $x \in (a, b)$. This gives

$$\{f^* \in \{U < U_c\}\} \subseteq \{\omega \in \Omega \mid J \text{ is not differentiable in } \lambda\varphi(\omega)\}.$$

The conjugate J is finite and convex on $(0, \infty)$ and thus differentiable there except for at most countably many points $y_i, i \in \mathbb{N}$. This yields (5.8) since

$$P[\{\omega \in \Omega \mid J(\cdot) \text{ is not differentiable in } \lambda\varphi(\omega)\}] = \sum_i P[\{\lambda\varphi = y_i\}] = 0,$$

where continuity of the distribution of φ is used in the last equality. \square

Remark 5.8. 1) If the distribution of φ is continuous, Theorem 5.1 follows directly from Lemma 5.7. Furthermore, note that in comparison to the quantile approach developed in Jin and Zhou [64], Carlier and Dana [31] and He and Zhou [56], we obtain more explicit results for the structure of the maximizer by using the more explicit form of the preference functional.

2) Lemma 5.7 is also useful for the more general case with more than one pricing measure. It is mentioned in Remark 2.4 that for the concave utility function U_c , the solution to the problem with more than one pricing measure also solves the problem for only one well-chosen (generalized) pricing measure. If the distribution of the (generalized) pricing density is continuous, it follows as in Lemma 5.7 that the solution f^* to the problem with this pricing measure satisfies $P[f^* \in \{U < U_c\}] = 0$. This implies that it is also the solution to the non-concave problem with multiple pricing measures, see also part 1) of Remark IV.4.8.

This idea is used in He and Zhou [56] for one particular model. Bichuch and Sturm [17] explore the idea more thoroughly and derive conditions under which it can be applied for specific classes of models.

3) The case with (infinitely) many pricing measures having not necessarily continuous distributions is more subtle. As in part 2), the solution f^* for U_c satisfies $f^* \in -\partial J(\lambda^*\varphi)$ and $E[\varphi f^*] = x$ for a particular dual object φ . However, Westray and Zheng [108] show that these conditions are in general not sufficient for optimality of f^* in the concavified problem. This is in contrast to the case with one pricing measure where these assumptions are sufficient for optimality and this is the reason why the proof of Proposition 5.3 cannot be extended directly to the case with many pricing measures. \diamond

In a similar way as in Lemma 5.7, we also get uniqueness for the maximizer for $u(x, U)$ provided that U_c is strictly increasing.

Lemma 5.9. *Let U_c be strictly increasing, let φ have a continuous distribution and let f^* be a maximizer for $u(x, U)$ derived in Corollary 5.5. Then the maximizer is P -a.s. unique.*

Proof. Let $f_i \in -\partial J(\lambda_i\varphi)$ for $i = 1, 2$ be maximizers for $u(x, U)$. If $\lambda_1 > \lambda_2$, it follows from part (ix) of Lemma 7.1 that $f_1 \leq f_2$. Since $f_i \in \{U = U_c\}$ P -a.s. and U is strictly increasing there, we deduce that $f_1 = f_2$ P -a.s. Suppose now $\lambda_1 = \lambda_2$. If J is differentiable in $\lambda_1\varphi(\omega)$, then $\partial J(\lambda_1\varphi(\omega))$ is a singleton and hence $f_1(\omega) = f_2(\omega)$. But this implies

$$\{f_1 \neq f_2\} \subseteq \{\omega \in \Omega \mid J(\cdot) \text{ is not differentiable in } \lambda^*\varphi(\omega)\},$$

and the latter has probability 0 since the distribution of φ is continuous. \square

We finally discuss one example which makes the advantages of Theorem 5.1 transparent.

Example 5.10. We consider some time horizon $T \in (0, \infty)$ and a probability space (Ω, \mathcal{F}, P) on which there are a Poisson process $(N_t)_{t \geq 0}$ with intensity $\lambda > 0$ and a Brownian motion $(W_t)_{t \geq 0}$. Let $(\mathcal{F}_t)_{0 \leq t \leq T}$ be the augmented filtration generated by N and W . We consider a (discounted) market consisting of a savings account $B \equiv 1$ and two risky stocks S^1, S^2 described by

$$\begin{aligned} dS_t^1 &= \alpha S_t^1 dt + \sigma_1 S_t^1 d\tilde{N}_t, & S_0^1 &= 1, \sigma_1 > -1, \sigma_1 \neq 0, \alpha/\sigma_1 < \lambda, \\ dS_t^2 &= \sigma_2 S_t^2 dW_t, & S_0^2 &= 1, \sigma_2 > 0, \end{aligned}$$

where $\tilde{N}_t := N_t - \lambda t$ is the compensated Poisson process. This model defines a complete financial market and the unique martingale measure is defined by $dQ/dP = e^{(\lambda - \tilde{\lambda})T} (\frac{\tilde{\lambda}}{\lambda})^{N_T}$ for $\tilde{\lambda} := \lambda - \alpha/\sigma$ (for details, see for instance Section 10.6.1 in Jeanblanc et al. [61]). So instead of maximizing expected utility over dynamic strategies in the market, we may as well solve (2.2) for $\varphi = dQ/dP$. The underlying probability space is atomless, so it follows from Theorem 5.1 that $u(x, U) = u(x, U_c)$. However, the distribution of φ is not continuous, so $u(x, U)$ cannot be solved directly via Lemma 5.7.

Note that the martingale measure here is the same as in Example 2.8 even though the financial markets are not identical. Looking at the static problem (2.2), the difference is reflected in the underlying probability space. Here it is atomless while it is not atomless in Example 2.8; the set $\{N_T = 0\}$ is an atom in that case. The economic interpretation is as follows. The market from Example 2.8 consisting only of B and S^1 defines a complete market in continuous time in which one *cannot* generate an arbitrary probability distribution. An agent having a non-concave utility U can therefore not always generate the same expected non-concave utility as an agent with the concave utility U_c . Introducing as above the third asset S^2 does not increase the optimal expected utility for the agent with the utility U_c since dQ/dP does not change; his optimal final position is therefore the same as before. However, the new asset increases the optimal expected utility for the agent with the non-concave utility U . Because the underlying space is now atomless and the financial market is complete, the agent is able to generate by self-financing trading any distribution; in this way the agent generates the same optimal expected non-concave utility as the agent with the concave utility U_c .

To put the present section into perspective, let us finally briefly discuss some results in the literature which belong to the class of models analyzed here. Berkelaar et al. [13], Larsen [77], Carassus and Pham [29], Rieger [94] and Basak and Makarov [9] consider non-concave utility maximization problems under the additional assumption that φ has a continuous distribution and use (sometimes implicitly) the argument given in Lemma 5.7. In

addition, Carassus and Pham [29] consider Example 2.7 for $\mu = 0$ and use dynamic programming tools to show that a particular non-concave optimization problem satisfies $u(x, U) = u(x, U_c)$. Our Theorem 5.1 and Corollary 5.5 provide a unified derivation and perspective for these results. Moreover, our result also allows us to deal with other examples, such as Example 5.10, that are not possible to solve with the existing results in the literature.

II.6 Conclusion

In this chapter, we study the problem of non-concave utility maximization from terminal wealth for a budget set given by a single pricing measure. In the literature, the classical approach is to reduce the non-concave problem to a concavified problem and to apply the classical techniques from the concave case. While this approach is powerful for specific examples, it cannot be applied in the general case.

The present chapter analyzes the non-concave utility maximization problem directly. We first show the existence and several properties of a maximizer for $u(x, U)$. As in the classical concave case, this requires some additional assumptions that can be formulated in terms of the conjugate J of the non-concave utility function U . In contrast to the literature, we also study the value function $u(x, U)$ in detail. In particular, we show that $u(x, U_c)$ is the concave envelope of $u(x, U)$, which gives the relation between the non-concave maximization problem $u(x, U)$ and the concavified problem $u(x, U_c)$.

In contrast to the concave case, the underlying probability space crucially affects the non-concave optimization problem. While the value function $u(x, U)$ is not necessarily concave in general, it is shown to be concave and to equal $u(x, U_c)$ if the underlying probability space is atomless. This also gives a structural, easily verifiable assumption under which one can use the concavified problem $u(x, U_c)$ to analyze $u(x, U)$.

The significant impact of the underlying probability space raises a natural *stability* question: Do small deviations from the atomless structure change the results in the atomless model drastically? As an example, take a sequence of binomial models approximating the Black–Scholes model and for every model, consider the non-concave utility maximization problem for a fixed non-concave utility U . It is then shown in Chapter III that the sequences of value functions and optimal final positions converge to the corresponding quantities in the limit model. These results complement the present chapter and provide additional intuition for the optimal behaviour of agents with a non-concave utility function.

Finally, we have to point out that our results depend crucially on the setup with one single pricing measure. It remains an interesting and challenging problem to study the case with infinitely many pricing measures. We take a step in this direction in Chapter IV.

II.7 Appendix A: Facts from convex analysis

Let $f : \mathbb{R} \rightarrow \bar{\mathbb{R}} = \mathbb{R} \cup \{\pm\infty\}$ be an extended real-valued function. The *effective domain* of f , denoted by $\text{dom}(f)$, consists of all $x \in \mathbb{R}$ such that $f(x) < \infty$, and its interior is denoted by $\text{int}(\text{dom}(f))$. The function f is called *proper* if both $\text{dom}(f) \neq \emptyset$ and $f(x) > -\infty$ for all x . The *conjugate* of f is the extended real-valued function f^* on \mathbb{R} defined by

$$f^*(y) := \sup_{x \in \mathbb{R}} (xy - f(x))$$

for all $y \in \mathbb{R}$. The *biconjugate* f^{**} is defined by $f^{**} := (f^*)^*$. If f is proper, lower-semicontinuous and convex, then its *subdifferential* ∂f is the multivalued mapping defined by $\partial f(x) = \emptyset$ if $f(x) = \infty$ and

$$\partial f(x) := \{y \in \mathbb{R} \mid f(\xi) \geq f(x) + y(\xi - x) \text{ for all } \xi \in \mathbb{R}\} \quad \text{if } x \in \text{dom}(f).$$

The *convex envelope* \bar{f} of f is the largest convex function $\bar{f} \leq f$.

Lemma 7.1. *Suppose that f is proper and lower-semicontinuous and its convex envelope \bar{f} is proper as well. Then:*

- (i) f^* is convex, proper and lower-semicontinuous.
 - (ii) f and its convex envelope have the same conjugate.
 - (iii) f^{**} is the lower-semicontinuous envelope of the convex envelope \bar{f} .
 - (iv) Fix $x_0 \in \text{dom}(f)$ and $y_0 \in \partial f^{**}(x_0)$. If f^* is differentiable in y_0 , then $f(x_0) = f^{**}(x_0)$.
- Let f , in addition, be convex and extend the right and left derivative functions f'_+ and f'_- beyond the interval $\text{dom}(f)$ by setting both $= \infty$ for points lying to the right of $\text{dom}(f)$ and both $= -\infty$ for points lying to the left. Then:
- (v) $y \in \partial f(x)$ if and only if $x \in \partial f^*(y)$.
 - (vi) $y \in \partial f(x)$ if and only if $f(x) + f^*(y) = xy$.
 - (vii) Let $x \in \text{dom}(f)$. f has a unique subgradient at x if and only if f is differentiable at x .
 - (viii) $\partial f(x) = \{y \in \mathbb{R} \mid f'_-(x) \leq y \leq f'_+(x)\}$.
 - (ix) $f'_+(z_1) \leq f'_-(x) \leq f'_+(x) \leq f'_-(z_2)$ when $z_1 < x < z_2$.

Proof. Most properties can be found in Rockafellar [95]; see Theorem 23.5 for (v) and (vi) Theorem 25.1 for (vii) and Theorem 24.1 and its discussion for (viii) and (ix). Statements (i), (ii) and (iii) are part of Theorem 11.1 of Rockafellar and Wets [96]. Part (iv) can be found in a similar form in Theorem 2 of Strömberg [104]. \square

We now apply Lemma 7.1 with $f(x) := -U(-x)$ to prove Lemmas 2.11, 2.12 and 2.13 and derive some additional properties. For convenience, we write $\text{dom}(f)$, $\text{dom}(f^{**})$ and ∂f^{**} in terms of the concave envelope U_c , i.e., $\text{dom}(U) := \{x \in \mathbb{R} \mid U(x) > -\infty\}$, $\text{dom}(U_c) := \{x \in \mathbb{R} \mid U_c(x) > -\infty\}$ and

$$\partial U_c(x) := \{p \in \mathbb{R} \mid U_c(z) \leq U_c(x) + p(z - x) \text{ for all } z \in \mathbb{R}\}$$

for $x \in \text{dom}(U_c)$. The function U is finite and upper-semicontinuous on $(0, \infty)$, satisfies $U(x) = -\infty$ for $x < 0$ and $U(0) = \lim_{x \searrow 0} U(x)$. Hence $-U(-x)$ is lower-semicontinuous on \mathbb{R} and proper. Moreover, $-U_c(-x)$ is proper due to the growth condition (2.1) and $-U_c(-x)$ is also lower-semicontinuous on all of \mathbb{R} . Indeed, in the case that $U(0) > -\infty$, this is a direct consequence of Corollary 17.2.1 of Rockafellar [95]; otherwise it results from Corollary 17.1.5 of Rockafellar [95] together with $\lim_{x \searrow 0} U(x) = -\infty$ and the growth condition (2.1). For completeness, we provide the detailed argument in Lemma 7.4. It therefore follows from part (iii) of Lemma 7.1 that $-U_c(-x) = f^{**}(x)$. Finally, note that $J(y) = f^*(y)$; so the convexity of J follows from part (i) of Lemma 7.1.

Proof of Lemma 2.11. U_c is finite and continuous: The condition (2.1) and the assumption $U(\infty) > 0$ imply that for any $\epsilon > 0$, there is $x_0 > 0$ such that $U(x) < \epsilon x$ for all $x \geq x_0$ and $U(x_0) > 0$. Therefore, $g(x) := U(x_0) + \epsilon x$ dominates U . Clearly, g is finite and concave. Hence $U_c \leq g$ and we conclude that U_c is on $(0, \infty)$ finite and concave, hence continuous.

U_c is locally affine on $\{U < U_c\}$: If $\{U < U_c\}$ is empty, there is nothing to prove. Otherwise choose $x_0 \in \{x \in \mathbb{R}_+ \mid U(x) < U_c(x)\}$ and observe that $\partial U_c(x_0)$ is non-empty since U_c is concave. Fix $y_0 \in \partial U_c(x_0)$. Note that $x_0 \in -\partial J(y_0)$ and $y_0 \in \text{dom}(J)$. Let us first consider the case that $y_0 \in (0, \infty)$. By way of contradiction, assume that J is differentiable in y_0 . By Lemma 7.1, (iv), it then follows that $U(x_0) = U_c(x_0)$. But x_0 is chosen in such a way that $U(x_0) < U_c(x_0)$. So J is not differentiable in y_0 and it follows from Lemma 7.1, (vii), that $\partial J(y_0)$ is multivalued. Thus, we have $y_0 \in \partial U_c(x)$ for all $x \in (-J'_+(y_0), -J'_-(y_0))$, which means that U_c is affine on that interval which also contains x_0 . The case $y_0 \in \text{dom}(J) \setminus (0, \infty)$ corresponds to $y_0 = 0 \in \partial U_c(x_0)$, which means that U_c is constant (hence affine) on $(-J'_+(0), \infty)$ and $-J'_+(0) \leq x_0$.

Structure of $\{U < U_c\}$: The set $\{U < U_c\} = \{x \in \mathbb{R}_+ \mid U(x) - U_c(x) < 0\}$ is open since $U - U_c$ is upper-semicontinuous. It can therefore be written as a union of disjoint open intervals. By way of contradiction, assume that $\{U < U_c\}$ contains (x_0, ∞) for some $x_0 \geq 0$. Then $U_c(x_0) = U(x_0)$ and U_c is affine on (x_0, ∞) . This contradicts the growth condition (2.1). \square

Proof of Lemma 2.12. Monotonicity of J as well as $J > -\infty$ on $(0, \infty)$ and $J = \infty$ on $(-\infty, 0)$ follow from the definition. To see that $J < \infty$ on $(0, \infty)$, we fix some $\epsilon > 0$. The growth condition (2.1) allows us to choose x_0 in such

a way that $U(x) < x\epsilon$ on (x_0, ∞) . This gives $\sup_{x \geq x_0} (U(x) - x\epsilon) \leq 0$ and so $J(\epsilon) \leq \max(0, U(x_0))$. Thus $J < \infty$ on $(0, \infty)$ follows from the monotonicity of J and since ϵ was arbitrary. The equivalence follows from Lemma 7.1, (v) and (vi). Lemma 7.1 (ii) gives that U and U_c have the same conjugate. \square

Lemma 7.2. $U(-J'_\pm(y)) = U_c(-J'_\pm(y))$ for $y > 0$. If $-J'_+(0) < \infty$, then also $U(-J'_+(0)) = U_c(-J'_+(0))$.

Proof. Recall that J is decreasing on $(0, \infty) = \text{int}(\text{dom}(J))$. So $-J'_\pm(y)$ is positive and finite. By way of contradiction, assume that there is $y > 0$ with $U(-J'_\pm(y)) < U_c(-J'_\pm(y))$. By Lemma 2.11, $\{U < U_c\} = \bigcup_i (a_i, b_i)$ is open. Hence there are i, z_1, z_2 such that $a_i < z_1 < -J'_\pm(y) < z_2 < b_i$. By Lemma 2.11, U_c is affine on (a_i, b_i) and by Lemma 2.12, we know that $U'_c(-J'_\pm(y)) = y$. This means that $z_i \in -\partial J(y)$ which contradicts Lemma 7.1, (viii). If $-J'_+(0) < \infty$, the same idea gives $U(-J'_+(0)) = U_c(-J'_+(0))$. \square

Lemma 7.3. The conjugate J satisfies $\lim_{y \rightarrow \infty} \sup_{x \in -\partial J(y)} x = 0$.

Proof. Since $x \leq -J'_-(y)$ for all $x \in -\partial J(y)$, it is sufficient to show that

$$\lim_{y \rightarrow \infty} -J'_-(y) = 0. \quad (7.1)$$

If $\tilde{y} := \lim_{x \rightarrow 0} \sup_{y \in \partial U_c(x)} y < \infty$, we have that $J(y) = U(0)$ for $y > \tilde{y}$ and this implies $\partial J(y) = 0$ for $y > \tilde{y}$. If $\lim_{x \rightarrow 0} \sup_{y \in \partial U_c(x)} y = \infty$, fix $\epsilon > 0$ and $y_0 := \sup_{y \in \partial U_c(\epsilon)} y$. For $q \in \partial U_c(\epsilon)$ and $y > y_0$, we have $q \leq y_0 < y$. Lemma 2.12 gives $-\epsilon \in \partial J(q)$ and Lemma 7.1, (ix) implies $-J'_-(y) \leq -J'_-(y_0) \leq \epsilon$, which finishes the proof of (7.1). \square

It remains to prove Lemma 2.13. The equivalence is proved in Lemma 4.1 of Deelstra et al. [37]. We only prove the second part.

Proof of Lemma 2.13. We argue that there are $x_0 > 0$ and $\gamma < 1$ such that

$$0 \leq U_c(x) \leq U(x) + \gamma U_c(x) \quad \text{on } (x_0, \infty). \quad (7.2)$$

The result then follows for $k = 1/(1 - \gamma)$.

Proposition 4.1. of Deelstra et al. [37] shows that $\text{AE}_0(J) < \infty$ and the growth condition (2.1) imply the existence of two constants x_0 and $\gamma < 1$ such that $\sup_{y \in \partial U_c(x)} yx \leq \gamma U_c(x)$ holds on (x_0, ∞) . By moving x_0 to the right if necessary, we may assume that $U(x_0)$ is positive. Moreover, recall that U_c is locally affine on $\{U < U_c\} = \bigcup_i (a_i, b_i)$. So, for $x \in (a_i, b_i) \cap (x_0, \infty)$, we rewrite $U_c(x)$, use that $U_c(a_i) = U(a_i) \leq U(x)$ and apply for $y = U'_c(x)$ the above inequality $U'_c(x)x \leq \gamma U_c(x)$ to get

$$U_c(x) = U_c(a_i) + U'_c(x)(x - a_i) \leq U(x) + U'_c(x)x \leq U(x) + \gamma U_c(x).$$

For $x \in \{U = U_c\} \cap (x_0, \infty)$, (7.2) follows since $U(x) = U_c(x)$ is positive. \square

Lemma 7.4. *If $U(0) = -\infty$, then $\lim_{x \rightarrow 0} U_c(x) = -\infty$.*

Proof. By way of contradiction, assume that $a := \lim_{x \rightarrow 0} U_c(x) > -\infty$. Fix a sequence $x_n \rightarrow 0$. By Corollary 17.1.5 in Rockafellar [95], it follows that there are sequences x_1^n, x_2^n, λ^n such that $x_1^n \leq x_n \leq x_2^n$ and

$$x_n = \lambda^n x_1^n + (1 - \lambda^n)x_2^n, \quad (7.3)$$

$$a - 1 < \lambda^n U(x_1^n) + (1 - \lambda^n)U(x_2^n). \quad (7.4)$$

(7.3) gives $x_2^n = (x_n - \lambda^n x_1^n)/(1 - \lambda^n)$. If $\lim_n \lambda^n = 0$, then it follows that $\lim_n x_2^n = \lim_n x_n = 0$ and therefore $\lambda^n U(x_1^n) + (1 - \lambda^n)U(x_2^n) \rightarrow -\infty$, which contradicts (7.4). Otherwise, consider a subsequence (again relabelled by n) with $\lim_n \lambda^n > 0$. If $\limsup_n x_2^n \leq b < \infty$, then it follows that

$$\lambda^n U(x_1^n) + (1 - \lambda^n)U(x_2^n) \leq \lambda^n U(x_1^n) + (1 - \lambda^n)U(b+1) \leq \lambda^n U(x_1^n) + |U(b+1)|$$

for n large enough, and this converges to $-\infty$ which again gives a contradiction. It remains to consider the case that $\limsup_n x_2^n = \infty$. Consider a subsequence (again relabelled by n) with $\lim_n x_2^n = \infty$. By the growth condition (2.1), for every $\epsilon > 0$ we can find x_0 such that $U(x) \leq \epsilon x$ on (x_0, ∞) . Fix n_0 such that $x_2^n > x_0$ for $n \geq n_0$. For $n \geq n_0$, we then get

$$\begin{aligned} \lambda^n U(x_1^n) + (1 - \lambda^n)U(x_2^n) &\leq \lambda^n U(x_1^n) + (1 - \lambda^n)\epsilon x_2^n \\ &= \lambda^n U(x_1^n) + (1 - \lambda^n)\epsilon \frac{x_n - \lambda^n x_1^n}{1 - \lambda^n} \\ &\leq \lambda^n U(x_1^n) + \epsilon x_n. \end{aligned}$$

The last part converges again to $-\infty$ since $\liminf \lambda^n > 0$ and $x_1^n \rightarrow 0$. This gives the required contradiction. \square

II.8 Appendix B: Non-smooth utility maximization

This appendix contains the results on non-smooth (concave) utility maximization which are relevant for the proofs in Section II.5. It combines non-smooth versions of results of Jin et al. [62] and known results on non-smooth utility maximization as presented in Bouchard et al. [22] and Westray and Zheng [107]. We use the same notation as in Sections II.3–II.4.

Theorem 3.4 shows that Assumption 3.1 is sufficient to guarantee the existence of a maximizer for $u(x, U_c)$ for all $x > 0$. Under the slightly weaker Assumption 2.3, the existence of a maximizer is still obtained for *some* $x > 0$.

Proposition 8.1. *Assume that $u(x, U_c) < \infty$ for some $x > 0$. Then there exists $\tilde{x} > 0$ such that for $x \in (0, \tilde{x})$, the problem $u(x, U_c)$ admits a maximizer. Every maximizer f^* is of the form $f^* \in -\partial J(\lambda^* \varphi)$ for some $\lambda^* > 0$.*

This result as well as its proof are non-smooth versions of Corollary 3.1 and its proof in Jin et al. [62]. The proof is based on the following lemma which is the non-smooth version of Theorem 3.1 of Jin et al. [62].

Lemma 8.2. *If $E[-J'_+(\lambda\varphi)\varphi] = \infty$ for some $\lambda > 0$, there is some $x > 0$ such that $u(x, U_c) = \infty$.*

Proof. 1) We start with the case $U_c(0) := \lim_{x \rightarrow 0} U_c(x) \geq 0$. Fix some $\lambda_0 > 0$ and $x > 0$. Since $E[-J'_+(\lambda_0\varphi)\varphi] = \infty$, one can find a set $A \in \mathcal{F}$ such that $E[-J'_+(\lambda_0\varphi)\varphi 1_A] \in (x, \infty)$. Define $h(\lambda) := E[-J'_+(\lambda\varphi)\varphi 1_A]$ for $\lambda \in [\lambda_0, \infty)$. Since $-J'_-$ is decreasing, the function $h(\lambda)$ is decreasing. Moreover, the monotone convergence theorem and Lemma 7.3 yield $\lim_{\lambda \rightarrow \infty} h(\lambda) = 0$. Therefore, there is some $\lambda_1 \geq \lambda_0$ such that

$$E[-J'_+(\lambda_1\varphi)\varphi 1_A] = h(\lambda_1+) \leq x \leq h(\lambda_1-) = E[-J'_-(\lambda_1\varphi)\varphi 1_A].$$

If the first inequality is an equality, define $f^* := -J'_+(\lambda_1\varphi)1_A$. If the second inequality is an equality, define $f^* := -J'_-(\lambda_1\varphi)1_A$. If both inequalities are equalities, the candidates are P -a.s. identical. If both inequalities are strict, define

$$\mu := \frac{x - h(\lambda_1+)}{h(\lambda_1-) - h(\lambda_1+)}$$

and $f^* := (-\mu J'_-(\lambda_1\varphi) - (1-\mu)J'_+(\lambda_1\varphi))1_A$. In all cases, we have $E[\varphi f^*] = x$ and $f^* \in -\partial J(\lambda_1\varphi)$ on A (in the third case, this uses that the subdifferential is a convex set). Lemma 2.12 gives $U_c(f^*(\omega)) = J(\lambda_1\varphi(\omega)) + f^*(\omega)\lambda_1\varphi(\omega)$ for $\omega \in A$ and $U_c(f^*(\omega)) = U_c(0)$ for $\omega \in A^c$. The assumption $U_c \geq 0$ implies $J \geq 0$ and we conclude that

$$E[U_c(f^*)] \geq E[(J(\lambda_1\varphi) + f^*\lambda_1\varphi)1_A] \geq \lambda_1 E[\varphi f^* 1_A] = \lambda_1 x \geq \lambda_0 x.$$

Since $\lambda_0 > 0$ is arbitrary, we arrive at $u(x, U_c) \geq \lim_{\lambda \rightarrow \infty} \lambda x = \infty$.

2) For the general case, define a shifted utility $U_c^{x_0}(x) := U_c(x_0 + x)$ on \mathbb{R}_+ , where $x_0 := \inf\{x > 0 \mid U(x) \geq 0\}$, and denote the conjugate of $U_c^{x_0}$ by J^{x_0} . If we show (as we do below) that

$$E[-J'_+(\lambda\varphi)\varphi] = \infty \quad \forall \lambda > 0 \implies E[-(J^{x_0})'_+(\lambda\varphi)\varphi] = \infty \quad \forall \lambda > 0, \quad (8.1)$$

we can use the fact that $U_c^{x_0} \geq 0$ and part 1) to see that

$$\sup\{E[U_c^{x_0}(f)] \mid f \in L_+^0, E[\varphi f] \leq \tilde{x} - x_0\} = \infty$$

for \tilde{x} large enough. Rewriting $U_c^{x_0}$ in terms of U_c then gives

$$u(\tilde{x}, U_c) \geq \sup\{E[U_c(f)] \mid f \in L_+^0, f \geq x_0, E[\varphi f] \leq \tilde{x}\} = \infty,$$

which is the statement of the lemma.

It remains to prove (8.1). In order to relate ∂J^{x_0} and ∂J , note that

$$J^{x_0}(y) - x_0 y = \sup_{x>0} \{U_c(x_0 + x) - (x + x_0)y\} \leq J(y) \quad (8.2)$$

holds for arbitrary $y > 0$. For $y \leq y_0 := \inf_{q \in \partial U_c(x_0)} q$, fix some $\tilde{x} \in -\partial J(y)$ and note that $\tilde{x} \geq x_0$ by Lemma 7.1, (ix). Together with the conjugacy relation between \tilde{x} and y and the definition of J^{x_0} , we obtain

$$J(y) = U_c(\tilde{x}) - \tilde{x}y = U_c(x_0 + (\tilde{x} - x_0)) - (\tilde{x} - x_0)y - x_0 y \leq J^{x_0}(y) - x_0 y,$$

which implies that there is equality in (8.2) for $y \leq y_0$. We use this together with the definition of $(J^{x_0})'_-$ and (8.2) (in this order) to see that for $y < y_0$,

$$\begin{aligned} J(y) + ((J^{x_0})'_-(y) - x_0)(z - y) \\ = J^{x_0}(y) - x_0 y + (J^{x_0})'_-(y)(z - y) - x_0(z - y) \leq J^{x_0}(z) - x_0 z \leq J(z) \end{aligned}$$

holds for $z > 0$ which means that $(J^{x_0})'_-(y) - x_0 \in \partial J(y)$ for $y < y_0$. Since $E[-J'_+(\lambda\varphi)\varphi 1_{\{\lambda\varphi > y_0\}}] \leq -J'_+(y_0) < \infty$, the assumption $E[-J'_+(\lambda\varphi)\varphi] = \infty$ for all $\lambda > 0$ also implies

$$\begin{aligned} E[-J'_+(\lambda\varphi)\varphi 1_{\{\lambda\varphi \leq y_0\}}] &= E[-J'_+(\lambda\varphi)\varphi(1 - 1_{\{\lambda\varphi > y_0\}})] \\ &\geq E[-J'_+(\lambda\varphi)\varphi] - J'_+(y_0) = \infty \end{aligned} \quad (8.3)$$

for all $\lambda > 0$. But then, plugging $(J^{x_0})'_-(y) - x_0 \in \partial J(y)$ into (8.3) gives

$$\begin{aligned} \infty &= E[-J'_+(\lambda\varphi)\varphi 1_{\{\lambda\varphi \leq y_0\}}] \leq E[-((J^{x_0})'_-(\lambda\varphi) - x_0)\varphi 1_{\{\lambda\varphi \leq y_0\}}] \\ &\leq E[-(J^{x_0})'_-(\lambda\varphi)\varphi] + x_0 \end{aligned}$$

for every $\lambda > 0$, which completes the proof of (8.1). \square

Proof of Proposition 8.1. Assumption 2.3 and the concavity of $u(\cdot, U_c)$ imply that $u(x, U_c) < \infty$ for all $x > 0$. Thus we apply Lemma 8.2 to get $\lambda_0 > 0$ satisfying $\tilde{x} := E[-J'_+(\lambda_0\varphi)\varphi] < \infty$. Similarly to the proof of Lemma 8.2, define $h(\lambda) := E[-J'_+(\lambda\varphi)\varphi]$ for $\lambda \in [\lambda_0, \infty)$. Since $-J'_+$ is decreasing, the function $h(\lambda)$ is decreasing. Moreover, the monotone convergence theorem and Lemma 7.3 yield $\lim_{\lambda \rightarrow \infty} h(\lambda) = 0$. Hence for all $x \in (0, \tilde{x}]$, there is some $\lambda_1 \geq \lambda_0$ such that

$$E[-J'_+(\lambda_1\varphi)\varphi] = h(\lambda_1+) \leq x \leq h(\lambda_1-) = E[-J'_-(\lambda_1\varphi)\varphi].$$

If one inequality is an equality, set $f^* := -J'_+(\lambda_1\varphi)$ or $f^* := -J'_-(\lambda_1\varphi)$, respectively. Otherwise, define $\mu := (x - h(\lambda_1+))/(h(\lambda_1-) - h(\lambda_1+))$ and set $f^* := (-\mu J'_-(\lambda_1\varphi) - (1 - \mu)J'_+(\lambda_1\varphi))$. In all three cases, f^* satisfies $E[\varphi f^*] = x$ and $f^* \in -\partial J(\lambda_1\varphi)$. The conjugacy relation (Lemma 2.12) gives

$$U_c(f^*(\omega)) = J(\lambda_1\varphi(\omega)) + f^*(\omega)\lambda_1\varphi(\omega). \quad (8.4)$$

In order to show optimality of f^* , consider $f \in C(x)$ and note that $E[\varphi f] \leq x$ by the definition of $C(x)$. It follows from the definition of J , relation (8.4) and $E[\varphi f^*] = x$ that

$$\begin{aligned} E[U_c(f)] &= E[U_c(f) - \lambda_1 f \varphi] + E[\lambda_1 f \varphi] \\ &\leq E[J(\lambda_1 \varphi)] + \lambda_1 x \\ &= E[U_c(f^*)] - E[f^* \varphi \lambda_1] + \lambda_1 x = E[U_c(f^*)]. \end{aligned} \tag{8.5}$$

Hence $u(x, U_c) = E[U_c(f^*)]$ since $f \in C(x)$ was arbitrary. Note that the inequality in (8.5) is strict if $f \notin -\partial J(\lambda_1 \varphi)$ on a set with strictly positive measure. Thus any other maximizer \hat{f} satisfies $\hat{f} \in -\partial J(\lambda_1 \varphi)$ as well. \square

Under Assumption 3.1, Theorem 3.2 in Bouchard et al. [22] or Theorem 5.1 in Westray and Zheng [107] give that Proposition 8.1 holds for any $x > 0$:

Proposition 8.3. *Suppose Assumption 3.1 is satisfied. Then the concave problem $u(x, U_c)$ has a solution $f^* \in C(x)$ for every $x > 0$. Every solution satisfies $f^* \in -\partial J(\lambda^* \varphi)$ for some $\lambda^* \geq 0$.*

Chapter III

Stability and asymptotics along a sequence of models

This chapter corresponds to the article [92]. We consider a sequence of models, optimize behavioural preference functionals in each model and show that key quantities of the demand problem (I.1.2) converge to the corresponding quantities in a limit model.

III.1 Introduction

The analysis in Chapter II as well as all the work in the literature study the demand problem for a fixed underlying model. Since one is never exactly sure of the accuracy of a proposed model, it is important to know whether the behavioural predictions generated by a model change drastically if one slightly perturbs the model.¹ To the best of our knowledge, results on the stability of behavioural portfolio selection problems have not been available in the literature so far, and the main purpose of this article is to study this issue in detail. Formally, we consider a sequence of models, each represented by some probability space $(\Omega^n, \mathcal{F}^n, P^n)$ and some pricing measure Q^n , and we assume that this sequence converges weakly in a suitable sense (to be made precise later) to a limit model $(\Omega^0, \mathcal{F}^0, P^0, Q^0)$. For each model, we are interested in the demand problem

$$v^n(x) := \sup\{V_n(f) \mid f \in L_+^0(\Omega^n, \mathcal{F}^n, P^n), E_{Q^n}[f] \leq x\}, \quad (1.1)$$

where the functional $V_n : (\Omega^n, \mathcal{F}^n, P^n) \rightarrow \mathbb{R} \cup \{-\infty\}$ is defined by

$$V_n(f) := \int U(f) d(T \circ P^n) \quad (1.2)$$

¹One motivating example is the pricing kernel puzzle. While non-concave utility functions might explain the pricing kernel puzzle in discrete time, they cannot do so in several continuous-time models (see Section V.6 for a detailed discussion). This indicates that the particular choice of the underlying model might fundamentally influence the behavioural predictions of the model.

for a *non-concave* and *non-smooth* utility function U on \mathbb{R}_+ and a strictly increasing function $T : [0, 1] \rightarrow [0, 1]$ representing the *probability distortion* of the beliefs. We are then interested in the asymptotics of the value (indirect utility) $v^n(x)$ and its maximizer $f^n = \arg \max V_n(f)$, and we want to compare them with the analogous quantities in the limit model.

The fundamental ingredients of each model are described by the quantities $(\Omega^n, \mathcal{F}^n, P^n, Q^n)$. The assumption that the sequence of models converges (in a suitable sense) to a limit model means that the economic situation described by the n -th model is for sufficiently large n close (in a suitable sense) to the one described by the limit model. Our main contribution is to give easily verifiable assumptions such that similar economic situations also imply similar behavioural predictions for the agent, in the sense that the values $v^n(x)$ as well as (along a subsequence) the optimal final positions f^n converge to the corresponding quantities in the limit model.

In *concave* utility maximization, the (essentially) sufficient condition for these stability results is the weak convergence of the *pricing density* (or *pricing kernel*) dQ^n/dP^n to dQ^0/dP^0 (see for instance He [55] and Prigent [88]). However, in our non-concave setting, we present an example of a sequence of financial markets for which dQ^n/dP^n converges weakly to dQ^0/dP^0 , but where the limit $\lim_{n \rightarrow \infty} v^n(x)$ and $v^0(x)$ as well as the corresponding final positions differ substantially. We discuss these new effects in detail and give sufficient conditions to prevent such unpleasant phenomena.

In order to illustrate the main results, we provide several applications. First, we consider a sequence of binomial models approximating the Black–Scholes model; this is the typical example for the transition from discrete- to continuous-time models. Apart from its purely theoretical interest, this example is also of practical relevance since the discrete-time analysis provides numerical procedures for the explicit computation of the optimal consumption. This allows one to numerically determine the value function for (computationally difficult) continuous-time models via the value functions for (computationally tractable) discrete-time models. The second application is motivated by the practical difficulties to calibrate an underlying model. As one example, we therefore study whether a (small) misspecification of the drift in the Black–Scholes model significantly influences the optimal behaviour of the agent. In both examples above, we use a fixed time horizon T for the portfolio optimization problem. In practical applications, however, the time horizon might be uncertain or changing. In the third application, we therefore analyze whether or not a (marginal) misspecification of the investment horizon significantly influences the optimal behaviour of the agent.

These examples show the necessity of our analysis: Our models are at best approximations to the reality, so if we perturb one model slightly in a reasonable way and the behavioural predictions generated by the model change drastically, we may suspect that the model cannot tell us much about the real world behaviour. Our convergence results demonstrate that, for a

fairly broad class of preference functionals and models, the optimal behaviour is stable with respect to such small perturbations.

This chapter is structured as follows. In Section III.2, we abstractly describe the sequence of models, preference functionals and optimization problems. We also formulate and discuss the main result. In Section III.3, we present three applications of the main result. This also allows us to discuss the connections to the existing literature in more detail. In addition, we provide a concrete numerical example to illustrate the results. We prove the main result in Sections III.4 and III.5.

III.2 Problem formulation and main results

The following notation is used. If $x, y \in \mathbb{R}$, we denote $x^\pm = \max\{\pm x, 0\}$, $x \wedge y = \min\{x, y\}$ and $x \vee y = \max\{x, y\}$. For a function G and a random variable X , we write $G(X)^\pm$ for the positive/negative parts $(G(X))^\pm$. For a sequence (f^n) of random variables, we denote weak convergence of (f^n) to f^0 by $f^n \Rightarrow f^0$. A *quantile function* q_F of a distribution function F is a generalized inverse of F , i.e., a function $q_F : (0, 1) \rightarrow \mathbb{R}$ satisfying

$$F(q_F(s) -) \leq s \leq F(q_F(s)) \quad \text{for all } s \in (0, 1).$$

Quantile functions are not unique, but any two for a given F coincide a.e. on $(0, 1)$. Thus, we sometimes blur the distinction between “the” and “a” quantile function. A quantile function q_f of a random variable f is understood to be a quantile function q_F of the distribution F of the random variable f . If the sequence (f^n) converges weakly to f^0 , then any corresponding sequence (q_{f^n}) of quantile functions converges a.e. on $(0, 1)$ to q_{f^0} (see for instance Theorem 25.6 of Billingsley [19]). More properties of quantile functions can be found in Appendix A.3 of Föllmer and Schied [43].

III.2.1 Sequence of models and optimization problems

We consider a sequence of probability spaces $(\Omega^n, \mathcal{F}^n, P^n)_{n \in \mathbb{N}_0}$, where the probability space $(\Omega^0, \mathcal{F}^0, P^0)$ is atomless (see Definition A.26 and Proposition A.27 in Föllmer and Schied [43] for a precise definition and equivalent formulations). On each probability space, there is a probability measure Q^n equivalent to P^n with density $\varphi^n = dQ^n/dP^n \in L^1_+(\Omega^n, \mathcal{F}^n, P^n)$. We refer to Q^n as *pricing measure* and to φ^n as *pricing density* (or *pricing kernel*). We assume that the sequence (φ^n) converges weakly to the pricing density in the atomless model, i.e., $\varphi^n \Rightarrow \varphi^0$. To ensure that the atomic structure tend to the atomless structure, we assume that the atoms disappear in the following sense. Let \mathcal{G}^n be the set of atoms in \mathcal{F}^n (with respect to Q^n).

Assumption 2.1. $\lim_{n \rightarrow \infty} \sup_{A \in \mathcal{G}^n} Q^n[A] = 0$.

We impose the following integrability condition on $(\varphi^n)_{n \in \mathbb{N}}$.

Assumption 2.2. *The family $((\varphi^n)^\xi)_{n \in \mathbb{N}}$ is uniformly integrable for all $\xi < 0$.*

Having specified the sequence of models, we turn to the preference functionals. One cornerstone is the concept of a non-concave utility function.

Definition 2.3. A *non-concave utility* is a function $U : (0, \infty) \rightarrow \mathbb{R}$, which is strictly increasing, continuous and satisfies the growth condition

$$\lim_{x \rightarrow \infty} \frac{U(x)}{x} = 0. \quad (2.1)$$

We consider only non-concave utility functions defined on \mathbb{R}_+ . To avoid any ambiguity, we set $U(x) = -\infty$ for $x < 0$ and define $U(0) := \lim_{x \searrow 0} U(x)$ and $U(\infty) := \lim_{x \rightarrow \infty} U(x)$. Without loss of generality, we may assume that $U(\infty) > 0$. Observe that we do not assume that U is concave. In the concave case, the growth condition (2.1) not only implies, but is even equivalent to, the Inada condition at ∞ that $U'(\infty) = 0$.

Definition 2.4. The *concave envelope* of U is the smallest concave function $U_c : \mathbb{R} \rightarrow \mathbb{R} \cup \{-\infty\}$ such that $U_c(x) \geq U(x)$ holds for all $x \in \mathbb{R}$.

Some basic properties of the concave envelope U_c as well as of the set $\{U < U_c\} := \{x \in \mathbb{R}_+ \mid U(x) < U_c(x)\}$ can be found in Lemma II.2.11. A key tool to study the relation between U and U_c is the conjugate of U defined by

$$J(y) := \sup_{x > 0} \{U(x) - xy\}.$$

Because of the non-concavity of U , the conjugate J is no longer smooth; we therefore work with the subdifferential which is denoted by ∂J for the convex function J and by ∂U_c for the concave function U_c . The right- and left-hand derivatives of J are denoted by J'_+ and J'_- . Our proofs (mainly in the appendix) use the classical duality relations between U_c , J , ∂J and ∂U_c . Precise statements and proofs can be found in Lemma II.2.12.

In classical concave utility maximization, the asymptotic elasticity (AE) of the utility function is of importance. In particular, many results impose an upper bound on $\text{AE}(U)$. For a non-concave utility function, we impose a similar condition via the asymptotic elasticity of the conjugate J ,

$$\text{AE}_0(J) := \limsup_{y \rightarrow 0} \sup_{q \in \partial J(y)} \frac{|q|y}{J(y)}.$$

In order to define our preference functionals, we introduce an additional function T which represents the distortion of the distribution of the beliefs.

Definition 2.5. A *distortion* is a function $T : [0, 1] \rightarrow [0, 1]$ that is strictly increasing and satisfies $T(0) = 0$ and $T(1) = 1$.

In the literature, one can find several explicit functional forms for T . The most prominent example is

$$T(p) = \frac{p^\alpha}{(p^\alpha + (1-p)^\alpha)^{\frac{1}{\alpha}}} \quad (2.2)$$

suggested by Kahneman and Tversky [69]; they use the parameter $\alpha = 0.61$. For each model, we now define a preference functional V_n on $(\Omega^n, \mathcal{F}^n, P^n)$.

Definition 2.6. We consider one of the following cases:

Case 1: The preference functional V_n is defined by

$$V_n(f) := \int_0^\infty T(P^n[U(f) > x]) dx \quad (2.3)$$

for a distortion T and a non-concave utility U satisfying $U(0) \geq 0$. We refer to this case as *rank-dependent expected utility* (RDEU).

Case 2: The preference functional V_n is defined by

$$V_n(f) := E_n[U(f)] \quad (2.4)$$

for a non-concave utility U , where we set $E_n[U(f)] := -\infty$ if $U(f)^- \notin L^1$. We refer to this case as *expected non-concave utility* (ENCU).

The functional V_n defined in (2.3) can be seen as a Choquet integral $\int U(f) d(T \circ P^n)$ with respect to the monotone set function $T \circ P^n$; see Chapter 5 of Denneberg [39] for an exposition of this concept. In the case $T(p) = p$, the functional V_n in (2.3) coincides with the classical expected utility $E_n[U(f)]$ in (2.4) for a positive non-concave utility U . We distinguish the two cases since the conditions for their treatments will be different.

Finally, we formulate the sequence of optimization problems. For a fixed (initial capital) $x > 0$, the (budget) set $C^n(x)$ in the n -th model is

$$C^n(x) := \{f \in L_+^0(\Omega^n, \mathcal{F}^n, P^n) \mid E_{Q^n}[f] \leq x\}.$$

For each model, we are interested in the demand problem

$$v^n(x) := \sup\{V_n(f) \mid f \in C^n(x)\}. \quad (2.5)$$

An element $f \in C^n(x)$ is *optimal* if $V_n(f) = v^n(x)$. By a *maximizer* for $v^n(x)$, we mean an optimal element for the optimization problem (2.5).

III.2.2 Main results

Even in the classical case of expected concave utility, the stability of the utility maximization problem is only obtained under suitable growth conditions on U (or its conjugate J). In the case of the RDEU functional in (2.3), the corresponding assumption has to be imposed jointly on U and T .

Assumption 2.7. *We suppose that*

$$U(x) \leq k_1 x^\gamma + k_2, \quad (2.6)$$

$$T(p) \leq k_3 p^\alpha, \quad (2.7)$$

$$\gamma < \alpha \quad (2.8)$$

with $\gamma, \alpha \in (0, 1]$ and $k_1, k_2, k_3 > 0$. This allows us to find and fix λ such that $\lambda\alpha > 1$ and $\lambda\gamma < 1$.

This assumption is inspired by Assumption 4.1 in Carassus and Rásonyi [30]. For the example distortion in (2.2), condition (2.7) is satisfied. In the case without distortion, $T(p) = p$, (2.7) is satisfied for $k_3 = \alpha = 1$. A sufficient condition for (2.6) is $AE_0(J) < \infty$ (see Lemma 7.3). For later reference, we summarize the case-dependent assumptions.

Assumption 2.8. *We assume we have one of the following cases:*

RDEU: *Let V_n be defined as in (2.3). In this case, we suppose that the distribution of φ^0 is continuous and that Assumption 2.7 is satisfied.*

ENCU: *Let V_n be defined as in (2.4). In this case, we suppose that Assumption 2.1, $U(\infty) > 0$ and $AE_0(J) < \infty$ are satisfied.*

We are now in a position to formulate the main result of this chapter. Note that this covers simultaneously both cases.

Theorem 2.9. *Let Assumptions 2.2 and 2.8 be satisfied. Then*

$$\lim_{n \rightarrow \infty} v^n(x) = v^0(x),$$

and for any sequence of maximizers f^n for $v^n(x)$, there are a subsequence (n_k) and a maximizer \bar{f} for $v^0(x)$ such that $f^{n_k} \Rightarrow \bar{f}$ as $k \rightarrow \infty$.

The maximizers for $v^0(x)$ are not necessarily unique (see Example II.3.7). Weak convergence along a subsequence of maximizers is therefore the best we can hope for. Moreover, note that for the second statement in Theorem 2.9, we start with a sequence of maximizers f^n for $v^n(x)$. For the ENCU functional in (2.4), the existence of a maximizer f^n for $v^n(x)$ is guaranteed under the present assumptions (see Theorem II.3.4). For the RDEU functional in (2.3), on the other hand, the existence of a maximizer for $v^n(x)$ has to be verified in any given setting. One sufficient criterion is that $(\Omega^n, \mathcal{F}^n, P^n)$ (or $(\Omega^n, \mathcal{F}^n, Q^n)$ due to the equivalence of Q^n and P^n) is atomless (see Remark 4.5 below). Another sufficient criterion is that $(\Omega^n, \mathcal{F}^n, P^n)$ consists of finitely many atoms. The latter in fact implies that any maximizing sequence (f^n) for $v^n(x)$ is bounded; this allows us to extract a subsequence a.s. converging to some limit \bar{f} , and arguments similar to the ones in Proposition 4.4 show that \bar{f} is a maximizer for $v^0(x)$. These two criteria cover all the examples discussed in Section III.3.

The assumption that U is strictly increasing and continuous is not strictly necessary; it avoids some (more) technical details. Let us shortly discuss a relevant excluded special case.

Remark 2.10. The ENCU functional defined in (2.4) do not cover the piecewise constant function $U(x) := 1_{\{x \geq 1\}}$ which describes the *goal-reaching problem* initiated by Kulldorff [76] and investigated extensively by Browne [27, 28]. But under the assumption that φ^0 has a continuous distribution, one can slightly adapt the arguments in its proof to show that the results of Theorem 2.9 also hold for the goal-reaching problem. We provide a detailed argument at the end of Section III.5. \diamond

III.2.3 The need for Assumption 2.1

For expected *concave* utilities, Assumption 2.1 is not necessary to obtain Theorem 2.9; see Proposition 5.4 below. However, for non-concave utilities, Assumption 2.1 cannot be dropped. The difference between these cases can be explained as follows. For a risk-averse agent with a concave U , the optimal final position is (essentially) $\sigma(\varphi)$ -measurable, and so it is enough to have convergence in distribution of the sequence of pricing densities. For risk-seeking agents, the optimal final position is not necessarily $\sigma(\varphi)$ -measurable. Additional information (if available) is used by the agent to avoid the non-concave part $\{U < U_c\}$ of U . In the atomless limit model, every payoff distribution can be supported, and Assumption 2.1 ensures that also the models along the sequence become sufficiently rich as $n \rightarrow \infty$. More concretely, Assumption 2.1 excludes the (pathological) behaviour illustrated in the next example.

Example 2.11. Consider a non-concave utility U with $\{U < U_c\} = (a, b)$ which is strictly concave on $(0, a)$ and (b, ∞) . The initial capital x is in (a, b) , but not exactly in the middle of the interval (a, b) . The probability spaces $(\Omega^n, \mathcal{F}^n, P^n)_{n \in \mathbb{N}}$ are all given by the same probability space consisting of two states with $P^n[\{\omega_1\}] = P^n[\{\omega_2\}] = 1/2$; and $(\Omega^0, \mathcal{F}^0, P^0)$ is an arbitrary atomless probability space. Set $\varphi^n \equiv 1$ for all $n \in \mathbb{N}_0$. Jensen's inequality and Theorem II.5.1 give $v^0(x) = U_c(x)$. On the other hand, we have $v^n(x) = v^1(x)$ for every $n \in \mathbb{N}$ and we now show that $v^1(x) < U_c(x)$ for x chosen above. First, note that $v^1(x)$ admits a maximizer \hat{f} since the model consists of two atoms (see the discussion following Theorem 2.9). The maximizer \hat{f} satisfies $E_Q[\hat{f}] = x$ since U is strictly increasing, so we can replace $\hat{f}(\omega_2)$ by $2x - \hat{f}(\omega_1)$. We therefore get the inequality

$$v^1(x) = \frac{1}{2}U(\hat{f}(\omega_1)) + \frac{1}{2}U(2x - \hat{f}(\omega_1)) \leq \frac{1}{2}U_c(\hat{f}(\omega_1)) + \frac{1}{2}U_c(2x - \hat{f}(\omega_1)) \leq U_c(x).$$

The first inequality is an equality if and only if both values $\hat{f}(\omega_1)$ and $\hat{f}(\omega_2)$ are not in $\{U < U_c\} = (a, b)$; the second inequality is an equality if and

only if the two value $\hat{f}(\omega_1)$ and $\hat{f}(\omega_2)$ are in $[a, b]$. But these two conditions cannot be satisfied at the same time by our choice of x . This shows that $v^n(x) = v^1(x) < U_c(x) = v^0(x)$. We conclude that $v^n(x)$ converges as $n \rightarrow \infty$ (it is constant), but the limit is not $v^0(x)$.

III.2.4 Anti-comonotonicity of the optimal investment

In Theorem 2.9, we describe the asymptotics of v^n and the corresponding maximizer f^n for $v^n(x)$. In some recent applications in financial economics (see the discussion Chapter V), however, one is interested in the relation between f^n and φ^n . In the limit model, it turns out that every maximizer f^0 for $v^0(x)$ and the pricing density φ^0 are anti-comonotonic, which is a very strong form of dependence. The approximating models, however, are allowed to contain atoms and for such models, the maximizer f^n for $v^n(x)$ and φ^n are not necessarily anti-comonotonic (see Example II.3.9). The goal of this section is to introduce a measure $\delta(X, Y)$ to quantify the dependence between two random variables X and Y and to show that $\delta_n(\varphi^n, f^n)$ converges to $\delta_0(\varphi^0, f^0)$.

In order to define the measure δ_n , let $n \in \mathbb{N}_0$ be fixed and X and Y two positive random variables on $(\Omega^n, \mathcal{F}^n, P^n)$ with quantiles q_X and q_Y . We define

$$\delta_n(X, Y) := E_n[XY] - \int_0^1 q_X(s)q_Y(1-s)ds.$$

If X and Y have finite mean, the measure $\delta_n(X, Y)$ can be rewritten as $\text{COV}_n(X, Y) - (\int_0^1 q_X(s)q_Y(1-s)ds - E_n[X]E_n[Y])$ which shows that up to normalisation (and assuming that X and Y have strictly positive and finite variance), $\delta_n(X, Y)$ is equal to the difference between the correlation between X and Y and the minimal attainable correlation between the laws of X and Y (see Section 5.2 of McNeil et al. [83] for a definition and applications of the minimal attainable correlation). Let us first formulate two useful properties.

Lemma 2.12. *Let $n \in \mathbb{N}_0$ be fixed and X and Y two positive random variables on $(\Omega^n, \mathcal{F}^n, P^n)$ with quantiles q_X and q_Y . Then $\delta_n(X, Y) \geq 0$, and equality holds if and only if X and Y are anti-comonotonic.*

Proof. The inequality is proved in Theorem A.24 of Föllmer and Schied [43]; the second part is a well-known property of anti-comonotonic random variables (see, for instance, Theorem 5.25 of McNeil et al. [83]). \square

The next result, a corollary of Theorem 2.9, shows that the dependence of φ^n and f^n in terms of δ_n converges to the dependence of φ^0 and f^0 .

Corollary 2.13. *Under the assumptions of Theorem 2.9, we have that*

$$\lim_{n \rightarrow \infty} \delta_n(\varphi^n, f^n) = \delta_0(\varphi^0, f^0) = 0. \quad (2.9)$$

Proof. The second equality follows from the “if and only if” part of Lemma 2.12 and anti-comonotonicity of f^0 and φ^0 . The latter is proved in Corollary 5.5 for the preference functional in (2.4). For the one in (2.3), it essentially follows from part (ii) of Theorem B.1 in Jin and Zhou [64]; Lemma 2.2 of Carlier and Dana [31] provides a similar statement.

For the first equality in (2.9), the inequality “ \geq ” follows from Lemma 2.12. We show below that

$$\liminf_{n \rightarrow \infty} \int_0^1 q_{\varphi^n}(1-s)q_{f^n}(s)ds \geq x. \quad (2.10)$$

We then get “ \leq ” in the first equality in (2.9) since for every $n \in \mathbb{N}$, the maximizer f^n for $v^n(x)$ satisfies $E_n[\varphi^n f^n] = x$; this uses that U as well as T are strictly increasing.

We now show (2.10). By passing to a (relabelled) subsequence that realizes the lim inf, we can assume that the sequence $\int_0^1 q_{\varphi^n}(1-s)q_{f^n}(s)ds$ converges. By assumption, we know that $\varphi^n \Rightarrow \varphi^0$, hence also $q_{\varphi^n}(s) \rightarrow q_{\varphi^0}(s)$ a.e. Moreover, applying Theorem 2.9 to the relabelled subsequence yields $f^{n_k} \Rightarrow f^0$ for a further subsequence (n_k) ; this also gives $q_{f^{n_k}}(s) \rightarrow q_{f^0}(s)$ a.e. Then we apply Fatou’s lemma on $(q_{\varphi^{n_k}}(1-s)q_{f^{n_k}}(s))_k$ to obtain

$$\begin{aligned} \liminf_{n \rightarrow \infty} \int_0^1 q_{\varphi^n}(1-s)q_{f^n}(s)ds &= \lim_{k \rightarrow \infty} \int_0^1 q_{\varphi^{n_k}}(1-s)q_{f^{n_k}}(s)ds \\ &\geq \int_0^1 q_{\varphi^0}(1-s)q_{f^0}(s)ds. \end{aligned}$$

This completes the proof of (2.10) since the last term is equal to $E_0[\varphi^0 f^0] = x$ due to anti-comonotonicity of f^0 and φ^0 and Lemma 2.12. \square

Remark 2.14. Let us mention one special case where our results yield another property that can be interpreted as *almost anti-comonotonicity*. We consider the case where the distribution of φ^0 is continuous and the preference functional is defined by $V_n(f) := E_n[U(f)]$ for a strictly increasing and continuously differentiable non-concave utility function U with $\lim_{x \searrow 0} U'(x) = \infty$ and $\{U < U_c\} = \cup_{i=1}^m (a_i, b_i)$ for some finite m . This includes, for instance, the classical Friedman–Savage utility (Friedman and Savage [45]). We now show that *for each $\epsilon > 0$ and all $n > n_0(\epsilon)$, there exists a set $A^n \in \mathcal{F}^n$ with $P^n[A^n] \geq 1 - \epsilon$ such that φ^n and f^n are anti-comonotonic on A^n* . We start with the observation that the maximizer f^n for v^n is strictly positive (because of the assumption that $\lim_{x \searrow 0} U'(x) = \infty$) and satisfies $U'(f^n) = \lambda^n \varphi^n$ for some $\lambda^n > 0$ (Lemma II.3.8). For $f^n(\omega) \notin \{U < U_c\}$, we have that $U'(f^n(\omega)) = U'_c(f^n(\omega))$. Indeed, this is clear for the interior of $\{U < U_c\}$ and extends to the boundary points by continuous differentiability of U (and U_c). But U'_c is decreasing, so we find that f^n and φ^n are anti-comonotonic on $A^n := \{f^n \notin \{U < U_c\}\}$. We claim that $P^n[A^n]$

converges to 1 as n tends to ∞ . First, note that $U'(f^0) = \lambda^0 \varphi^0$ gives $\{f^0 = a\} \subset \{\lambda^0 \varphi^0 = U'_c(a)\}$ for all $a > 0$. Continuity of the distribution of φ^0 therefore implies continuity of the one of f^0 . Theorem 2.9, the continuity of the distribution of f^0 and $P^0[f^0 \in \{U < U_c\}] = 0$ (see Theorem II.5.1) finally give

$$\begin{aligned} \lim_{n \rightarrow \infty} P^n[f^n \in \{U < U_c\}] &= \sum_{i=1}^m \lim_{n \rightarrow \infty} (P^n[f^n < b_i] - P^n[f^n \leq a_i]) \\ &\leq \sum_{i=1}^m P^0[f^0 < b_i] - P^0[f^0 \leq a_i] = 0, \end{aligned}$$

which shows that $P^n[A^n] = 1 - P^n[f^n \in \{U < U_c\}] \rightarrow 1$ as $n \rightarrow \infty$. \diamond

III.3 Applications

So far, our analysis has been conducted for an abstract sequence of models. We now present three types of application to illustrate the main results. We also provide a numerical example in Section III.3.4 to visualize our results.

III.3.1 (Numerical) computation of the value function

In recent years, there has been remarkable progress in the problem of behavioural portfolio selection. In particular, there are several new results for complete markets in continuous time. Most of them give the existence of a solution and describe the structure of the optimal final position as a decreasing function of the pricing density. While these results are interesting from a theoretical point of view, they are less helpful for explicit computations. In this section, we show how Theorem 2.9 can be used to determine the value function numerically for a complete model in continuous time.

The idea is to approximate the (computationally difficult) continuous-time model by a sequence of (computationally tractable) discrete-time models. We illustrate this for the Black–Scholes model which can be approximated by a sequence of binomial models. This is the typical example for the transition from a discrete- to a continuous-time setting. In this example, the limit model is atomless while the approximating models are not. He [55] and Prigent [88] analyze the stability for expected *concave* utilities for this setting by directly analyzing the sequence of optimal terminal wealths f^n as a function of φ^n . This is possible due to the concavity of their utility function, but cannot be used here.

To fix ideas, we briefly recall the (classical) binomial approximation of the Black–Scholes model to verify that our assumptions in Section III.2 are satisfied. We consider a time horizon $T \in (0, \infty)$, a probability space $(\Omega^0, \mathcal{F}^0, P^0)$ on which there is a standard Brownian motion $W = (W_t)_{t \geq 0}$, and a (discounted) market consisting of a savings account $B \equiv 1$ and one stock S

described by

$$\frac{dS_t}{S_t} = \mu dt + \sigma dW_t, \quad S_0 = s_0 > 0, \sigma > 0,$$

in the filtration generated by W . The pricing density is then given by

$$\varphi^0 := \exp\left(-\frac{\mu}{\sigma}W_T - \frac{1}{2}\left(\frac{\mu}{\sigma}\right)^2 T\right).$$

For the construction of the n -th approximation, we start with a probability space $(\Omega^n, \tilde{\mathcal{F}}^n, P^n)$ on which we have independent and identically distributed random variables $(\tilde{\epsilon}_k)_{k=1, \dots, n}$ taking values 1 and -1 , both with probability $\frac{1}{2}$. For any n , we consider the n -step market consisting of the savings account $B^{(n)} \equiv 1$ and the stock given by $S_t^{(n)} = S_0$ for $t \in [0, \frac{T}{n})$ and

$$S_t^{(n)} = S_0 \prod_{j=1}^k \left(1 + \frac{\mu T}{n} + \sqrt{\frac{T}{n}} \sigma \tilde{\epsilon}_j\right), \quad \frac{kT}{n} \leq t < \frac{(k+1)T}{n}, \quad k = 1, \dots, n.$$

This process is right-continuous with left limits. The filtration generated by $S^{(n)}$ is denoted by $\mathbb{F}^n = (\mathcal{F}_t^n)_{0 \leq t \leq T}$ and we take $\mathcal{F}^n := \mathcal{F}_T^n = \sigma(\tilde{\epsilon}_1, \dots, \tilde{\epsilon}_n)$. The market is active at the times $0, \frac{T}{n}, \frac{2T}{n}, \dots, T$. It is well known that this market is complete, and we denote by Q^n the unique martingale measure. The martingale condition implies that

$$Q^n[\tilde{\epsilon}_k = 1] \left(1 + \frac{\mu T}{n} + \frac{\sigma \sqrt{T}}{\sqrt{n}}\right) + (1 - Q^n[\tilde{\epsilon}_k = 1]) \left(1 + \frac{\mu T}{n} - \frac{\sigma \sqrt{T}}{\sqrt{n}}\right) = 1,$$

and solving gives $Q^n[\tilde{\epsilon}_k = 1] = \frac{1}{2}(1 - \frac{\mu \sqrt{T}}{\sigma \sqrt{n}})$. This is positive for n large enough and we only consider such n from now on. The measures Q^n and P^n are equivalent on $(\Omega^n, \mathcal{F}^n)$ and we denote the pricing density by $\varphi^n := dQ^n/dP^n$. It is shown in He [54], Theorem 1 that $\varphi^n \Rightarrow \varphi^0$; this is a consequence of the central limit theorem. The set \mathcal{G}^n of atoms in $(\Omega^n, \mathcal{F}^n)$ can be identified with the paths of $S^{(n)}$. The Q^n -probability for a path is of the form

$$\left(\frac{1}{2}\right)^n \left(1 - \frac{\mu \sqrt{T}}{\sigma \sqrt{n}}\right)^k \left(1 + \frac{\mu \sqrt{T}}{\sigma \sqrt{n}}\right)^{n-k}$$

for some $k \in \{0, 1, \dots, n\}$ (which is the number of up moves in the path). For n large, we have $|\mu/(\sigma \sqrt{n})| < \frac{1}{2}$ and we see that $\sup_{A \in \mathcal{G}^n} Q^n[A] < (3/4)^n$. Taking the limit $n \rightarrow \infty$ gives $\lim_{n \rightarrow \infty} \sup_{A \in \mathcal{G}^n} Q^n[A] = 0$, which means that Assumption 2.1 is satisfied. The distribution of φ^0 is continuous if $\mu \neq 0$. Finally, a proof of uniform integrability of $((\varphi^n)^\xi)_{n \in \mathbb{N}}$ for $\xi < 0$ can be found in Prigent [88, Page 172, Lemma c].

We conclude that all the assumptions of Section III.2 are satisfied. We can therefore apply Theorem 2.9 to relate the optimization problem in the

Black–Scholes model with the sequence of optimization problems in the sequence of binomial models. More precisely, Theorem 2.9 shows that the sequence of value functions (v^n) in the binomial models converges to the value function v^0 in the Black–Scholes model and that the sequence of maximizers converges along a subsequence. Particularly in the case of the preference functional (2.3), this turns out to be useful for computational purposes: While there are (abstract) results on the existence of a maximizer in the Black–Scholes model in the literature, these results are less helpful to determine a maximizer and the corresponding value explicitly. In the binomial model, however, the (numerical) computation of the value function and its maximizers is straightforward since the model consists of a finite number of atoms. In this context, Theorem 2.9 provides the insight that we can use the value functions in the binomial model to approximate the value function in the Black–Scholes model. This gives a method to determine numerically the value function v_0 in the Black–Scholes model. Note that for *rank-dependent expected utility*, we have no dynamic programming and hence no description of v_0 by a (HJB) PDE we could solve numerically.

Remark 3.1. The Black–Scholes model can also be approximated by a sequence of trinomial models. One possible choice for the approximating sequence is to choose $S_{k+1}^n = S_k^n(1 + \frac{\mu T}{n} + \xi_{k+1}^n)$ for $k = 1, \dots, n-1$ and independent and identically distributed random variables $(\xi_1^n, \dots, \xi_n^n)$ which assumes the three values

$$U^n := \frac{\sigma\sqrt{3T}}{\sqrt{2n}}, \quad M^n := 0, \quad D^n := -\frac{\sigma\sqrt{3T}}{\sqrt{2n}}$$

each with a probability of $1/3$. The subfamily of those equivalent martingale measures Q^n under which the process $(S_k^n)_{k=1, \dots, n}$ is again a homogeneous trinomial model can be characterized by the probabilities

$$Q^n[\xi_k^n = U^n] = \frac{1 - \lambda^n}{2} - \frac{\mu\sqrt{T}}{\sigma\sqrt{6n}}, \quad Q^n[\xi_k^n = D^n] = \frac{1 - \lambda^n}{2} + \frac{\mu\sqrt{T}}{\sigma\sqrt{6n}}$$

and $Q^n[\xi_k^n = M^n] = \lambda^n$ with $0 < \lambda^n < 1 - 2\mu\sqrt{T}/\sigma\sqrt{6n}$. A sequence (Q^n) of such martingale measures now corresponds to a sequence (λ^n) . For a generic sequence (λ^n) , the associated sequence (dQ^n/dP^n) need not converge weakly to the pricing density dQ^0/dP^0 in the limit model.

Bizid et al. [20], Jouini and Napp [67] and Jouini [65] suggest considering only those martingale measures whose densities are “in reverse order” of the total wealth of the economy since these martingale measures are consistent with a “completion” of the market with assets that are in zero net supply and whose prices are determined in a standard equilibrium framework. For the above choice of the approximating sequence, this leads to the condition

$$\frac{1 - \lambda^n}{2} - \frac{\mu\sqrt{T}}{\sigma\sqrt{6n}} \leq \lambda^n \leq \frac{1 - \lambda^n}{2} + \frac{\mu\sqrt{T}}{\sigma\sqrt{6n}}.$$

This implies $\lambda^n \rightarrow 1/3$ as $n \rightarrow \infty$. The corresponding one-step conditional densities have mean 1 and variance $\frac{\mu^2 T}{\sigma^2 n} + \frac{9}{2}(\lambda^n - \frac{1}{3})^2$. If the convergence $\lambda^n \rightarrow 1/3$ is sufficiently fast (e.g. $|\lambda^n - 1/3| \leq C/n^p$ for some $p > 1/2$ and some constant C), then one can check via the central limit theorem that the associated sequence of martingale measures converges weakly to the pricing density in the limit model. Similarly to the binomial case above, one can also check that Assumption 2.2 is satisfied for such a sequence.

The economic interpretation of this remark is as follows. We approximate a complete financial market by a sequence of incomplete financial markets. In every approximating model, there are infinitely many martingale measures and any of these corresponds to a “completion” of the market. If we choose for each approximating model an arbitrary “completion”, the resulting sequence of models generally has bad convergence properties, in the sense that (not surprisingly) the sequence of portfolio selection problems need not converge in general. However, if one chooses an economically reasonable “completion”, one does obtain good convergence properties for the sequence of portfolio selection problems, in the sense that all of our results from Section III.2 can be applied. \diamond

III.3.2 Stability results

In this section, we use Theorem 2.9 to show, as explained in the introduction, the stability of the portfolio choice results for a model with respect to small perturbations. While Section III.3.1 can be seen as perturbation of the underlying model itself, we are interested here in perturbations of a model’s parameters.

III.3.2.1 Misspecifications of the market model

The first example is motivated by the practical difficulties one encounters when trying to precisely calibrate an underlying model. In this section, we analyze how the optimal final position and the corresponding value are affected by a (small) misspecification of the underlying market model.

This question is well studied for expected *concave* utilities; see for instance Larsen and Žitković [79] and Kardaras and Žitković [72]. Their sequence of model classes is very general in the sense that they need not restrict the setup to a single pricing density. However, the key to solving their problem is the classical duality theory which can be applied since their utilities are (strictly) concave. In our setting, this is not possible.

As one example in our framework, we can think of the Black–Scholes model where it is generally difficult to measure the drift. To formalize this situation, we fix some time horizon $T \in (0, \infty)$ and a probability space $(\Omega^0, \mathcal{F}^0, P^0)$ on which there is a Brownian motion $W = (W_t)_{t \geq 0}$. We introduce the sequence of probability spaces by $(\Omega^n, \mathcal{F}^n, P^n) := (\Omega^0, \mathcal{F}^0, P^0)$ for

$n \in \mathbb{N}$. In order to define a sequence of price processes, we consider a sequence (μ^n) converging to some drift parameter $\mu^0 \in \mathbb{R}$ in the limit model. For each n , we consider a (discounted) market consisting of a savings account $B \equiv 1$ and one stock S^n described by

$$\frac{dS_t^n}{S_t^n} = \mu^n dt + \sigma dW_t, \quad S_0^n = s_0 > 0, \sigma > 0,$$

in the filtration generated by W . The pricing density for the n -th model is then given by

$$\varphi^n := \exp\left(-\frac{\mu^n}{\sigma} W_T - \frac{1}{2} \left(\frac{\mu^n}{\sigma}\right)^2 T\right).$$

In this example, each model is atomless. Assumption 2.1 is therefore trivially satisfied. Moreover, looking at the explicit form of φ^n shows that $\varphi^n(\omega) \rightarrow \varphi^0(\omega)$ for all $\omega \in \Omega$, so in particular $\varphi^n \Rightarrow \varphi^0$. In order to check the uniform integrability of $((\varphi^n)^\xi)$ for any $\xi < 0$, we use

$$\begin{aligned} E_n[|(\varphi^n)^\xi|^{1+\epsilon}] &= E_n\left[\exp\left(-\frac{\mu^n}{\sigma}(1+\epsilon)\xi W_T - \frac{1}{2}\left(\frac{\mu^n}{\sigma}(1+\epsilon)\xi\right)^2 T\right)\right] \\ &\quad \times \exp\left(\frac{T}{2}\left(\frac{\mu^n}{\sigma}\right)^2((1+\epsilon)^2\xi^2 - (1+\epsilon)\xi)\right) \\ &= \exp\left(\frac{T}{2}\left(\frac{\mu^n}{\sigma}\right)^2((1+\epsilon)^2\xi^2 - 1 + \epsilon)\xi\right) \end{aligned}$$

to obtain $\sup_n E_n[|(\varphi^n)^\xi|^{1+\epsilon}] < \infty$. Hence Assumption 2.2 is satisfied.

We conclude that the assumptions of Section III.2 are satisfied and we can apply the results there. Theorem 2.9 tells us that the value functions (as well as the corresponding maximizers along a subsequence) for the model with drift μ^n converge to the corresponding quantities in the model with drift μ^0 . The economic interpretation of this result is that the behavioural prediction does not change drastically if we slightly perturb the drift.

It is also worth mentioning that the above arguments only use convergence of the market price of risk μ^n/σ to μ^0/σ . If we consider more generally a stochastic market price of risk $\lambda_t^n = \mu_t^n/\sigma_t^n$, then assuming $E_n[\int_0^T (\lambda_t^n)^2 dt] \rightarrow E_0[\int_0^T (\lambda_t^0)^2 dt]$ gives weak convergence of the stochastic exponential $\mathcal{E}(\int_0^\cdot \lambda_t^n dW_t)_T$ to $\mathcal{E}(\int_0^\cdot \lambda_t^0 dW_t)_T$ (Proposition A.1 in Larsen and Žitković [79]). In addition, one then needs some integrability condition on λ^n to ensure that the family $((\mathcal{E}(\int_0^\cdot \lambda_t^n dW_t)_T)^\xi)_{n \in \mathbb{N}}$ is uniformly integrable for $\xi < 0$; for instance a nonrandom upper bound for all the $\int_0^T (\lambda_t^n)^2 dt$ is sufficient.

In the present setting, the limit model as well as the approximating sequence are given by atomless models. For this class of models and for the ENCU functional (2.4), the optimization problem v^n can be reduced to the concavified utility maximization problem (see Theorem II.5.1). In this way, the stability result can also be obtained via stability results for expected *concave* utilities. For the RDEU functional (2.3) with distortion, however, the results are new.

III.3.2.2 Horizon dependence

In Section III.3.1 as well as in the first example in this section, we have started with a fixed time horizon T . In practical applications, however, the time horizon might be uncertain or changing. The goal of this section is to use Theorem 2.9 to study whether a (marginal) misspecification of the investment horizon significantly influences the optimal behaviour of the agent. For expected *concave* utilities, Larsen and Yu [78] analyze this question in an incomplete Brownian setting. The key to solving their problem is again the duality theory which cannot be used in our setup.

In order to formalize a similar situation in our framework, we start again with a probability space $(\Omega^0, \mathcal{F}^0, P^0)$ on which there is a standard Brownian motion $W = (W_t)_{t \geq 0}$, and we introduce the sequence of probability spaces by setting $(\Omega^n, \mathcal{F}^n, P^n) := (\Omega^0, \mathcal{F}^0, P^0)$ for $n \in \mathbb{N}$. We now fix a sequence $T^n \rightarrow T^0 \in (0, \infty)$ representing the time horizons. For each n , we consider the Black–Scholes model with time horizon T^n as described in Section III.3.1. The pricing density for the n -th model is therefore given by

$$\varphi^n := \exp\left(-\frac{\mu}{\sigma}W_{T^n} - \frac{1}{2}\left(\frac{\mu}{\sigma}\right)^2 T^n\right).$$

Assumption 2.1 is again trivially satisfied since each model is atomless. Moreover, adapting the arguments from Section III.3.2.1 shows that $\varphi^n \Rightarrow \varphi^0$ and Assumption 2.2 are satisfied as well. As in Section III.3.2.1, we can therefore use Theorem 2.9 to conclude that behavioural predictions of the model are stable with respect to small misspecifications in the time horizon.

Remark 3.2. The examples so far constitute complete financial markets. For these examples, it is well known that and how the portfolio optimization problem can be reduced to a problem of the form (2.5). However, there are also more general incomplete settings where our main result can be applied. He and Zhou [56] analyze the portfolio optimization problem for an incomplete Brownian setting with a deterministic opportunity set and show that the solution for the portfolio optimization problem can also be obtained by only working with the density of the minimal martingale measure. The results of Sections III.3.2.1 and III.3.2.2 then carry over to the new setup as soon as we have weak convergence of the densities of the minimal martingale measure in the approximating model to the one in the limit model. In particular, this shows that our main result is not limited to complete market models. \diamond

III.3.3 Theoretical applications

For the ENCU functional defined in (2.4) without distortion, the problem v^n turns out to be very tractable if φ^n has a continuous distribution. In this section, we apply Theorem 2.9 to approximate a pricing density with a

general distribution by a pricing density with a continuous distribution. To explain the idea in more detail, we use the notation

$$u^n(x, U) := \sup\{E_n[U(f)] \mid f \in C^n(x)\}.$$

Every maximizer f^n for $u^n(x, U_c)$ satisfies $f^n \in -\partial J(\lambda^n \varphi^n)$ for some λ^n ; see Proposition 7.1. If φ^n has a continuous distribution, then we have $P^n[f^n \in \{U < U_c\}] = 0$ (see Lemma II.5.7 for details) and it follows that $u^n(x, U) = u^n(x, U_c)$. In this way, the existence of a maximizer as well as several properties of $u^n(x, U)$ can be derived directly via the concavified problem. If the limit model is atomless but the distribution of the pricing density φ^0 is not continuous, then this reduction does not follow directly.

The idea now is to construct a sequence (φ^n) weakly converging to φ^0 for which each φ^n has a continuous distribution. For this approach, we assume that $E_0[(\varphi^0)^\eta] < \infty$ for all $\eta < 0$. Since $(\Omega^0, \mathcal{F}^0, P^0)$ is atomless, we can find a uniformly distributed random variable \mathcal{U} such that $q_{\varphi^0}(\mathcal{U}) = \varphi^0$ P^0 -a.s. (Lemma A.28 in Föllmer and Schied [43]). Moreover, we choose another random variable $Y > 0$ with $E_0[Y] = 1$ having a continuous distribution (e.g. $Y = \mathcal{U} + 1/2$). Now we define the sequence (φ^n) by

$$\varphi^n := (1 - \frac{1}{n})q_{\varphi^0}(\mathcal{U}) + \frac{1}{n}q_Y(\mathcal{U}) = (1 - \frac{1}{n})\varphi^0 + \frac{1}{n}q_Y(\mathcal{U}).$$

Every element satisfies $E_0[\varphi^n] = (1 - \frac{1}{n})E_0[q_{\varphi^0}(\mathcal{U})] + \frac{1}{n}E_0[q_Y(\mathcal{U})] = 1$ by construction. Moreover, the function $h_n(x) := (1 - \frac{1}{n})q_{\varphi^0}(x) + \frac{1}{n}q_Y(x)$ converges pointwise to $h(x) := q_{\varphi^0}(x)$. The set D_h of all points where h is not continuous is at most countable since h is increasing; and \mathcal{U} has a continuous distribution. So it follows that $P^0[\mathcal{U} \in D_h] = 0$ and we obtain $\varphi^n = h_n(\mathcal{U}) \Rightarrow h(\mathcal{U}) = \varphi^0$; see Theorem 5.1 of Billingsley [18].

With the arguments so far, we have a sequence of probability spaces defined by $(\Omega^n, \mathcal{F}^n, P^n) := (\Omega^0, \mathcal{F}^0, P^0)$ for $n \in \mathbb{N}$ together with a sequence of pricing measures (φ^n) weakly converging to φ^0 . To verify Assumption 2.2, note that $\varphi^n \geq \varphi^0(1 - \frac{1}{n})$ gives $(\varphi^n)^{\eta(1+\epsilon)} \leq ((1 - \frac{1}{n})\varphi^0)^{\eta(1+\epsilon)}$ for every $\eta < 0$, which gives a uniformly integrable upper bound due to our assumption that $E_0[(\varphi^0)^\eta] < \infty$ for all $\eta < 0$.

It remains to show that the distribution of $\varphi^n = h_n(\mathcal{U})$ is continuous. Since h_n is increasing, it follows that $h_n(x_1) = h_n(x_2) = k$ if and only if $q_{\varphi^0}(x_1) = q_{\varphi^0}(x_2)$ and $q_Y(x_1) = q_Y(x_2)$. But $q_Y(\cdot)$ is strictly increasing since Y has a continuous distribution, so we infer that $P^0[h_n(\mathcal{U}) = k] = 0$ for $k \in \mathbb{R}$.

Theorem 2.9 now gives $u^n(x, U) \rightarrow u^0(x, U)$ as $n \rightarrow \infty$. Since the distribution of φ^n is continuous for each n , we have that $u^n(x, U) = u^n(x, U_c)$ for all n and we also get $u^0(x, U) = u^0(x, U_c)$ in the limit. In this way, we recover Theorem II.5.1 under slightly less general assumptions, but with completely different techniques. Instead of rearrangement techniques as in Chapter II, here we approximate the mass points in the distribution of φ^0 by continuous distributions and apply Theorem 2.9.

III.3.4 A numerical illustration

The goal of this section is to illustrate the convergence result numerically. We consider the functional $V_n(f) := E_n[U(f)]$ (for a specific non-concave utility) in the framework presented in Section III.3.1 where we can derive $v^0(x)$ explicitly so that we can compare $v^0(x)$ with the value functions $v^n(x)$ in the approximating models. As in Section III.3.3, we use the notation $u^n(x, U) := v^n(x)$.

The utility function in this example is given by

$$U(x) := \begin{cases} \ln x, & x < 1, \\ x - \frac{4 + \cos(x-1)}{3}, & 1 \leq x < 2\pi + 1, \\ 2\pi + \ln(x - 2\pi), & 2\pi + 1 \leq x. \end{cases}$$

This function is strictly increasing, continuous, in C^1 and satisfies the Inada conditions at 0 and ∞ . Its concave envelope is given by

$$U_c(x) = \begin{cases} \ln x, & x < 1, \\ x - 1, & 1 \leq x < 2\pi + 1, \\ 2\pi + \ln(x - 2\pi), & 2\pi + 1 \leq x, \end{cases}$$

and the conjugate of U (and U_c) is

$$J(y) = \begin{cases} -\ln y - 1 + 2\pi(1 - y), & y < 1, \\ -\ln y - 1, & y \geq 1. \end{cases}$$

The conjugate satisfies $AE_0(J) < \infty$. On $(0, 1) \cup (1, \infty)$, the conjugate is differentiable and ∂J is a singleton. More precisely, we have

$$-\partial J(y) = \begin{cases} 2\pi + \frac{1}{y}, & y < 1, \\ (1, 1 + 2\pi), & y = 1, \\ \frac{1}{y}, & y > 1. \end{cases}$$

Figure III.1 shows U and U_c as well as the conjugate J .

Let us now determine $u^0(x, U)$. Recall from Section III.3.1 that

$$\varphi^0 = \exp(-\zeta W_T - \zeta^2 T/2),$$

where $\zeta = \mu/\sigma$ and T are fixed. For simplicity, we assume that $\mu \neq 0$. We now consider some $f \in -\partial J(\lambda\varphi^0)$ for some $\lambda > 0$. Plugging in the above particular form of ∂J , using the fact that $\{\lambda\varphi^0 = 1\}$ has P^0 -measure 0 for any $\lambda > 0$, and doing some elementary calculations gives

$$\begin{aligned} E_0[\varphi^0 f] &= E_0 \left[\varphi^0 \left(1_{\{\lambda\varphi^0 < 1\}} \left(2\pi + \frac{1}{\lambda\varphi^0} \right) + 1_{\{\lambda\varphi^0 = 1\}} f + 1_{\{\lambda\varphi^0 > 1\}} \frac{1}{\lambda\varphi^0} \right) \right] \\ &= \frac{1}{\lambda} + 2\pi E_0[\varphi^0 1_{\{\lambda\varphi^0 < 1\}}]. \end{aligned}$$




Figure III.1: Panel (A) shows the non-concave utility U and its concave envelope U_c (the dotted line). Panel (B) shows the conjugate J of U and U_c . The conjugate J has a kink in 1 which is the slope of the affine part in U_c in Panel (A).

In the next step, we rewrite the set $\{\lambda\varphi^0 < 1\}$ in a suitable way and use that $(W_t + \zeta t)_{t \geq 0}$ is a Q^0 -Brownian motion (by Girsanov's theorem) to obtain

$$\begin{aligned} E_0[\varphi^0 1_{\{\lambda\varphi^0 < 1\}}] &= Q^0[\lambda\varphi^0 < 1] \\ &= Q^0\left[\frac{W_T + \zeta T}{\sqrt{T}} > -\frac{1}{\zeta\sqrt{T}}\left(\ln\frac{1}{\lambda} + \frac{\zeta^2 T}{2}\right) + \zeta\sqrt{T}\right] \\ &= 1 - \Phi\left(-\frac{1}{\zeta\sqrt{T}}\left(\ln\frac{1}{\lambda} + \frac{\zeta^2 T}{2}\right) + \zeta\sqrt{T}\right), \end{aligned}$$

where $\Phi(\cdot)$ denotes the cumulative distribution function of the standard normal distribution. From this explicit form, we see that $E_0[\varphi^0 1_{\{\lambda\varphi^0 < 1\}}]$ is a continuous and decreasing function of λ with limits 1 and 0 at 0 and ∞ , respectively. The equation

$$x = E_0[\varphi^0 f] = \frac{1}{\lambda} + 2\pi \left(1 - \Phi\left(-\frac{1}{\zeta\sqrt{T}}\left(\ln\frac{1}{\lambda} + \frac{\zeta^2 T}{2}\right) + \zeta\sqrt{T}\right)\right)$$

therefore has a unique solution λ^* . Fix $\hat{f} \in -\partial J(\lambda^*\varphi^0)$. By the definition of λ^* , \hat{f} satisfies $E_0[\varphi^0 \hat{f}] = x$, which means that $\hat{f} \in C^0(x)$. Moreover, \hat{f} satisfies

$$P^0[\hat{f} \in \{U < U_c\}] = P^0[\hat{f} \in (1, 1 + 2\pi)] = P^0[\lambda^*\varphi^0 = 1] = 0,$$

and this gives $E_0[U_c(\hat{f})] = E_0[U(\hat{f})]$. The conjugacy relation between U and J and the explicit form of \hat{f} give

$$E_0[U(f)] \leq E_0[J(\lambda^*\varphi^0)] + x\lambda^* = E_0[U_c(\hat{f})] = E_0[U(\hat{f})]$$




Figure III.2: Value functions $u^n(x, U)$ for parameters $T = 1$, $\mu = 5\%$, $\sigma = 20\%$ and $n = 0, 1, 2, 4$. Recall that $n = 0$ is the limit case.

for all $f \in C^0(x)$ which, together with $\hat{f} \in C^0(x)$, gives optimality of \hat{f} for $u^0(x, U)$. In order to determine $u^0(x, U) = E_0[U(\hat{f})]$, we recall the explicit expression for $\partial J(y)$ and use the fact that $\{\lambda^* \varphi^0 = 1\}$ has measure 0 to get

$$E_0[U(\hat{f})] = E_0[-\ln(\lambda^* \varphi^0)] + 2\pi P^0[\lambda^* \varphi^0 < 1].$$

Elementary calculations show that $E_0[-\ln(\lambda^* \varphi^0)] = -\ln \lambda^* + \zeta^2 T/2$ and

$$\begin{aligned} P^0[\lambda^* \varphi^0 < 1] &= P^0 \left[W_T > -\frac{1}{\zeta} \left(\ln \frac{1}{\lambda^*} + \frac{\zeta^2 T}{2} \right) \right] \\ &= 1 - \Phi \left(-\frac{1}{\zeta \sqrt{T}} \left(\ln \frac{1}{\lambda^*} + \frac{\zeta^2 T}{2} \right) \right). \end{aligned}$$

We conclude that

$$\begin{aligned} u^0(x, U) &= E_0[U(\hat{f})] \\ &= -\ln \lambda^* + \frac{\zeta^2 T}{2} + 2\pi \left(1 - \Phi \left(-\frac{1}{\zeta \sqrt{T}} \left(\ln \frac{1}{\lambda^*} + \frac{\zeta^2 T}{2} \right) \right) \right). \end{aligned}$$

In order to illustrate the convergence result, we determine the parameter λ^* for $u^0(x, U)$. For comparison purposes, we compute $u^n(x, U)$ numerically for particular $n \in \mathbb{N}$ by backward recursion. Figure III.2 shows the value functions for some approximations as well as the value function for the Black–Scholes model.

III.4 Stability of the demand problem for RDEU

In this section, we analyze the stability for the *rank-dependent expected utility* for which the functional V_n is defined in (2.3) by

$$V_n(f) := \int_0^\infty T(P^n[U(f) > x])dx.$$

The goal is to prove Theorem 2.9 for this case, that is, to prove that

$$\lim_{n \rightarrow \infty} v^n(x) = v^0(x) \tag{4.1}$$

and to show that any given sequence of maximizers f^n for $v^n(x)$ contains a subsequence that converges weakly to a maximizer for $v^0(x)$.

III.4.1 Weak convergence of maximizers

We start with a convergence result for the maximizers. For later purposes, we prove a slightly more general statement; in particular, our proof needs no assumption on the distribution of φ^0 so that we can use Proposition 4.1 also in Section III.5.

Proposition 4.1. *For every sequence (f^n) with $f^n \in C^n(x)$, there are a subsequence (n_k) and some $\bar{f} \in C^0(x)$ such that $f^{n_k} \Rightarrow \bar{f}$ as $k \rightarrow \infty$.*

Let us outline the main ideas of the proof. We first use Helly's selection principle to get a limit distribution \bar{F} . In order to find a final position with distribution \bar{F} , we then follow the path of Jin and Zhou [64], He and Zhou [56], Carlier and Dana [31] and define the candidate payoff in the limit model as a quantile function $q_{\bar{F}}$ applied to a uniformly distributed random variable. This ensures that the distribution of this final position is \bar{F} . In order to find the cheapest final position with the given distribution \bar{F} , one has to choose the "right" uniformly distributed random variable. If the pricing density φ^0 has a continuous distribution (as assumed in [64, 56, 31] mentioned above), then $1 - F_{\varphi^0}(\varphi^0)$ turns out to be the good choice. In the general case where the distribution of φ^0 is not necessarily continuous, one can work with a uniformly distributed random variable \mathcal{U} satisfying $\varphi^0 = q_{\varphi^0}(\mathcal{U})$ P^0 -a.s. and then proceed similarly as in the first case. We also make use of the Hardy–Littlewood inequality, which states that any two random variables $f, g \in L_+^0(\Omega, \mathcal{F}, P)$ satisfy

$$E[fg] \geq \int_0^1 q_f(s)q_g(1-s)ds, \tag{4.2}$$

see Theorem A.24 of Föllmer and Schied [43] for a proof.

In the proof of Proposition 4.1, we use the following tightness result to apply Helly's selection principle. Its proof is given at the end of this subsection.

Lemma 4.2. *Let F^n be the distribution of f^n . Then $(F^n)_{n \in \mathbb{N}}$ is tight, i.e.,*

$$\lim_{c \rightarrow \infty} \sup_{n \in \mathbb{N}} P^n[f^n > c] = 0.$$

We are now in a position to prove Proposition 4.1.

Proof of Proposition 4.1. Let F^n be the distribution function of f^n . Since the sequence (F^n) is tight (Lemma 4.2), we may apply Helly's selection theorem (Billingsley [18, Theorem 6.1 and p.227]) to get a subsequence (n_k) and a distribution function \bar{F} such that $\lim_{k \rightarrow \infty} F^{n_k}(a) = \bar{F}(a)$ holds for all continuity points a of \bar{F} .

Since $(\Omega^0, \mathcal{F}^0, P^0)$ is atomless, it is possible to find on $(\Omega^0, \mathcal{F}^0, P^0)$ a random variable \mathcal{U} uniformly distributed on $(0, 1)$ such that $\varphi^0 = q_{\varphi^0}(\mathcal{U})$ P^0 -a.s. (Lemma A.28 in Föllmer and Schied [43]). Define $\bar{f} := q_{\bar{F}}(1 - \mathcal{U})$. Since $1 - \mathcal{U}$ is again uniformly distributed on $(0, 1)$, the candidate \bar{f} has distribution \bar{F} (Lemma A.19 in [43]). This gives $f^{n_k} \Rightarrow \bar{f}$ as $k \rightarrow \infty$.

The proof is completed by showing that $\bar{f} \in C^0(x)$, as follows. We rewrite φ^0 and \bar{f} in terms of \mathcal{U} , and combine Fatou's lemma and the fact that weak convergence implies convergence of any quantile functions to get a first inequality. A second one follows by applying the Hardy–Littlewood inequality (4.2). Finally, we make use of $f^{n_k} \in C^{m_k}(x)$. These steps together give

$$\begin{aligned} E_0[\varphi^0 \bar{f}] &= E_0[q_{\varphi^0}(\mathcal{U})q_{\bar{F}}(1 - \mathcal{U})] = \int_0^1 q_{\varphi^0}(s)q_{\bar{F}}(1 - s)ds \\ &\leq \liminf_{k \rightarrow \infty} \int_0^1 q_{\varphi^{n_k}}(s)q_{F^{n_k}}(1 - s)ds \leq \liminf_{k \rightarrow \infty} E_{n_k}[\varphi^{n_k} f^{n_k}] \leq x, \end{aligned}$$

which proves that $\bar{f} \in C^0(x)$. □

It remains to give the

Proof of Lemma 4.2. We show below that

$$\alpha := \lim_{c \rightarrow \infty} \limsup_{n \rightarrow \infty} P^n[f^n > c] = 0. \quad (4.3)$$

This allows us for every $\epsilon > 0$ to choose c_0 and $n_0(c_0)$ in such a way that $P^n[f^n > c_0] < \epsilon$ for $n > n_0(c_0)$. For any $c \geq c_0$ and any $n \geq n_0(c_0)$, we then obtain $0 \leq P^n[f^n > c] \leq P^n[f^n > c_0] \leq \epsilon$ and therefore

$$0 \leq \sup_{n \geq n_0(c_0)} P^n[f^n > c] \leq \epsilon \quad \text{for } c \geq c_0.$$

By increasing c_0 to c_1 to account for the finitely many $n < n_0(c_0)$, we get

$$\sup_{n \in \mathbb{N}} P^n[f^n > c] \leq \epsilon \quad \text{for } c \geq c_1.$$

Because $\epsilon > 0$ was arbitrary, this means that the family $(F^n)_{n \in \mathbb{N}}$ is tight.

We now show (4.3). First, note that the assumption $\varphi^0 > 0$ implies $F_{\varphi^0}(0) = 0$ and, by the definition of a quantile function, that q_{φ^0} is positive and satisfies $F_{\varphi^0}(q_{\varphi^0}(\epsilon)) \geq \epsilon > 0$ for every $\epsilon > 0$. Thus, q_{φ^0} must be strictly positive on $(0, \epsilon)$ for $\epsilon > 0$ which implies that $\int_0^\epsilon q_{\varphi^0}(t)dt$ is strictly positive for $\epsilon > 0$. Assume by way of contradiction that $\alpha > 0$. For $\epsilon > 0$ small enough, choose a constant c_0 in such a way that $\limsup_n P^n[f^n > c_0] \geq \alpha - \epsilon$ and $c_0 \int_0^{\alpha-2\epsilon} q_{\varphi^0}(t)dt > x+1$. Weak convergence gives convergence of the quantile functions and so we have $q_{\varphi^n}(t) \leq q_{\varphi^0}(\alpha - 2\epsilon) + \epsilon$ on $(0, \alpha - 2\epsilon)$ for sufficiently large n , so dominated convergence gives $\int_0^{\alpha-2\epsilon} q_{\varphi^n}(t)dt \rightarrow \int_0^{\alpha-2\epsilon} q_{\varphi^0}(t)dt$. Because the limit is strictly positive, this and the choice of c_0 allow us to choose n_0 in such a way that

$$c_0 \int_0^{\alpha-2\epsilon} q_{\varphi^{n_0}}(t)dt > x \quad (4.4)$$

and $P^{n_0}[f^{n_0} > c_0] \geq \alpha - 2\epsilon$. The latter implies $P^{n_0}[f^{n_0} \leq c_0] \leq 1 - \alpha + 2\epsilon$ which can be used to control $q_{f^{n_0}}$ on $(1 - \alpha + 2\epsilon, 1)$. Indeed, the last inequality and the definition of a quantile give $F^{n_0}(c_0) < t \leq F^{n_0}(q_{f^{n_0}}(t))$ for any $t \in (1 - \alpha + 2\epsilon, 1)$ which implies that

$$q_{f^{n_0}}(t) > c_0 \quad (4.5)$$

on $(1 - \alpha + 2\epsilon, 1)$. Finally, we use the Hardy–Littlewood inequality (4.2) to rewrite $E_{n_0}[\varphi^{n_0} f^{n_0}]$ in terms of quantiles, plug in (4.5) and use (4.4) to obtain

$$E_{n_0}[\varphi^{n_0} f^{n_0}] \geq \int_0^1 q_{\varphi^{n_0}}(t)q_{f^{n_0}}(1-t)dt \geq c_0 \int_0^{\alpha-2\epsilon} q_{\varphi^{n_0}}(t)dt > x,$$

which contradicts $f^{n_0} \in C^{n_0}(x)$. \square

III.4.2 Upper-semicontinuity of $v^n(x)$

In this section, we prove the first inequality of (4.1), namely that

$$\limsup_{n \rightarrow \infty} v^n(x) \leq v^0(x). \quad (4.6)$$

Having proved weak convergence along a subsequence for any sequence (f^n) with $f^n \in C^n(x)$, the remaining step is to show that the corresponding sequence of values $V_n(f^n)$ converges as well. For this, we use the growth condition imposed on U and T as well as of the integrability condition imposed on $(\varphi^n)_{n \in \mathbb{N}}$.

Throughout this section, we assume that Assumptions 2.2 and 2.7 hold true.

Lemma 4.3. *Let $f^n \in C^n(x)$. Then the family $(T(P^n[U(f^n) > y]))_{n \in \mathbb{N}}$ is uniformly integrable.*

Proof. Since $T(P^n[U(f^n) > y])$ is nonnegative for every $n \in \mathbb{N}$, it is sufficient to find an integrable upper bound independent of n . We first apply (2.7), the Chebyshev inequality and (2.6) and then use that $|x + y|^\eta \leq c(\eta)(|x|^\eta + |y|^\eta)$ for some constant $c(\eta)$ (Lemma 7.4) to obtain

$$\begin{aligned}
 T(P^n[U(f^n) > y]) &\leq k_3(P^n[U(f^n) > y])^\alpha \\
 &\leq k_3 \frac{E_n[U(f^n)^\lambda]^\alpha}{y^{\lambda\alpha}} \\
 &\leq \frac{k_3}{y^{\lambda\alpha}} E_n \left[(k_1(f^n)^\gamma + k_2)^\lambda \right]^\alpha \\
 &\leq \frac{k_3}{y^{\lambda\alpha}} E_n \left[c(\lambda)(k_1^\lambda(f^n)^{\gamma\lambda} + k_2^\lambda) \right]^\alpha \\
 &\leq \frac{k_3}{y^{\lambda\alpha}} (c(\lambda))^\alpha \left(k_1^\lambda E_n[(f^n)^{\gamma\lambda}] + k_2^\lambda \right)^\alpha,
 \end{aligned} \tag{4.7}$$

where λ is the one fixed in Assumption 2.7. In the next step, we estimate the term $E_n[(f^n)^{\gamma\lambda}]$. Recall that $\gamma\lambda < 1$ by Assumption 2.7, so the conjugate of the function $x \mapsto x^{\gamma\lambda}$ is $y \mapsto c_1 y^{\gamma\lambda/(\gamma\lambda-1)}$ for some constant c_1 . Since $f^n \in C^n(x)$, this gives

$$E_n[(f^n)^{\gamma\lambda}] \leq E_n[(f^n)^{\gamma\lambda} - f^n \varphi^n] + x \leq c_1 E_n[(\varphi^n)^{\gamma\lambda/(\gamma\lambda-1)}] + x. \tag{4.8}$$

Recall that $\varphi^n \Rightarrow \varphi^0$ by assumption, so also $(\varphi^n)^{\gamma\lambda/(\gamma\lambda-1)} \Rightarrow (\varphi^0)^{\gamma\lambda/(\gamma\lambda-1)}$. Since the family $(\varphi^n)^{\gamma\lambda/(\gamma\lambda-1)}$ is uniformly integrable by Assumption 2.2, we therefore obtain $E_n[(\varphi^n)^{\gamma\lambda/(\gamma\lambda-1)}] \rightarrow E_0[(\varphi^0)^{\gamma\lambda/(\gamma\lambda-1)}]$ as $n \rightarrow \infty$. Together with (4.8), this gives

$$E_n[(f^n)^{\gamma\lambda}] \leq c_1 E_n \left[(\varphi^n)^{\frac{\gamma\lambda}{\gamma\lambda-1}} \right] + x \leq c_1 E_0 \left[(\varphi^0)^{\frac{\gamma\lambda}{\gamma\lambda-1}} \right] + x + 1 =: k_4 \tag{4.9}$$

for sufficiently large n . Combining (4.7) and (4.9) finally yields

$$0 \leq T(P^n[U(f^n) > y]) \leq \frac{k_3}{y^{\lambda\alpha}} c(\lambda) \left(k_1^\lambda k_4 + k_2^\lambda \right)^\alpha 1_{\{y \geq 1\}} + 1_{\{y < 1\}},$$

which gives an integrable upper bound since $\lambda\alpha > 1$ by Assumption 2.7. \square

We now combine Proposition 4.1 and Lemma 4.3 to prove the upper-semicontinuity of $v^n(x)$.

Proposition 4.4. *Let (f^n) be a sequence with $f^n \in C^n(x)$ and (f^{n_k}) , \bar{f} the subsequence and its limit constructed in Proposition 4.1. Then we have*

$$\lim_{k \rightarrow \infty} V_{n_k}(f^{n_k}) = V_0(f^0). \tag{4.10}$$

Consequently, we have $\limsup_{n \rightarrow \infty} v^n(x) \leq v^0(x)$.

Proof. Starting from an arbitrary sequence (f^n) with $f^n \in C^n(x)$, Proposition 4.1 gives a subsequence (n_k) and a weak limit \bar{f} such that $f^{n_k} \Rightarrow \bar{f}$. The function U is continuous, hence $U(f^{n_k}) \Rightarrow U(f^0)$ and therefore we get $P^{n_k}[U(f^{n_k}) > y] \rightarrow P^0[U(f^0) > y]$ for all points y where $P^0[U(f^0) > y]$ is continuous. But $P^0[U(f^0) > y]$ is a decreasing function of y ; hence there are at most countably many points where $P^0[U(f^0) > y]$ is not continuous, and we deduce that $P^{n_k}[U(f^{n_k}) > y] \rightarrow P^0[U(f^0) > y]$ for a.e. y . Moreover, T is increasing; hence it is continuous a.e. and we infer that we have $T(P^{n_k}[U(f^{n_k}) > y]) \rightarrow T(P^0[U(f^0) > y])$ for a.e. y . By Lemma 4.3, the family $(T(P^n[U(f^n) > y]))_{n \in \mathbb{N}}$ is uniformly integrable and we arrive at

$$\begin{aligned} \lim_{k \rightarrow \infty} V_{n_k}(f^{n_k}) &= \lim_{k \rightarrow \infty} \int_0^\infty T(P^{n_k}[U(f^{n_k}) > y]) dy \\ &= \int_0^\infty T(P^0[U(f^0) > y]) dy \\ &= V_0(f^0). \end{aligned}$$

For the proof of upper-semicontinuity of v^n (in n), assume by way of contradiction that $\limsup_n v^n(x) > v^0(x)$. This allows us to choose a sequence (f^n) with $f^n \in C^n(x)$ and $\limsup_n V_n(f^n) > v^0(x)$. We can then pass to a subsequence realizing the lim sup and apply the first part of proof to the subsequence to get a further subsequence (f^{n_k}) and a weak limit $\bar{f} \in C^0(x)$ with

$$v^0(x) < \limsup_{n \rightarrow \infty} V_n(f^n) = \lim_{k \rightarrow \infty} V_{n_k}(f^{n_k}) = V_0(\bar{f}),$$

which gives the required contradiction. \square

Remark 4.5. Proposition 4.4 can also be used to prove the existence of a maximizer for $v^0(x)$, as follows. We formally introduce a sequence of models by setting $(\Omega^n, \mathcal{F}^n, P^n) := (\Omega^0, \mathcal{F}^0, P^0)$ for all $n \in \mathbb{N}$ and fix a maximizing sequence (f^n) for $v^0(x)$. Proposition 4.4 then shows that the limit \bar{f} constructed in Proposition 4.1 is a maximizer. As a by-product, we also see that $v^0(x) < \infty$. Note that so far we have not used the assumption that φ^0 has a continuous distribution. Jin and Zhou [64] and Carlier and Dana [31] prove the existence of a maximizer for $v^0(x)$ under the assumption that φ^0 has a continuous distribution. Proposition 4.4 (together with Proposition 4.1) shows how to extend their results to an atomless underlying model with a pricing density which is not necessarily continuous. \diamond

III.4.3 Lower-semicontinuity of $v^n(x)$

The purpose of this section is to show the second inequality “ \geq ” of (4.1). The natural idea is to approximate payoffs in the limit model by a sequence of payoffs in the approximating models. For a generic payoff, this might

be difficult; but we argue in the first step that it suffices to consider pay-offs of the form $h(\varphi^0)$ for a bounded function h . Those elements can be approximated by the sequence $(h(\varphi^n))$. Since h is bounded, the sequence $T(P[U(h(\varphi^n)) > y])$ as well as $(\varphi^n h(\varphi^n))$ have nice integrability properties, so that one obtains the desired convergence results for $(V_n(h(\varphi^n)))$ as well as for $(E_n[\varphi^n h(\varphi^n)])$.

Proposition 4.6. *Suppose that φ^0 has a continuous distribution. Then*

$$\liminf_{n \rightarrow \infty} v^n(x) \geq v^0(x).$$

Proof. 1) *Reduction to a bounded payoff:* Suppose by way of contradiction that $\liminf_n v^n(x) < v^0(x)$. In Lemma 4.7 below, we show lower-semicontinuity of the function $v^0(x)$. This allows us to choose $\epsilon > 0$ such that $\liminf_n v^n(x) < v^0(x - \epsilon)$. Therefore we can find and fix $\tilde{f} \in C^0(x - \epsilon)$ satisfying $\liminf_n v^n(x) < V_0(\tilde{f})$. Next we define an additional sequence (f^m) by $f^m := \tilde{f} \wedge m$. By construction, this sequence is increasing to \tilde{f} , and this gives $P^0[U(f^m) > y] \nearrow P^0[U(\tilde{f}) > y]$ for all y . The function T is increasing, hence continuous a.e., and we thus have $T(P^0[U(f^m) > y]) \nearrow T(P^0[U(\tilde{f}) > y])$ for a.e. y . Monotone convergence then yields $V_0(f^m) \rightarrow V_0(\tilde{f})$. This allows us to find and fix m_0 such that $\liminf_n v^n(x) < V_0(\tilde{f}^{m_0})$.

2) *Reduction to a payoff $h(\varphi^0)$:* We define $h(s) := q_{\tilde{f}^{m_0}}(1 - F_{\varphi^0}(s))$. By the definition of a quantile, $q_{\tilde{f}^{m_0}}$ is increasing, so h is decreasing. Moreover, since \tilde{f}^{m_0} is bounded by m_0 , the quantile $q_{\tilde{f}^{m_0}}$ is bounded by m_0 as well. Recall now that the distribution of φ^0 is assumed to be continuous. Thus $F_{\varphi^0}(\varphi^0)$ as well as $1 - F_{\varphi^0}(\varphi^0)$ are uniformly distributed on $(0, 1)$ and $h(\varphi^0)$ therefore has the same distribution as \tilde{f}^{m_0} . But the preference functional V_0 only depends on the distribution of its argument, and so we get

$$V_0(h(\varphi^0)) = V_0(\tilde{f}^{m_0}) > \liminf_{n \rightarrow \infty} v^n(x). \quad (4.11)$$

Finally, we use the monotonicity of h together with (4.2), $\tilde{f}^{m_0} \leq \tilde{f}$ and $\tilde{f} \in C^0(x - \epsilon)$ to obtain

$$E_0[\varphi^0 h(\varphi^0)] \leq E_0[\varphi^0 \tilde{f}^{m_0}] \leq E_0[\varphi^0 \tilde{f}] \leq x - \epsilon.$$

This gives $h(\varphi^0) \in C^0(x - \epsilon)$.

3) *Convergence of $(V_n(h(\varphi^n)))$:* Let D_h denote the set of all points where h is not continuous. The function h is decreasing and so D_h is at most countable; but φ^0 has a continuous distribution and it follows that $P^0[\varphi^0 \in D_h] = 0$. Hence we get $h(\varphi^n) \Rightarrow h(\varphi^0)$ which then implies $T(P^n[h(\varphi^n) > y]) \rightarrow T(P^0[h(\varphi^0) > y])$ for every y . By construction, h is positive and bounded by m_0 ; hence $1_{\{y \leq m_0\}}$ is an integrable upper bound for $(T(P^n[h(\varphi^n) > y]))_{n \in \mathbb{N}}$ and dominated convergence gives $V_n(h(\varphi^n)) \rightarrow V_0(h(\varphi^0)) = V_0(\tilde{f}^{m_0})$.

4) *Convergence of $(E_n[\varphi^n h(\varphi^n)])$* : We define $g(x) := xh(x)$ and note that the set of points where g is not continuous is again D_h . As in step 3), we therefore have $g(\varphi^n) \Rightarrow g(\varphi^0)$. Recall now that $E_n[\varphi^n] = 1$ and $\varphi^n \Rightarrow \varphi^0$. This gives uniform integrability of $(\varphi^n)_{n \in \mathbb{N}}$ which in turn implies uniform integrability of $(g(\varphi^n))_{n \in \mathbb{N}}$ since $0 \leq g(\varphi^n) \leq m_0 \varphi^n$. Together with $h(\varphi^0) \in C^0(x - \epsilon)$ as proved in part 3), we obtain

$$E_n[\varphi^n h(\varphi^n)] = E_n[g(\varphi^n)] \longrightarrow E_0[g(\varphi^0)] = E_0[\varphi^0 h(\varphi^0)] \leq x - \epsilon$$

as $n \rightarrow \infty$. This implies $h(\varphi^n) \in C^n(x)$ (for sufficiently large n). But on the other hand, (4.11) and part 3) give

$$\lim_{n \rightarrow \infty} V_n(h(\varphi^n)) = V_0(\tilde{f}^{m_0}) > \liminf_{n \rightarrow \infty} v^n(x),$$

which gives the required contradiction. \square

Lemma 4.7. *Suppose that φ^0 has a continuous distribution. Then*

$$\liminf_{x \nearrow x_0} v^0(x) \geq v^0(x_0).$$

Proof. We assume to the contrary that $\liminf_{x \rightarrow x_0} v^0(x) < v^0(x_0)$. Fix $f^0 \in C^0(x_0)$ with $\liminf_{x \rightarrow x_0} v^0(x) < V_0(f^0)$ and define $f^\epsilon := (1 - \epsilon)f^0$. It follows that $f^\epsilon \in C^0((1 - \epsilon)x_0)$ and $f^\epsilon \nearrow f^0$ a.s. The latter implies $P^0[U(f^\epsilon) > y] \nearrow P^0[U(f^0) > y]$ for a.e. y as $\epsilon \rightarrow 0$, and since T is increasing, we have $T(P^0[U(f^\epsilon) > y]) \nearrow T(P^0[U(f^0) > y])$ a.e. Monotone convergence then yields $V_0(f^\epsilon) \nearrow V_0(f^0)$ as $\epsilon \rightarrow 0$, which gives a contradiction. \square

Proof of Theorem 2.9 for RDEU. It follows from Propositions 4.4 and 4.6 that $\lim_n v^n(x) = v^0(x)$. For the second part, we apply Proposition 4.1 to (f^n) . This gives a subsequence (f^{n_k}) weakly converging to some $\bar{f} \in C^0(x)$, and Proposition 4.4 implies that \bar{f} is a maximizer for $v^0(x)$. \square

III.5 Stability of the demand problem for ENCU

In this section, we analyze the case that the functional V_n is defined by

$$V_n(f) := E_n[U(f)]$$

for a non-concave utility function U . Except for non-concavity, this coincides with the classical expected utility where the value function is usually denoted by u^n instead of v^n . We follow that tradition and switch to u^n from now on. Moreover, the analysis in Chapter II, particularly in Sections II.4 and II.5, shows that the optimization problem for the non-concave utility function U is closely linked to the optimization problem for its concave envelope U_c , and

both of them are useful for the analysis in this section. Therefore we use in this section the notation

$$u^n(x, U) := v^n(x)$$

for the value function. The goal is to prove Theorem 2.9 for the present case, that is, to prove

$$\lim_{n \rightarrow \infty} u^n(x, U) = u^0(x, U) \quad (5.1)$$

and to show that a given sequence of maximizers f^n for $u^n(x, U)$ contains a subsequence that converges weakly to a maximizer for $u^0(x, U)$. In contrast to Section III.4, we prove the stability results here without further assumptions on the distribution of φ^0 . This is necessary for the theoretical applications described in Section III.3.3. A model of an atomless limit model with a unique (nontrivial) pricing density which does not have a continuous distribution can be found in Example II.5.10.

Throughout this section, we assume that $\mathbf{AE}_0(J) < \infty$ and that Assumption 2.2 is satisfied.

III.5.1 Upper-semicontinuity of $u^n(x, U)$

The main idea is similar to the proof of upper-semicontinuity of $v^n(x)$ in Section III.4.2. Starting from a sequence (f^n) with $f^n \in C^n(x)$, we use the results of Section III.4.1 to obtain weak convergence along a subsequence (n_k) to an element $\bar{f} \in C^0(x)$. Using Fatou's lemma for $(E_{n_k}[U(f^{n_k})^-])$, the remaining step is then to show that the corresponding sequence $(E_{n_k}[U(f^{n_k})^+])$ converges as well. This requires the uniform integrability of the family $(U(f^n)^+)_{n \in \mathbb{N}}$ which can be proved with the help of the following lemma.

Lemma 5.1. *Assumption 2.2 and $\mathbf{AE}_0(J) < \infty$ imply the uniform integrability of the family $(J(\lambda\varphi^n)^+)_{n \in \mathbb{N}}$ for any $\lambda > 0$.*

Proof. We show below that Assumption 2.2 and $\mathbf{AE}_0(J) < \infty$ ensure that

$$\text{the family } (J(\varphi^n)^+)_{n \in \mathbb{N}} \text{ is uniformly integrable.} \quad (5.2)$$

The statement is then clear for $\lambda \geq 1$ since J is decreasing. For $\lambda < 1$, the assumption that $\mathbf{AE}_0(J) < \infty$ (in combination with Lemma 7.3) can be used to obtain constants $\gamma > 0$ and $y_0 > 0$ such that $J(\mu y) \leq \mu^{-\gamma} J(y)$ for all $\mu \in (0, 1]$ and $y \in (0, y_0]$. Applying this estimate on the set $\{\varphi^n \leq y_0\}$ and using monotonicity of J on the complement gives

$$J(\lambda\varphi^n)^+ \leq \lambda^{-\gamma} J(\varphi^n)^+ 1_{\{\varphi^n \leq y_0\}} + J(\lambda y_0)^+ 1_{\{\varphi^n > y_0\}} \leq \lambda^{-\gamma} J(\varphi^n)^+ + J(\lambda y_0)^+.$$

The second term in the last line is constant, and uniform integrability of the first one is due to (5.2).

It remains to prove (5.2). The second part of Lemma 7.3 in the appendix shows that $AE_0(J) < \infty$ implies the existence of constants k_1, k_2 and $\gamma < 1$ such that $U(x) \leq k_1 + k_2x^\gamma$ for $x \geq 0$. Plugging this inequality into the definition of J and doing some elementary computations gives

$$J(y) \leq k_1 + \sup_{x \geq 0} (k_2x^\gamma - xy) = k_1 + Cy^{\gamma/(\gamma-1)}$$

for some constant C . But then it follows that $0 \leq J(\varphi^n)^+ \leq k_1 + C(\varphi^n)^{\gamma/(\gamma-1)}$, and Assumption 2.2 yields the uniform integrability of $(J(\varphi^n)^+)_{n \in \mathbb{N}}$. \square

We now describe, as outlined above, the limit behaviour of $(E_n[U(f^n)])_n$. This gives upper-semicontinuity in \bar{n} for $u^n(x, U)$ and it can also be used (later) to deduce the optimality of \bar{f} constructed in Proposition 4.1.

Proposition 5.2. *The sequence (f^n) of maximizers for $u^n(x, U)$ contains a subsequence (f^{n_k}) weakly converging to some limit $\bar{f} \in C^0(x)$, and it satisfies $\limsup_{n \rightarrow \infty} E_n[U(f^n)] \leq E_0[U(\bar{f})]$. Consequently, we have*

$$\limsup_{n \rightarrow \infty} u^n(x, U) \leq u^0(x, U).$$

Proof. We consider a (relabelled) subsequence (f^n) realizing the \limsup γ , say, for $E_n[U(f^n)]$ (or equivalently for $u^n(x, U)$, since the f^n are maximizers). Proposition 4.1 gives a further subsequence (f^{n_k}) with weak limit \bar{f} . Since $U, \max(\cdot, 0)$ and $\min(\cdot, 0)$ are continuous, we infer that $U(f^{n_k})^\pm \Rightarrow U(\bar{f})^\pm$. Fatou's lemma then gives

$$E_0[U(\bar{f})^-] \leq \liminf_{k \rightarrow \infty} E_{n_k}[U(f^{n_k})^-].$$

We show below that $(U(f^n)^+)_{n \in \mathbb{N}}$ is uniformly integrable, which implies that $(E_{n_k}[U(f^{n_k})^+])$ converges to $E_0[U(\bar{f})^+]$ as $k \rightarrow \infty$. Combining this with the inequality for the negative parts yields

$$\begin{aligned} E_0[U(\bar{f})] &\geq \lim_{k \rightarrow \infty} E_{n_k}[U(f^{n_k})^+] - \liminf_{k \rightarrow \infty} E_{n_k}[U(f^{n_k})^-] \\ &\geq \limsup_{k \rightarrow \infty} E_{n_k}[U(f^{n_k})] \\ &= \limsup_{n \rightarrow \infty} u^n(x, U), \end{aligned}$$

where we use in the last step that the (f^{n_k}) form a subsequence of the sequence (f^n) for which we have $\lim_n E_n[U(f^n)] = \gamma$ from above.

It remains to show uniform integrability of $(U(f^n)^+)_{n \in \mathbb{N}}$. This family is, by the definition of J , dominated by $(J(\epsilon\varphi^n)^+ + \epsilon\varphi^n f^n)_{n \in \mathbb{N}}$. Uniform integrability of the first summand family follows from Lemma 5.1, and since $(\varphi^n f^n)_{n \in \mathbb{N}}$ is bounded in L^1 , the sequence $(\epsilon\varphi^n f^n)_{n \in \mathbb{N}}$ can be made arbitrarily small in expectation by choosing ϵ small. So uniform integrability of $(U(f^n)^+)_{n \in \mathbb{N}}$ follows and the proof is complete. \square

III.5.2 Lower-semicontinuity of $u^n(x, U)$

The goal of this section is to show lower-semicontinuity in n for $u^n(x, U)$.

Theorem 5.3. *Suppose that Assumption 2.1 holds true. Then*

$$\liminf_{n \rightarrow \infty} u^n(x, U) \geq u^0(x, U).$$

The approach to prove this statement is as follows. Observe that

$$\liminf_{n \rightarrow \infty} u^n(x, U) \geq \liminf_{n \rightarrow \infty} (u^n(x, U) - u^n(x, U_c)) + \liminf_{n \rightarrow \infty} u^n(x, U_c).$$

If one shows (as we do below in Section III.5.2.1), that

$$\liminf_{n \rightarrow \infty} u^n(x, U_c) = \lim_{n \rightarrow \infty} u^n(x, U_c) = u^0(x, U_c), \quad (5.3)$$

it only remains to show that

$$u^n(x, U_c) - u^n(x, U) \rightarrow 0 \quad \text{as } n \rightarrow \infty. \quad (5.4)$$

While the proof of (5.3) follows (essentially) from non-smooth versions of known stability results on *concave* utility maximization, the proof of (5.4) requires a careful analysis of the *non-concave* problem which will be explained in detail in Section III.5.2.2. Note that the additional Assumption 2.1 is only used to prove (5.4). We start with the proof of (5.3).

III.5.2.1 Continuity in n of $u^n(x, U_c)$

Instead of lower-semicontinuity, we prove slightly more than needed, namely

Proposition 5.4. $\lim_{n \rightarrow \infty} u^n(x, U_c) = u^0(x, U_c)$.

In the case of strictly concave utility functions, this result follows by directly analyzing the sequence of optimal terminal wealths f^n as a function of φ^n . In the non-concave framework, U_c is not strictly concave; hence its conjugate J is non-smooth and f^n cannot be written as a function of φ^n (f^n only lies in the subgradient of $-J$ at φ^n). Instead, we use the fact that $u^n(x, U_c)$ can be written (see Lemma 5.7 below) in a dual form as

$$u^n(x, U_c) = \inf_{\lambda > 0} E_n[J(\lambda \varphi^n) + x\lambda] = E_n[J((\lambda(n)\varphi^n) + x\lambda(n))]$$

for some dual minimizer $\lambda(n) \geq 0$. Continuity in n of $u^n(x, U_c)$ can then be shown by proving that the sequence $(\lambda(n))$ converges (along a subsequence) to a dual minimizer in the limit model and that the sequence $(E_n[J(\lambda(n)\varphi^n)])$ converges to the corresponding value in the limit model. The latter requires uniform integrability of the family $(J(\lambda(n)\varphi^n))_{n \in \mathbb{N}}$. For the positive parts, this can be proved via Lemma 5.1. We now show that the family of negative parts is uniformly integrable as well.

Lemma 5.5. *For each $s > 0$, the family $\{J(\lambda\varphi^n)^- \mid n \in \mathbb{N}, \lambda \in [0, s]\}$ is uniformly integrable.*

Proof. The idea for this result goes back to Kramkov and Schachermayer [73]; the extension to the non-smooth case is proved in Lemma 6.1 of Bouchard et al. [22]. A modified version of their proof works for our setup, as follows.

Since the conjugate J is decreasing, it is enough to check uniform integrability of $(J(s\varphi^n)^-)_n$. If $J(\infty) > -\infty$, all the $J(s\varphi^n)^-$ are bounded by a uniform constant and the statement is clear. So assume $J(\infty) = -\infty$. To use the de la Vallée–Poussin characterization of uniform integrability, we need to find a convex increasing function $\Phi : [0, \infty) \rightarrow \mathbb{R}$ such that $\lim_{x \rightarrow \infty} \frac{\Phi(x)}{x} = \infty$ and $\sup_n E_n[\Phi(J(s\varphi^n)^-)] < \infty$. The function J is convex, decreasing and finite on $(0, \infty)$ (see Lemma II.2.12); so for $J(\infty) = -\infty$, J is strictly decreasing and J as well as $-J$ have a classical inverse. Let $\Phi : (-J(0), +\infty) \rightarrow (0, \infty)$ be the inverse of $-J$. Since $-J$ is increasing and concave, its inverse Φ is increasing and convex. In order to prove that

$$\infty = \lim_{x \rightarrow \infty} \frac{\Phi(x)}{x} = \lim_{y \rightarrow \infty} \frac{y}{-J(y)}, \quad (5.5)$$

note first that $\lim_{y \rightarrow \infty} \sup_{q \in -\partial J(y)} q = 0$ (see Lemma II.7.3) implies

$$\liminf_{y \rightarrow \infty} \inf_{q \in -\partial J(y)} \frac{1}{q} = \liminf_{y \rightarrow \infty} \left(\sup_{q \in -\partial J(y)} q \right)^{-1} = \infty.$$

Hence for all M , there is y_0 such that $\inf_{q \in -\partial J(y)} \frac{1}{q} \geq M$ for all $y \geq y_0$. Fix some y_1 and y_2 satisfying $y_0 < y_1 < y_2$ and set $z := (J(y_2) - J(y_1))/(y_2 - y_1)$. The mean value theorem gives $\tau \in [y_1, y_2]$ such that $z \in \partial J(\tau)$. This implies by the definition of the subdifferential ∂J that

$$M \leq \inf_{q \in -\partial J(y_0)} \frac{1}{q} \leq \frac{y_2 - y_1}{-J(y_2) + J(y_1)} \leq \frac{\frac{y_2}{J(y_2)} - \frac{y_1}{J(y_2)}}{-1 + \frac{J(y_1)}{J(y_2)}}.$$

Taking the \liminf as $y_2 \rightarrow \infty$ gives $M \leq \liminf_{y_2 \rightarrow \infty} -y_2/J(y_2)$. The proof of (5.5) is complete since the constant M is arbitrary.

It remains to prove that $\sup_n E_n[\Phi(J(s\varphi^n)^-)] < \infty$. Recall that J is convex and finite on $(0, \infty)$ and hence continuous, and that $J(0) = U(\infty) > 0$ by the assumption on U . Moreover, $J(\infty) = -\infty$ in the present case, so there is $a \in (0, \infty)$ with $J(a) = 0$ and this implies $\Phi(0) = a < \infty$. By a direct computation, we see that for $s > 0$,

$$\begin{aligned} E_n[\Phi(J(s\varphi^n)^-)] &= E_n[\Phi(\max\{0, -J(s\varphi^n)\})] \\ &\leq E_n[\max\{\Phi(0), s\varphi^n\}] \\ &\leq \Phi(0) + E_n[s\varphi^n] = \Phi(0) + s \end{aligned}$$

which completes the proof. \square

We now show that weak convergence of $\lambda_n \varphi^n$ to $\lambda \varphi^0$ indeed implies convergence of $E_n[J(\lambda_n \varphi^n)]$ to $E_0[J(\lambda \varphi^0)]$.

Lemma 5.6. *Let $\lambda_n \rightarrow \lambda \in (0, \infty)$ be given. Then it holds that*

$$E_n[J(\lambda_n \varphi^n)] \longrightarrow E_0[J(\lambda \varphi^0)] \quad \text{as } n \rightarrow \infty.$$

Proof. The continuity of J together with $\lambda_n \rightarrow \lambda \in (0, \infty)$ and $\varphi^n \Rightarrow \varphi^0$ imply $J(\lambda_n \varphi^n) \Rightarrow J(\lambda \varphi^0)$ as $n \rightarrow \infty$. Since the limit λ is in $(0, \infty)$, the λ_n lie eventually in a compact set B of the form $[\epsilon, \frac{1}{\epsilon}]$ with $0 < \epsilon < 1$, and so it is enough to show the uniform integrability of $\{J(\mu \varphi^n) \mid n \in \mathbb{N}, \mu \in B\}$. For the negative parts $\{J(\mu \varphi^n)^- \mid n \in \mathbb{N}, \mu \in B\}$, this is a consequence of Lemma 5.5, and for the positive parts, it follows by Lemma 5.1. \square

For the n -th model, the classical dual representation of $u^n(x, U_c)$ for our setting with a fixed pricing density gives a dual minimizer $\lambda(n)$. The sequence $(\lambda(n))$ does not necessarily converge; however, every cluster point yields a dual minimizer in the limit model.

Lemma 5.7. *Given any $n \in \mathbb{N}_0$, the problem $u^n(x, U_c)$ admits a maximizer $f^n \in -\partial J(\lambda(n) \varphi^n)$, where $\lambda(n) \in (0, \infty)$ is a minimizer for*

$$\inf_{\lambda > 0} E_n[J(\lambda \varphi^n) + x\lambda]. \quad (5.6)$$

Any cluster point $\bar{\lambda}$ of the sequence $(\lambda(n))$ is a minimizer for the problem $\inf_{\lambda > 0} E_0[J(\lambda \varphi^0) + x\lambda]$ and satisfies $\bar{\lambda} \in (0, \infty)$.

Proof. Lemmas 5.5 and 5.1 give $E_n[J(\lambda \varphi^n)] < \infty$ for all $n \in \mathbb{N}_0$ and all $\lambda > 0$. The existence and the structure of the solution for $u^n(x, U_c)$ and the dual representation then follow by Proposition 7.1.

For the second part, we use the notation

$$H^n(\lambda) := E_n[J(\lambda \varphi^n)] + \lambda x$$

for $n \in \mathbb{N}_0$. Convexity of J implies convexity of H^n . Fix a minimizer $\lambda(0)$ for $\inf_{\lambda > 0} H^0(\lambda)$ and a cluster point $\bar{\lambda}$ of $(\lambda(n))$. We show below that any values between $\bar{\lambda}$ and $\lambda(0)$ are minimizers for $\inf_{\lambda > 0} H^0(\lambda)$. Since by Proposition 7.1, the minimizers for $\inf_{\lambda > 0} H^0(\lambda)$ are bounded away from 0 and ∞ , we therefore must have $\bar{\lambda} \in (0, \infty)$, and continuity of H^0 then implies that $\bar{\lambda}$ is also a minimizer.

We now argue that $\lambda(n_k) \rightarrow \bar{\lambda}$ implies

$$H^0(\lambda) = H^0(\lambda(0)) \quad \text{for any } \lambda \in (\lambda(0) \wedge \bar{\lambda}, \lambda(0) \vee \bar{\lambda}).$$

By way of contradiction, we assume that $H^0(\lambda) > H^0(\lambda(0)) + 2\epsilon$ holds for some $\lambda \in (\lambda(0) \wedge \bar{\lambda}, \lambda(0) \vee \bar{\lambda})$. Lemma 5.6 with $\lambda_n \equiv \lambda(0)$ implies that

$H^{n_k}(\lambda(0)) \rightarrow H^0(\lambda(0))$ as $k \rightarrow \infty$. Thus, for ϵ small enough, there is a constant k_0 such that

$$H^{n_k}(\lambda(0)) \leq H^0(\lambda(0)) + \epsilon < H^0(\lambda) - \epsilon \leq H^{n_k}(\lambda)$$

for all $k > k_0$. From the definition of the minimizer $\lambda(n_k)$, it holds that $H^{n_k}(\lambda(n_k)) \leq H^{n_k}(\lambda(0))$. Putting the two inequalities together gives

$$H^{n_k}(\lambda(n_k)) \leq H^{n_k}(\lambda(0)) < H^{n_k}(\lambda) \quad (5.7)$$

for $k > k_0$. Since $\lambda(n_k) \rightarrow \bar{\lambda}$, the number λ is between $\lambda(n_k)$ and $\lambda(0)$ for large enough values of k . Thus, (5.7) contradicts the convexity of H^{n_k} . \square

We finally have all the ingredients to prove the convergence of $u^n(x, U_c)$.

Proof of Proposition 5.4. To obtain $\limsup_n u^n(x, U_c) \leq u^0(x, U_c)$, we apply Proposition 5.2 to U_c . For the other inequality, fix a relabelled sequence of maximizers (f^n) with $\gamma := \liminf_m u^m(x, U_c) = \lim_n E_n[U_c(f^n)]$. We use Lemma 5.7 to fix for each $n \in \mathbb{N}$ a corresponding dual minimizer $\lambda(n) \in (0, \infty)$ for (5.6). By classical duality theory and Lemma 5.7, any cluster point $\bar{\lambda}$ of $(\lambda(n))$ satisfies

$$u^0(x, U_c) = \inf_{\lambda > 0} E_0[J(\lambda\varphi^0) + x\lambda] = E_0[J(\bar{\lambda}\varphi^0)] + \bar{\lambda}x$$

and $\bar{\lambda} \in (0, \infty)$. Fix one cluster point $\bar{\lambda}$ and a converging subsequence $\lambda(n_k) \rightarrow \bar{\lambda}$. It follows from Lemma 5.6 that $E_{n_k}[J(\lambda(n_k)\varphi^{n_k})] \rightarrow E_0[J(\bar{\lambda}\varphi^0)]$ and we conclude again from the dual representation for $u^{n_k}(x, U_c)$ that $E_{n_k}[U_c(f^{n_k})] = u^{n_k}(x, U_c) \rightarrow u^0(x, U_c)$. But the full sequence $(E_n[U_c(f^n)])$ converges to γ ; so we obtain $u^0(x, U_c) = \gamma = \liminf_m u^m(x, U_c)$. This completes the proof. \square

III.5.2.2 Controlling the difference $u^n(x, U_c) - u^n(x, U)$

Let us now turn to (5.4) and prove that $u^n(x, U_c) - u^n(x, U) \rightarrow 0$. The idea here is as follows. In general, $u^n(x, U)$ is smaller than $u^n(x, U_c)$ since U_c dominates U . For some initial values x , however, the maximizer for $u^n(x, U_c)$ does not have probability mass in $\{U < U_c\}$, i.e. $P[f^* \in \{U < U_c\}] = 0$, and thus also maximizes $u^n(x, U)$. Consequently, the values $u^n(x, U_c)$ and $u^n(x, U)$ coincide for such “good” initial values, and the key is to analyze the complement of these x more carefully. For the n -th model, the “good” initial values induce a (n -dependent) partition of $(0, \infty)$ and its (n -dependent) mesh size, the maximal distance between two successive partition points, goes to 0 as $n \rightarrow \infty$ thanks to Assumption 2.1. The next result formalizes this idea.

Proposition 5.8. *Let Assumption 2.1 be satisfied and let $x_0 > 0$ and $\delta > 0$ be fixed. For every $n \in \mathbb{N}_0$, there is a set $\mathcal{B}^n \subseteq (0, \infty)$ such that*

i) $u^n(x, U) = u^n(x, U_c)$ for $x \in \mathcal{B}^n$ and

ii) there is n_0 such that $\mathcal{B}^n \cap [x_0 - \delta, x_0]$ is non-empty for $n \geq n_0$.

As a consequence, we have $\lim_{n \rightarrow \infty} (u^n(x, U_c) - u^n(x, U)) = 0$.

Let us first outline the two main ideas. The problem $u^n(x, U_c)$ admits (under our conditions) a maximizer $f^n \in -\partial J(\lambda(n)\varphi^n)$ for some $\lambda(n)$. The right- and left-hand derivatives $-J'_\pm$ satisfy $-J'_\pm \notin \{U < U_c\}$ (see Lemma II.7.2). So in order to have no probability mass in the area $\{U < U_c\}$, it is sufficient if the maximizer value $f^n(\omega)$ is equal to $-J'_-(\lambda(n)\varphi^n(\omega))$ or $-J'_+(\lambda(n)\varphi^n(\omega))$. Therefore, the initial values given by

$$E_{Q^n} [-J'_+(\lambda(n)\varphi^n)1_{D^c} - J'_-(\lambda(n)\varphi^n)1_D] \quad (5.8)$$

for $D \in \mathcal{F}^n$ are good candidates for initial values satisfying property i).

In order to also have property ii), we need to control the distance between any two points defined by (5.8). This boils down to controlling terms of the form $-J'_-(y) + J'_+(y)$. These are non-zero if y is the slope of an affine part of U_c . The distance between the points defined by (5.8) is therefore dominated by the product of the length of the longest affine part and the Q^n -probability of the biggest atom in \mathcal{F}^n . In the case of a single affine part in U_c , this goes to 0 by Assumption 2.1. In general, there is no upper bound for the length of the affine parts, but we can estimate the tails with Lemma 5.9 below. Recall that \mathcal{G}^n is the set of Q^n -atoms in \mathcal{F}^n and that Assumption 2.1 ensures that the maximal Q^n -probability of all elements in \mathcal{G}^n goes to 0.

Proof of Proposition 5.8. In order to define the set \mathcal{B}^n for Proposition 5.8, we start with some preliminary definitions and remarks. For all $n \in \mathbb{N}$, fix a maximizer $f^n_{x_0}$ for $u^n(x_0, U_c)$ and the corresponding minimizer $\lambda(n) \in (0, \infty)$ given in Lemma 5.7. This lemma also yields $\liminf_{n \rightarrow \infty} \lambda(n) > 0$. So fix $\epsilon > 0$ such that $\lambda(n) \geq \epsilon > 0$ for all n . Using Lemma 5.9 below, we obtain

$$\begin{aligned} 0 &\leq \lim_{\alpha \rightarrow \infty} \sup_{n \in \mathbb{N}} E_n [\varphi^n (-J'_-(\lambda(n)\varphi^n)) 1_{\{-J'_-(\lambda(n)\varphi^n) \geq \alpha\}}] \\ &= \lim_{\alpha \rightarrow \infty} \sup_{n \in \mathbb{N}} \frac{1}{\lambda(n)} E_n [\varphi^n \lambda(n) (-J'_-(\lambda(n)\varphi^n)) 1_{\{-J'_-(\lambda(n)\varphi^n) \geq \alpha\}}] \\ &\leq \frac{1}{\epsilon} \lim_{\alpha \rightarrow \infty} \sup_{n \in \mathbb{N}} E_n [\varphi^n \lambda(n) (-J'_-(\lambda(n)\varphi^n)) 1_{\{-J'_-(\lambda(n)\varphi^n) \geq \alpha\}}] = 0. \end{aligned}$$

Hence we may choose α_0 such that

$$\sup_{n \in \mathbb{N}} E_n [\varphi^n (-J'_-(\lambda(n)\varphi^n)) 1_{\{-J'_-(\lambda(n)\varphi^n) > \alpha_0\}}] < \delta. \quad (5.9)$$

Define the set

$$\mathcal{D}^n := \{\omega \in \Omega^n \mid -J'_+(\lambda(n)\varphi^n(\omega)) < -J'_-(\lambda(n)\varphi^n(\omega)) \leq \alpha_0\} \in \mathcal{F}^n.$$

Now we are in a position to define the set \mathcal{B}^n by

$$\mathcal{B}^n := \{E_{Q^n}[-J'_+(\lambda(n)\varphi^n)1_{D^c} - J'_-(\lambda(n)\varphi^n)1_D] \mid D \in \mathcal{F}^n, D \subset \mathcal{D}^n\}.$$

We claim that this \mathcal{B}^n satisfies the assumptions of Proposition 5.8.

1) *Property i)*: For any $n \in \mathbb{N}$ and $x \in \mathcal{B}^n$, there is some $D \in \mathcal{F}^n$ such that

$$g^n := -J'_+(\lambda(n)\varphi^n)1_{D^c} - J'_-(\lambda(n)\varphi^n)1_D \in C^n(x).$$

Note that $g^n \in -\partial J(\lambda(n)\varphi^n)$ by definition and fix some $f \in C^n(x)$. Applying the definition of J together with $E_n[\varphi^n f] \leq x$ gives

$$E_n[U_c(f)] \leq E_n[J(\lambda(n)\varphi^n)] + \lambda(n)x = E_n[U_c(g^n)],$$

where the equality follows from the classical duality relation between U_c and J . Taking the sup over all $f \in C^n(x)$ gives optimality of g^n for $u^n(x, U_c)$. Since J'_\pm do not take values in $\{U < U_c\}$ (see Lemma II.7.2), g^n satisfies $P^n[g^n \in \{U < U_c\}] = 0$ and it follows that

$$u^n(x, U_c) = E_n[U_c(g^n)] = E_n[U(g^n)] = u^n(x, U)$$

because $E_n[U(g^n)] \leq u^n(x, U) \leq u^n(x, U_c)$.

2) *Property ii)*: For this part, we use Assumption 2.1 to choose n_0 large enough such that $\sup_{A \in \mathcal{G}^n} Q^n[A] \leq \delta/\alpha_0$ for $n \geq n_0$. Fix some $n \geq n_0$ and define the map $x : \mathcal{F}^n \rightarrow \mathbb{R}_+$ by

$$x(D) := E_{Q^n}[-J'_+(\lambda(n)\varphi^n)1_{D^c} - J'_-(\lambda(n)\varphi^n)1_D].$$

Monotonicity of ∂J (see Lemma II.7.1) implies $x(\emptyset) \leq x(\mathcal{D}^n) \leq x(\Omega^n)$. Moreover, recall that $f_{x_0}^n$ and $\lambda(n)$ are fixed in such a way that

$$f_{x_0}^n \in -\partial J(\lambda(n)\varphi^n) = [-J'_+(\lambda(n)\varphi^n), -J'_-(\lambda(n)\varphi^n)]$$

satisfies $E_{Q^n}[f_{x_0}^n] = x_0$. This gives $x(\emptyset) \leq x_0 \leq x(\Omega^n)$.

We first consider the case $x_0 < x(\mathcal{D}^n)$. In order to construct a grid contained in $\mathcal{B}^n \cap [x(\emptyset), x(\mathcal{D}^n)]$, we decompose \mathcal{D}^n into disjoint subsets D_1, \dots, D_m such that $Q^n[D_i] \leq \delta/\alpha_0$ and $\bigcup_{i=1}^m D_i = \mathcal{D}^n$; this uses that for $n \geq n_0$, the largest atom in \mathcal{F}^n has Q^n -probability at most δ/α_0 . The values $x(\bigcup_{i=1}^k D_i)$, $k = 1, \dots, m$, are contained in \mathcal{B}^n , and since $-J'_-(\lambda(n)\varphi^n) \leq \alpha_0$ on $D_k \subset \mathcal{D}^n$ and $J'_+ \leq 0$, these values satisfy

$$\begin{aligned} x\left(\bigcup_{i=1}^k D_i\right) - x\left(\bigcup_{i=1}^{k-1} D_i\right) &= E_{Q^n}[(J'_+(\lambda(n)\varphi^n) - J'_-(\lambda(n)\varphi^n))1_{D_k}] \\ &\leq \alpha_0 Q^n[D_k] \leq \alpha_0 \delta/\alpha_0 = \delta \end{aligned}$$

for $k = 1, \dots, m$. We deduce that $x(\emptyset)$ and $x(\bigcup_{i=1}^k D_i)$, $k = 1, \dots, m$, form a grid with starting point $x(\emptyset)$ and endpoint $x(\bigcup_{i=1}^m D_i) = x(\mathcal{D}^n)$ whose mesh size is smaller than δ .

It remains to consider the case $x_0 \in [x(\mathcal{D}^n), x(\Omega^n)]$. Since $x(\mathcal{D}^n) \in \mathcal{B}^n$, it is sufficient to show $x(\Omega^n) - x(\mathcal{D}^n) \leq \delta$. Observe first that

$$(\mathcal{D}^n)^c = \{-J'_+(\lambda(n)\varphi^n) = -J'_-(\lambda(n)\varphi^n)\} \cup \{-J'_-(\lambda(n)\varphi^n) > \alpha_0\}.$$

We rewrite $x(\Omega^n) - x(\mathcal{D}^n)$ in terms of J'_\pm and $(\mathcal{D}^n)^c$ and use $0 \leq -J'_+ \leq -J'_-$ to obtain

$$\begin{aligned} & x(\Omega^n) - x(\mathcal{D}^n) \\ &= E_{Q^n} [(-J'_-(\lambda(n)\varphi^n) + J'_+(\lambda(n)\varphi^n)) 1_{(\mathcal{D}^n)^c}] \\ &\leq E_{Q^n} [(-J'_-(\lambda(n)\varphi^n) + J'_+(\lambda(n)\varphi^n)) 1_{\{-J'_+(\lambda(n)\varphi^n) = -J'_-(\lambda(n)\varphi^n)\}}] \\ &\quad + E_{Q^n} [(-J'_-(\lambda(n)\varphi^n) + J'_+(\lambda(n)\varphi^n)) 1_{\{-J'_-(\lambda(n)\varphi^n) > \alpha_0\}}] \\ &\leq E_n [\varphi^n (-J'_-(\lambda(n)\varphi^n)) 1_{\{-J'_-(\lambda(n)\varphi^n) > \alpha_0\}}] < \delta, \end{aligned}$$

where the definition of α_0 in (5.9) is used in the last step.

3) *Proof of $\lim_{n \rightarrow \infty} (u^n(x, U_c) - u^n(x, U)) = 0$:* Fix $\epsilon > 0$. Because of the continuity of $u^0(x, U_c)$ in x and Proposition 5.4, we can fix $\delta > 0$ and n_1 such that $|x - x_0| \leq \delta$ implies $|u^n(x, U_c) - u^n(x_0, U_c)| < \epsilon$ for all $n \geq n_1$. Applying the first part of this proof for δ gives n_0 such that for all $n \geq n_0$, there is some set \mathcal{B}^n with property i) and ii). So for each $n \geq n_0$, there is some $x(n) \in \mathcal{B}^n \cap [x_0 - \delta, x_0]$. By the definition of \mathcal{B}^n , the relation $u^n(x(n), U_c) = u^n(x(n), U)$ holds for all $n \geq n_0$. Moreover, $u^n(x, U)$ is increasing in x , so adding and subtracting $u^n(x(n), U_c) = u^n(x(n), U)$ and using that $x(n) \in [x_0 - \delta, x_0]$ yields

$$u^n(x_0, U_c) - u^n(x_0, U) \leq u^n(x_0, U_c) - u^n(x(n), U_c) < \epsilon.$$

With the arguments so far, we have shown that for every $x_0 > 0$ we have

$$\limsup_{n \rightarrow \infty} (u^n(x_0, U_c) - u^n(x_0, U)) \leq 0.$$

The result follows since $u^n(x_0, U_c) \geq u^n(x_0, U)$ for each $n \in \mathbb{N}$. \square

It remains to state and prove

Lemma 5.9. *Let B be a compact set of the form $[\epsilon, \frac{1}{\epsilon}]$ for $\epsilon \in (0, 1)$. Then $\{-J'_-(\lambda\varphi^n)\varphi^n \lambda \mid n \in \mathbb{N}, \lambda \in B\}$ is uniformly integrable.*

Proof. $\text{AE}_0(J) < \infty$ implies by the definition of $\text{AE}_0(J)$ that there are a constant $M \in (0, \infty)$ and $y_0 > 0$ such that we have

$$\sup_{q \in \partial J(y)} |q|y \leq MJ(y)$$

for $0 < y \leq y_0$. An application of this inequality for $y = \lambda\varphi^n$ and $q = J'_-(\lambda\varphi^n)$ on $\{\varphi^n \lambda \leq y_0\}$, some elementary calculations and $\lambda \in [\epsilon, \frac{1}{\epsilon}]$ yield

$$\begin{aligned} 0 &\leq -J'_-(\lambda\varphi^n)\varphi^n \lambda 1_{\{\lambda\varphi^n \leq y_0\}} \\ &\leq M |J(\lambda\varphi^n)| 1_{\{\lambda\varphi^n \leq y_0\}} \\ &\leq MJ(\varphi^n\epsilon)^+ + MJ(\varphi^n/\epsilon)^-. \end{aligned}$$

The family $(J(\varphi^n\epsilon)^+)_{n \in \mathbb{N}}$ is uniformly integrable by Lemma 5.1, and so is the family $(J(\varphi^n/\epsilon)^-)_{n \in \mathbb{N}}$ by Lemma 5.5. With the arguments so far, we have shown that the family $\{-J'_-(\lambda\varphi^n)\lambda\varphi^n 1_{\{\lambda\varphi^n \leq y_0\}} \mid n \in \mathbb{N}, \lambda \in B\}$ is uniformly integrable. Now fix some $x_0 \in -\partial J(y_0)$ and recall that any $x \in -\partial J(y)$ for $y \geq y_0$ satisfies $x \leq x_0$ and thus also $U_c(x) \leq U_c(x_0)$. The classical conjugacy relation between ∂J and ∂U_c gives

$$xy = U_c(x) - J(y) \leq U_c(x_0) + J(y)^-$$

for $y \geq y_0$. Applying this inequality for $y = \lambda\varphi^n$ and $x = -J'_-(\lambda\varphi^n)$ on $\{\lambda\varphi^n \geq y_0\}$ shows that $\{-J'_-(\lambda\varphi^n)\lambda\varphi^n 1_{\{\lambda\varphi^n \geq y_0\}} \mid n \in \mathbb{N}, \lambda \in B\}$ is dominated by $\{(U_c(x_0) + J(\lambda\varphi^n)^-) 1_{\{\lambda\varphi^n \geq y_0\}} \mid n \in \mathbb{N}, \lambda \in B\}$. This completes the proof since the latter family is uniformly integrable by Lemma 5.5. \square

The lower-semicontinuity in n of $u^n(x, U)$ stated in Theorem 5.3 is now a straightforward consequence of Propositions 5.4 and 5.8. For completeness, we formally carry out the argument.

Proof of Theorem 5.3. Since $u^n(x, U_c) - u^n(x, U)$ converges to 0 by Proposition 5.8, since $u^n(x, U_c)$ converges to $u^0(x, U_c)$ by Proposition 5.4 and because $u^n(x, U) \leq u^n(x, U_c)$ holds true for all $n \in \mathbb{N}_0$, we deduce from $\liminf (a_n + b_n) \geq \liminf a_n + \liminf b_n$ that

$$\begin{aligned} \liminf_{n \rightarrow \infty} u^n(x, U) &\geq \liminf_{n \rightarrow \infty} (u^n(x, U) - u^n(x, U_c)) + \lim_{n \rightarrow \infty} u^n(x, U_c) \\ &\geq 0 + u^0(x, U_c) \geq u^0(x, U). \end{aligned}$$

This completes the proof. \square

III.5.3 Putting everything together

On the way, we have separately proved the second case of Theorem 2.9. For completeness, we summarize the main steps.

Proof of Theorem 2.9 for ENCU. Theorem 5.3 and Proposition 5.2 give the convergence $\lim_n u^n(x, U) = u^0(x, U)$. For the second part, fix a maximizer f_x^n for $u^n(x, U)$ for every n . Proposition 4.1 shows that the sequence (f_x^n) contains a subsequence weakly converging to some $\bar{f} \in C^0(x)$. It then follows from Proposition 5.2, the optimality of f_x^n and $\lim_n u^n(x, U) = u^0(x, U)$ that

$$E_0[U(\bar{f})] \geq \limsup_{n \rightarrow \infty} E_n[U(f_x^n)] = \limsup_{n \rightarrow \infty} u^n(x, U) = u^0(x, U).$$

This shows that \bar{f} is a maximizer for $u^0(x, U)$ since $\bar{f} \in C^0(x)$. \square

It remains to give the proof for the stability of the goal-reaching problem. Recall from Remark 2.10 that this is the case where $U(x) = 1_{\{x \geq 1\}}$, so that $U_c(x) = x \wedge 1$ for $x \in (0, \infty)$. In particular, U_c is strictly increasing on $(0, 1)$ and uniformly bounded by 1.

Proof of Remark 2.10. The statement is clear for $x \geq 1$ since $u^n(x, U) = 1$ there for each $n \in \mathbb{N}$; so we assume that $x \in (0, 1)$. In Section III.5.2.1, strict monotonicity of U_c is used via Proposition 7.1 to show the existence of the lower bound c_1^n . A closer inspection of the argument there shows that we only need strict monotonicity of $u^n(x, U_c)$. But $u^n(x, U_c)$ admits a maximizer f^n (see the discussion following Theorem 2.9) and the constraint $E_n[\varphi^n f^n] \leq x < 1$ implies $P^n[f^n \in [0, 1]] > 0$. This yields strict monotonicity of $u^n(x, U_c)$ for $x \in (0, 1)$ since U_c is strictly increasing on $[0, 1)$ and we can prove $\lim_n u^n(x, U_c) = u^0(x, U_c)$ for $x \in (0, 1)$ as in Proposition 5.4. This implies $\limsup_n u^n(x, U) \leq \limsup_n u^n(x, U_c) = u^0(x, U_c)$. For the lim inf, we first fix for each $n \in \mathbb{N}$ a maximizer $f^n \in -\partial J(\lambda(n)\varphi^n)$ for $u^n(x, U_c)$ and recall that $f^n \in -\partial J(\lambda(n)\varphi^n)$ implies $-J'_+(\lambda(n)\varphi^n) \leq f^n \leq -J'_-(\lambda(n)\varphi^n)$ and that $U(-J'_\pm(y)) = U_c(-J'_\pm(y))$ holds for $y > 0$ (Lemma II.7.2). This gives

$$\liminf_{n \rightarrow \infty} u^n(x, U) \geq \liminf_{n \rightarrow \infty} E_n[U(f^n)] \geq \liminf_{n \rightarrow \infty} E_n[U_c(-J'_+(\lambda(n)\varphi^n))]. \quad (5.10)$$

We now fix a subsequence (n_k) realizing the $\liminf_n u^n(x, U)$ and such that the associated sequence $(\lambda(n_k))$ converges to $\bar{\lambda}$. As in Lemma 5.7 (and again using the modified version of Proposition 7.1), this gives $\bar{\lambda} > 0$. The assumptions that φ^0 has a continuous distribution and $\varphi^n \Rightarrow \varphi^0$ imply then that $U_c(-J'_\pm(\lambda(n_k)\varphi^{n_k})) \Rightarrow U_c(-J'_\pm(\bar{\lambda}\varphi^0))$. Moreover, as the function U_c is uniformly bounded, the right-hand side of (5.10) converges to $E_0[U_c(-J'_+(\bar{\lambda}\varphi^0))]$. But since φ^0 has a continuous distribution, it follows that $-J'_+(\bar{\lambda}\varphi^0) = -J'_-(\bar{\lambda}\varphi^0)$ P -a.s. and applying similar arguments in the reverse order, we find that

$$\liminf_{n \rightarrow \infty} u^n(x, U) \geq \liminf_{n \rightarrow \infty} E_n[U_c(-J'_-(\lambda(n)\varphi^n))] \geq \liminf_{n \rightarrow \infty} u^n(x, U_c).$$

With the arguments so far, we have proved that

$$\limsup_{n \rightarrow \infty} u^n(x, U) \leq u^0(x, U_c) = \limsup_{n \rightarrow \infty} u^n(x, U_c) \leq \liminf_{n \rightarrow \infty} u^n(x, U),$$

which gives $\lim_n u^n(x, U) = u^0(x, U_c)$. But the limit model is atomless, so we have $u^0(x, U) = u^0(x, U_c)$ by Theorem II.5.1 and the result follows. Finally, the convergence of the maximizers along a subsequence follows as in Proposition 5.2. \square

III.6 Conclusion

In this chapter, we study the stability along a sequence of models for a class of behavioural portfolio selection problems. The analyzed preference functionals allow for non-concave and non-smooth utility functions as well as for probability distortions. While there are several explicit results in the literature for behavioural portfolio selection problems in complete continuous-time markets, there are no comparable results for the discrete-time analogue.

Our convergence results demonstrate that the explicit results from the continuous-time model are approximately valid also for the discrete-time setting if the latter is *sufficiently close* to the continuous-time setting. As illustrated by a counterexample, the required notion of *sufficiently close* is slightly but strictly stronger compared to the stability results for *concave* utility maximization problems. The convergence results can also be applied to other situations such as (marginal) drift misspecification or changing time horizons.

III.7 Appendix

III.7.1 Non-smooth utility maximization

This appendix contains the results on non-smooth (concave) utility maximization which are relevant for the proofs in Section III.5. Following the notation there, we use $u^n(x, U) := v^n(x)$ to denote the value function. Recall that J is the conjugate of U (as well as U_c).

Proposition 7.1. *Fix $n \in \mathbb{N}_0$. Suppose that $E_n[J(\lambda\varphi^n)] < \infty$ for all $\lambda > 0$. Then the concave problem $u^n(x, U_c)$ has a solution $f^n \in C^n(x)$ for every $x > 0$. Every solution satisfies $f^n \in -\partial J(\lambda^n\varphi)$, where $\lambda^n \in (c_1^n, c_2^n)$ is a minimizer for*

$$\inf_{\lambda > 0} E_n[J(\lambda\varphi^n) + x\lambda] \quad (7.1)$$

and c_1^n and c_2^n are strictly positive constants.

Most of the statements contained in Proposition 7.1 are proved in greater generality in Bouchard et al. [22] and Westray and Zheng [107]. For completeness, we include a proof. We make use of Lemma 6.1 in Bouchard et al. [22] which reads in our setup as follows.

Lemma 7.2. *There is a function $\Phi^n : (-J(0), +\infty) \rightarrow (0, \infty)$ which is convex and increasing with $\lim_{x \rightarrow \infty} \Phi^n(x)/x = \infty$ and*

$$E_n[\Phi^n(J(y\varphi^n)^-)] \leq C^n + y \quad \text{for all } y > 0. \quad (7.2)$$

Proof of Proposition 7.1. 1) The existence of a maximizer $f^n \in C^n(x)$ in the present setting is shown in Theorem II.3.4. Remark II.3.3 there also shows

that $E_n[J(\lambda\varphi^n)] < \infty$ for all λ implies $u^n(x, U_c) < \infty$ for some $x > 0$ so that we can use Theorem II.4.1 there to get

$$E_n[J(\lambda\varphi^n)] = \sup_{x>0} \{u^n(x, U_c) - \lambda x\} \quad \text{for all } \lambda > 0.$$

Moreover, $u^n(x, U_c)$ is on $(0, \infty)$ finite and concave, hence continuous. This implies that we also have $u^n(x, U_c) = \inf_{\lambda>0} \{E_n[J(\lambda\varphi^n)] + x\lambda\}$. In order to find the upper bound c_2^n , we consider a minimizing sequence (λ_k) for (7.1) and show that it is bounded by some constant. Since (λ_k) is minimizing, it holds that

$$-E_n[J(\lambda_k\varphi^n)^-] + x\lambda_k \leq E_n[J(\lambda_k\varphi^n)] + x\lambda_k \leq u^n(x, U_c) + 1 \quad (7.3)$$

for k large enough. We use the function Φ^n introduced in Lemma 7.2. Then for all $\epsilon > 0$, there is some $x_0 > 0$ such that $\Phi^n(x)/x \geq 1/\epsilon$ for $x \geq x_0$, and then $x \leq x_0 + \epsilon\Phi^n(x)\mathbf{1}_{\{x \geq x_0\}} \leq x_0 + \epsilon\Phi^n(x)$ for all $x \geq 0$. Using (7.2), we compute that for some $C^n > 0$,

$$E_n[J(\lambda_k\varphi^n)^-] \leq x_0 + \epsilon E_n[\Phi^n(J(\lambda_k\varphi^n)^-)] \leq x_0 + \epsilon(C^n + \lambda_k).$$

Combining this inequality and (7.3) gives $(x - \epsilon)\lambda_k \leq u^n(x, U_c) + 1 + x_0 + \epsilon C^n$. Choosing $\epsilon = x/2 > 0$ shows that (λ_k) is bounded by some constant.

In order to find the lower bound c_1^n , we start with the case $J(0) < \infty$. Thanks to the existence of a maximizer for $u^n(x, U_c)$ and the strict monotonicity of U_c , we also deduce strict monotonicity of $u^n(x, U_c)$ and we infer $J(0) = U_c(\infty) > u(x, U_c)$. Together with the continuity of the function $H^n(\lambda) := E_n[J(\lambda\varphi^n)] + x\lambda$ in 0, we can find c_1^n such that the minimization in (7.1) can be reduced to $\lambda > c_1^n$. In the case $J(0) = \infty$, we can again find c_1^n since $E_n[J(\lambda\varphi^n)] \rightarrow \infty$ for $\lambda \rightarrow 0$.

2) With the arguments so far, we find a maximizer $f^n \in C^n(x)$ and some parameter $\lambda^n \in (0, \infty)$ satisfying

$$E_n[U_c(f^n)] = u^n(x, U_c) = \inf_{\lambda>0} \{E_n[J(\lambda\varphi^n)] + x\lambda\} = E_n[J(\lambda^n\varphi^n)] + x\lambda^n.$$

Suppose by way of contradiction that there exists a set $A \in \mathcal{F}^n$ satisfying $P^n[A] > 0$ and $f^n \notin \partial J(\lambda^n\varphi^n)$ on the set A . The conjugacy relation between U_c and J then implies $u^n(x, U_c) = E_n[U_c(f^n)] < E_n[J(\lambda^n\varphi^n)] + x\lambda^n$, which is the required contradiction. \square

III.7.2 Auxiliary results

Lemma 7.3. *The asymptotic elasticity condition $AE_0(J) < \infty$ is equivalent to the existence of two constants $\gamma > 0$ and $y_0 > 0$ such that*

$$J(\mu y) \leq \mu^{-\gamma} J(y) \quad \text{for all } \mu \in (0, 1] \text{ and } y \in (0, y_0].$$

Moreover, if $AE_0(J) < \infty$ is satisfied, then there are k_1, k_2 and $\gamma < 1$ such that $U(x) \leq k_1 + k_2 x^\gamma$ for $x \geq 0$.

Proof of Lemma 7.3. The equivalence is proved in Lemma 4.1 of Deelstra et al. [37]. We only prove the last implication. Similarly to Lemma 4.1 of Deelstra et al. [37], we argue that $U_c(\lambda x) \leq \lambda^\gamma U_c(x)$ holds for $x \geq x_0$ and $\lambda > 1$. Note first that Proposition 4.1. of Deelstra et al. [37] shows that $\text{AE}(J) < \infty$ and the growth condition (2.1) imply the existence of two constants x_0 and $\gamma < 1$ such that

$$yx - \gamma U_c(x) < 0 \quad \text{for } x \geq x_0 \text{ and } y \in \partial U_c(x). \quad (7.4)$$

By moving x_0 to the right if necessary, we may assume that $U(x_0)$ is positive. Now choose some $x > x_0$ and observe that $\lambda x \geq x_0$ for all $\lambda \geq 1$. We want to compare the functions $U_c(\lambda x)$ and $\lambda^\gamma U_c(x)$ for $\lambda > 1$. Let F be the concave function on $[1, \infty)$ defined by $F(\lambda) := U_c(\lambda x)$. Fix some $q \in \partial F(\lambda)$. By definition, this implies $U_c(zx) \leq U_c(\lambda x) + q(z - \lambda)$ and therefore that $\frac{q}{x} \in \partial U_c(\lambda x)$. Thus, it follows from (7.4) that

$$\lambda q - \gamma F(\lambda) < 0 \quad \text{for all } \lambda \geq 1 \text{ and } q \in \partial F(\lambda). \quad (7.5)$$

Set $G(\lambda) := \lambda^\gamma U_c(x)$. In order to complete the proof, we have to check that $(F - G)(\lambda) \leq 0$ for all $\lambda \geq 1$. Clearly, the function G satisfies the equation

$$\lambda G'(\lambda) - \gamma G(\lambda) = 0 \quad (7.6)$$

for all $\lambda \geq 1$. Since $G(1) = F(1)$, it follows from (7.5) and (7.6) that $0 > q - \gamma F(1) = q - \gamma G(1) = q - G'(1)$. Hence we have $q < G'(1)$ for all $q \in \partial F(1)$. Since G is continuously differentiable, there exists $\epsilon > 0$ such that $q < G'(\lambda)$ for all $q \in \partial F(1)$ and $\lambda \in [1, 1 + \epsilon)$. This gives

$$F(\lambda) \leq F(1) + q(\lambda - 1) < G(1) + G'(\lambda)(\lambda - 1) \leq G(\lambda) \quad (7.7)$$

for all $q \in \partial F(1)$ and $\lambda \in [1, 1 + \epsilon)$. To show that $F(\lambda) < G(\lambda)$ holds for all $\lambda > 1$, let $\bar{\lambda} := \inf\{\lambda > 1 : F(\lambda) = G(\lambda)\}$ and suppose that $\bar{\lambda} < \infty$. By the definition of $\bar{\lambda}$ and (7.7), we have $(F - G) < 0$ on $[1, \bar{\lambda})$ and $(F - G)(\bar{\lambda}) = 0$. This implies that

$$q_0 \geq G'(\bar{\lambda}) \quad (7.8)$$

for some $q_0 \in \partial F(\bar{\lambda})$. On the other hand, combining (7.5) and (7.6) gives $0 > \bar{\lambda} q - \gamma F(\bar{\lambda}) = \bar{\lambda} q - \gamma G(\bar{\lambda}) = \bar{\lambda} q - \bar{\lambda} G'(\bar{\lambda})$ for all $q \in \partial F(\bar{\lambda})$. The latter is equivalent to $G'(\bar{\lambda}) > q$ and gives the required contradiction to (7.8).

Above, it is proved that there exist $\gamma < 1$ and $x_0 > 0$ such that we have $U_c(\lambda x) \leq \lambda^\gamma U_c(x)$ for $x \geq x_0$ and $\lambda > 1$. This gives

$$U(x) \leq U_c\left(\frac{x}{x_0} x_0\right) \leq \left(\frac{x}{x_0}\right)^\gamma U_c(x_0) = U_c(x_0) \left(\frac{1}{x_0}\right)^\gamma x^\gamma$$

for $x \geq x_0$. Thus, choosing $k_1 := U(x_0)$ and $k_2 := U_c(x_0) (1/x_0)^\gamma$ gives $U(x) \leq k_1 + k_2 x^\gamma$ which is the desired result. \square

Lemma 7.4. *For all $\eta > 0$, there exists some constant $c(\eta) \geq 0$ such that*

$$|x + y|^\eta \leq c(\eta)(|x|^\eta + |y|^\eta).$$

Proof. For $\eta \in (0, 1]$, we use $z^\eta \geq z$ for $z \in (0, 1)$ to obtain

$$\left(\frac{|x|}{|x| + |y|}\right)^\eta + \left(\frac{|y|}{|x| + |y|}\right)^\eta \geq \frac{|x|}{|x| + |y|} + \frac{|y|}{|x| + |y|} = 1,$$

which is equivalent to $|x|^\eta + |y|^\eta \geq (|x| + |y|)^\eta$ and the result follows since $(|x| + |y|)^\eta \geq |x + y|^\eta$.

For $\eta > 1$, we rewrite $(|x| + |y|)^\eta$ as a suitable convex combination and use the convexity of z^η for $z \geq 0$ to obtain

$$(|x| + |y|)^\eta = 2^\eta \left(\frac{1}{2}|x| + \frac{1}{2}|y|\right)^\eta \leq 2^\eta \left(\frac{1}{2}|x|^\eta + \frac{1}{2}|y|^\eta\right) = 2^{\eta-1}(|x|^\eta + |y|^\eta).$$

The result follows since $|x + y|^\eta \leq (|x| + |y|)^\eta$. □

Chapter IV

Examples in incomplete markets

In this chapter, we study the portfolio selection problem (I.1.3) in a general financial market. We give sufficient conditions for the existence of an optimal strategy, and study the associated optimal final position as well as the optimal expected non-concave utility.

IV.1 Introduction

Expected utility maximization is the predominant investment decision rule in financial portfolio selection. One standard formulation of this problem is as follows: Given some discounted price process S , a (rational) agent with initial capital $x > 0$ tries to find a portfolio ϑ that maximizes the expected utility of terminal wealth, i.e.

$$E \left[U \left(x + \int_0^T \vartheta dS \right) \right] \longrightarrow \max_{\vartheta} \quad (1.1)$$

The standard case is to assume that U is concave which corresponds to the situation in which the agent is risk-averse. Compared to realistic applications, this formulation has two shortcomings. First, typical agents often do not have time to manage their portfolio by themselves. Instead, an agent often employs a portfolio manager to make financial investments on his behalf. The manager invests the agent's money in the financial market, and in return he receives a compensation via a function g , depending on the final position that he generates over some period of time. One prominent example used in Carpenter [32] is to consider $g(x) = a(x - K)^+ + F$ which means that the salary of the manager consists of a fixed amount F and a variable part depending on his performance. The manager with utility U is maximizing the expected utility of his own final wealth which is $g(X_T)$. Therefore he solves the portfolio optimization problem for the utility $U(g(\cdot))$

of his compensation, which is usually non-concave. The second shortcoming of the standard theory is, as discussed in detail in Chapter I, that people (agents as well as managers) are not always risk-averse and rational.

We are therefore interested in the problem (1.1) for a non-concave utility function U (corresponding to $U \circ g$ in the above notation). We formulate the non-concave utility maximization problem in a general financial market where the discounted asset prices are (locally bounded) semimartingales. We first give sufficient conditions for the existence of a maximizer in terms of a closedness assumption, under weak convergence, on the set of final positions generated by allowed trading. This enables us to tackle the problem in a systematic and unified way. In particular, this allows us to explain the results on the existence of a maximizer obtained previously by other authors in various specific frameworks. We also verify the closedness assumption in some models that are not covered by the existing literature. Moreover, we argue that the same assumption is also sufficient to prove the existence of a maximizer for similar optimization problems for more general functionals with distorted beliefs as employed in behavioural finance.

Having proved the existence of a maximizer, we then study its properties and the resulting optimal expected utility more thoroughly. We start with models on a finite probability space in order to bring out the intuition and structure, abstracting them from technical complexities. We show that the optimal final position satisfies, as in the classical concave case, the first order condition for optimality in the sense that the marginal utility of the optimal final position defines (up to a constant) a local martingale measure. We also illustrate with a counterexample that the classical interpretation of the resulting martingale measure as *least-favourable completion* does not carry over to the non-concave case.

While the optimal expected utility, as a function of the initial wealth x , is non-concave in general, we present sufficient conditions on a general model such that this function is concave. These conditions involve the *utility price* introduced in Jouini and Kallal [66] and can be seen as a natural generalization of the results from the complete market to the general model. To round off the chapter, we illustrate the use of our results with several explicit examples and give further links to the literature.

This chapter is organized as follows. The next section specifies the optimization problem and the notation. We also state the closedness assumption and show the existence of a maximizer under this assumption. In Section IV.3, we analyze models on finite probability spaces in more detail. Section IV.4 gives sufficient conditions on the model such that the optimal expected utility is concave. Finally, several examples are presented in Section IV.5.

IV.2 Problem formulation and existence

The following notation is used. If $x, y \in \mathbb{R}$, we denote $x^\pm = \max\{\pm x, 0\}$. For a probability space (Ω, \mathcal{F}, P) , let $L^0(\Omega, \mathcal{F}, P)$ (and $L^1(\Omega, \mathcal{F}, P)$) be the space of (equivalence classes of) \mathcal{F} -measurable (and integrable) random variables. The space $L^0_+(\Omega, \mathcal{F}, P)$ (and $L^1_+(\Omega, \mathcal{F}, P)$) consists of all non-negative elements of $L^0(\Omega, \mathcal{F}, P)$ (and $L^1(\Omega, \mathcal{F}, P)$). We sometimes drop the dependence on the probability space if it is clear from the context. By a distribution, we always refer to the distribution under P . For a sequence (f^n) of random variables, we denote weak convergence of (f^n) to f^0 by $f^n \Rightarrow f^0$. We use \sim to denote equality in distribution. A *quantile function* q_F of a distribution function F is a generalized inverse of F , i.e., a function $q_F : (0, 1) \rightarrow \mathbb{R}$ satisfying

$$F(q_F(s) -) \leq s \leq F(q_F(s)) \quad \text{for all } s \in (0, 1).$$

Quantile functions are not unique, but any two for a given F coincide a.e. on $(0, 1)$. Thus, we sometimes blur the distinction between “the” and “a” quantile function. A quantile function q_f of a random variable f is understood to be a quantile function q_F of the distribution F of the random variable f . If the sequence (f^n) converges weakly to f^0 , then any corresponding sequence (q_{f^n}) of quantile functions converges a.e. on $(0, 1)$ to q_{f^0} (see for instance Theorem 25.6 of Billingsley [19]). More properties of quantile functions can be found in Appendix A.3 of Föllmer and Schied [43].

IV.2.1 The optimization problem

We consider a model of a security market that consists of one bond and d stocks. We denote by $S = (S^i)_{1 \leq i \leq d}$ the price process of the d stocks and suppose that the price of the bond is constant and equal to 1. The process S is assumed to be a d -dimensional, locally bounded semimartingale on a filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{0 \leq t \leq T}, P)$, where the filtration $(\mathcal{F}_t)_{0 \leq t \leq T}$ satisfies the usual conditions. For simplicity, we assume that $\mathcal{F}_T = \mathcal{F}$. We focus on portfolio processes with initial capital x and predictable and S -integrable strategies ϑ . The value process of such a portfolio is then given by

$$X_t = X_t^\vartheta = X_0 + \int_0^t \vartheta_u dS_u, \quad 0 \leq t \leq T. \quad (2.1)$$

For $x > 0$ we denote by $\mathcal{X}(x)$ the family of value processes $X = (X_t)_{0 \leq t \leq T}$ with nonnegative capital at any time and with initial capital equal to x , i.e.,

$$\mathcal{X}(x) = \{X \geq 0 \mid X \text{ is defined by (2.1) with } X_0 = x\}. \quad (2.2)$$

We denote the set of equivalent (resp. absolutely continuous) local martingale measures for S by \mathcal{M}^e (resp. \mathcal{M}^a) and require the no-arbitrage type

assumption

$$\mathcal{M}^e \neq \emptyset. \quad (2.3)$$

For later purposes, we define the set of processes dual to (2.2) by

$$\mathcal{Y}(y) := \{Y \geq 0 \mid E[YX] \leq xy \text{ for all } X \in \mathcal{X}(x)\}$$

as well as two sets of random variables related to $\mathcal{X}(x)$ and $\mathcal{Y}(y)$ by

$$\mathcal{C}(x) := \{f \in L_+^0(\Omega, \mathcal{F}, P) \mid 0 \leq f \leq X_T, \text{ for some } X \in \mathcal{X}(x)\}, \quad (2.4)$$

$$\mathcal{D}(y) := \{h \in L_+^0(\Omega, \mathcal{F}, P) \mid 0 \leq h \leq Y_T, \text{ for some } Y \in \mathcal{Y}(y)\}. \quad (2.5)$$

Remark 2.1. Recall that for $f \geq 0$, it holds that

$$f \in \mathcal{C}(x) \iff \sup_{Q \in \mathcal{M}^e} E_Q[f] \leq x \iff \sup_{Q \in \mathcal{M}^a} E_Q[f] \leq x \quad (2.6)$$

by the general duality relationships between the terminal position of strategies and the densities of martingale measures; see, for instance, Section 11.3 of Delbaen and Schachermayer [38] for details. \diamond

We consider an economic agent in our model who has a non-concave utility function $U : (0, \infty) \rightarrow \mathbb{R}$ for wealth, that is, an increasing and right-continuous function satisfying and

$$\lim_{x \rightarrow \infty} \frac{U(x)}{x} = 0. \quad (2.7)$$

Observe that we do not assume that U is concave. To avoid any ambiguity, we set $U(x) = -\infty$ for $x < 0$ and define $U(0) := \lim_{x \searrow 0} U(x)$ and $U(\infty) := \lim_{x \nearrow \infty} U(x)$. Without loss of generality, we assume that $U(\infty) > 0$. The concave envelope of U is denoted by U_c , the conjugate J of U (and U_c) is defined by

$$J(y) := \sup_{x > 0} \{U(x) - xy\},$$

and the asymptotic elasticity of J is defined by

$$AE_0(J) := \limsup_{y \rightarrow 0} \sup_{q \in \partial J(y)} \frac{|q|y}{J(y)},$$

where ∂J denotes the subdifferential of J ; see Chapter II for precise definitions and several properties. We impose the following condition on J .

Assumption 2.2. $AE_0(J) < \infty$.

For a fixed initial capital $x > 0$, the goal of the agent is to maximize the expected non-concave utility from terminal wealth $E[U(X_T)]$. We are therefore interested in studying the optimization problem

$$\sup_{X \in \mathcal{X}(x)} E[U(X_T)]. \quad (2.8)$$

In abstract notation, we can replace (2.8) by the problem

$$u(x, U) := \sup\{E[U(f)] \mid f \in \mathcal{C}(x)\} \quad (2.9)$$

where we define $E[U(f)] := -\infty$ if $U(f)^- \notin L^1$. To exclude the trivial case we impose

Assumption 2.3. $u(x, U) < \infty$ for some $x > 0$.

Remark 2.4. The arguments of Lemma II.3.2 show that Assumptions 2.2 and 2.3 imply $u(x, U_c) < \infty$ for some $x > 0$, which in turn implies $u(x, U_c) < \infty$ for all $x > 0$ by concavity of U_c . \diamond

Let us shortly discuss the relation between the problems (2.9) and (II.2.2). Recall from (2.6) that a random variable f is in $\mathcal{C}(x)$ if and only if $f \in L_+^0$ and $\sup_{Q \in \mathcal{M}^e} E_Q[f] \leq x$. In Chapter II, we work with a fixed martingale measure $Q \approx P$ and consider the set

$$C(x) = C(x; Q) := \{f \in L_+^0 \mid E_Q[f] \leq x\}.$$

With this notation, (2.6) can be written in the form

$$\mathcal{C}(x) = \bigcap_{Q \in \mathcal{M}^e} C(x; Q).$$

Problem (2.9) can therefore be written as optimization problem

$$u(x, U; \mathcal{M}^e) = \sup \left\{ E[U(f)] \mid f \in L_+^0(\Omega, \mathcal{F}, P) \text{ with } \sup_{Q \in \mathcal{M}^e} E_Q[f] \leq x \right\} \quad (2.10)$$

with many (linear) constraints, one for each martingale measure Q . Problem (II.2.2) can be written as

$$u(x, U; \{Q\}) = \sup \left\{ E[U(f)] \mid f \in L_+^0(\Omega, \mathcal{F}, P) \text{ with } E_Q[f] \leq x \right\}. \quad (2.11)$$

Compared to (2.10), we consequently optimize in problem (II.2.2) over a larger class of payoffs since they do not necessarily need to satisfy $E_Q[f] \leq x$ for all $Q \in \mathcal{M}^e$. In the special case $\mathcal{M}^e = \{Q\}$, however, the problem $u(x, U, \mathcal{M}^e)$ in (2.10) reduces to $u(x, U, \{Q\})$ in (2.11) and thus has the same structure as the problem (II.2.2). For the sake of brevity, we continue in this section with the notation $u(x, U) = u(x, U, \mathcal{M}^e)$. Note that this is a slight abuse of notation since $u(x, U)$ defined in this chapter does not necessarily coincide with $u(x, U)$ defined in Chapter II.

IV.2.2 Existence of a maximizer

The goal of this section is to prove the existence of a maximizer for the problem $u(x, U)$. Because of the non-concavity of U , proving the existence via a suitable dual minimization problem (as, for instance, in Kramkov and Schachermayer [73]) or via Komlós-type arguments (as, for instance, in Kramkov and Schachermayer [74]) is not possible. Instead, we follow the idea of Carlier and Dana [31] (and later also applied in Rásonyi and Rodrigues [90]) and use Helly's selection principle. In order to be able to use that approach, the set $\mathcal{C}(x)$ essentially needs to satisfy a *closedness assumption under weak convergence* as follows.

Assumption 2.5. *If (f^n) is a sequence in $\mathcal{C}(x)$, F^n denotes the distribution function of f^n and $F^n \Rightarrow F$ for some distribution function F , then there is $f^* \in \mathcal{C}(x)$ with distribution F .*

This assumption allows us to show the existence of a maximizer for several models in a unified way. We discuss below in Sections IV.3–IV.5 sufficient conditions for this assumption as well as several explicit examples where it is satisfied.

Theorem 2.6. *Let Assumptions 2.2, 2.3 and 2.5 be satisfied. Moreover, we assume that we have one of the following cases:*

Case 1: *The non-concave utility function U is continuous.*

Case 2: *The non-concave utility function U is positive on $(0, \infty)$ and uniformly bounded.*

Then the problem $u(x, U)$ admits a solution $f^ \in \mathcal{C}(x)$.*

Proof. 1) We consider a maximizing sequence (f^n) in $\mathcal{C}(x)$ and denote by F^n the distribution function of f^n . Let us fix some $Q \in \mathcal{M}^e$. Every element f^n satisfies $E[\frac{dQ}{dP} f^n] \leq x$; it thus follows from Lemma III.4.2 that the sequence (f^n) is tight (with respect to the measure P). Helly's selection principle then gives a subsequence (n_k) and a distribution \bar{F} such that $\lim_{k \rightarrow \infty} F^{n_k}(a) = \bar{F}(a)$ holds for all continuity points a of \bar{F} , i.e., $F^{n_k} \Rightarrow \bar{F}$ as $k \rightarrow \infty$. Due to Assumption 2.5, there is some $f^* \in \mathcal{C}(x)$ with distribution \bar{F} . We claim that this f^* is a maximizer for $u(x, U)$.

2) In order to prove the optimality of f^* , we start with the observation that $E[U(f)^\pm] = \int_0^1 U(q_f(s))^\pm ds$. To see this, let $(\tilde{\Omega}, \tilde{\mathcal{F}}, \tilde{P})$ be a probability space that supports a random variable \tilde{U} with a uniform distribution on $(0, 1)$. Then $\tilde{f} := q_f(\tilde{U})$ has the same distribution as f and this gives

$$E[U(f)^\pm] = \tilde{E}_{\tilde{P}}[U(\tilde{f})^\pm] = \int_0^1 U(q_f(s))^\pm ds.$$

By construction, we have that (f^{n_k}) converges weakly to f^* and this implies a.e. convergence of the associated quantile functions. If U is continuous,

as assumed in case 1, $\max(\pm U(\cdot), 0)$ is continuous as well and we find that $U(q_{f^{n_k}}(s))^\pm \rightarrow U(q_{f^*}(s))^\pm$ a.e. For the negative part, Fatou's lemma yields

$$\liminf_{k \rightarrow \infty} E[U(f^{n_k})^-] = \liminf_{k \rightarrow \infty} \int_0^1 U(q_{f^{n_k}}(s))^- ds \geq \int_0^1 U(q_{f^*}(s))^- ds. \quad (2.12)$$

We show below in step 3) that the family $(U(q_{f^{n_k}}(s))^+)_{k \in \mathbb{N}}$ is uniformly integrable; this exploits the asymptotic elasticity condition (2.2). This gives

$$\lim_{k \rightarrow \infty} E[U(f^{n_k})^+] = \lim_{k \rightarrow \infty} \int_0^1 U(q_{f^{n_k}}(s))^+ ds = \int_0^1 U(q_{f^*}(s))^+ ds. \quad (2.13)$$

Together with (2.12), we arrive at

$$\lim_{k \rightarrow \infty} E[U(f^{n_k})] \leq \lim_{k \rightarrow \infty} E[U(f^{n_k})^+] - \liminf_{k \rightarrow \infty} E[U(f^{n_k})^-] = E[U(f^*)],$$

which completes the proof for case 1 since (f^n) (and thus also (f^{n_k})) is a maximizing sequence for $u(x, U)$.

If U is positive on $(0, \infty)$ and bounded above by some constant \bar{U} , as assumed in case 2, right-continuity of U gives $\limsup_k U(q_{f^{n_k}}(s)) \leq U(q_{f^*}(s))$ and Fatou's lemma (formally applied on $(\bar{U} - U(q_{f^{n_k}}(s)))_{k \in \mathbb{N}}$) yields

$$\lim_{k \rightarrow \infty} E[U(f^{n_k})] = \lim_{k \rightarrow \infty} \int_0^1 U(q_{f^{n_k}}(s)) ds \leq \int_0^1 U(q_{f^*}(s)) ds = E[U(f^*)],$$

which completes the proof for case 2 since (f^n) (and thus also (f^{n_k})) is a maximizing sequence for $u(x, U)$.

3) It remains to show the uniform integrability of $(U(q_{f^n}(s))^+)_{n \in \mathbb{N}}$. We define $g(x) := x1_{\{x \geq K\}}$ and observe that

$$\int_0^1 U(q_{f^n}(s))^+ 1_{\{U(q_{f^n}(s))^+ \geq K\}} ds = E[g(U(f^n)^+)] = E[U(f^n)^+ 1_{\{U(f^n)^+ \geq K\}}].$$

It is therefore sufficient to show uniform integrability of $(U(f^n)^+)_{n \in \mathbb{N}}$ which in turn is satisfied if the family $(U_c(f^n)^+)_{n \in \mathbb{N}}$ is uniformly integrable since $0 \leq U(f^n)^+ \leq U_c(f^n)^+$. But the latter is a consequence of $u(x, U_c) < \infty$ (obtained via Remark 2.4) and $\text{AE}_0(J) < \infty$; see for instance Note 2 and Lemma 1 in Kramkov and Schachermayer [74] or Lemma 5.5 in Westray and Zheng [107] for the nonsmooth case. \square

The two assumptions in Theorem 2.6 imposed on U cover (to the best of our knowledge) all the applications of expected non-concave utilities on the positive half-line. Case 1 includes, for instance, the non-concave utilities used in behavioural finance as well as situations in which risk-averse managers obtain performance-based salaries (see, for instance, Carpenter [32] and Basak and Makarov [9]). Case 2 covers the class of goal-reaching problems initiated by Kulldorff [76] and investigated further by Browne [27, 28].

There are also interesting examples where the non-concave utility is defined on \mathbb{R} . A closer inspection of the proof of Theorem 2.6 shows that the domain of U is only used (indirectly) in part 3) to show uniform integrability of $(U(f^n)^+)_n$. If one can obtain the latter property by a suitable argument or assumption (e.g. if U is bounded by above), the proof of Theorem 2.6 is also valid for a non-concave utility defined on \mathbb{R} . However, one has to think carefully about the right notion of admissibility for the allowed strategies such that one is still able to verify Assumption 2.5.

The structure of the proof can also be applied on more general preferences than expected (non-concave) utility. The essential ingredient of the preference functional is law-invariance which allows for a *quantile formulation*. In addition, one needs some structure ensuring uniform integrability of a maximizing sequence. As one possible extension, we discuss below *rank-dependent expected utilities* (RDU or RDEU; see Quiggin [89]). More preference functionals with a quantile formulation can be found in He and Zhou [56].

Remark 2.7. 1) In addition to a non-concave utility function, some models in behavioural finance suggest using non-linear expectations to account for the observation that people tend to overweight extreme events having small probabilities (see, for instance, Kahneman and Tversky [69] and references therein). The main building block for these models are functionals of the form

$$V(f) := \int_0^1 T(P[U(f) > x]) dx \quad (2.14)$$

for a continuous non-concave utility $U \geq 0$ and a *distortion* T which is an increasing and continuous function $T : [0, 1] \rightarrow [0, 1]$ satisfying $T(0) = 0$ and $T(1) = 1$. Analogously to our problem (2.9), one then is interested in

$$v(x) := \sup\{V(f) \mid f \in \mathcal{C}(x)\}. \quad (2.15)$$

Assumption 2.5 is (under suitable technical assumptions on U and T) also sufficient for the existence of a maximizer f^* for $v(x)$. Indeed, as in part 1) of the proof of Theorem 2.6, one starts with a maximizing sequence (f^n) for $v(x)$ and extracts along a subsequence a weak limit $f^* \in \mathcal{C}(x)$. In order to show the optimality of f^* , one needs to show uniform integrability with respect to the Lebesgue measure dy on $(0, 1)$ of $(T(P[U(f^n) > y]))_n$. Imposing some growth condition on U and T (see Assumption III.2.7), this reduces to showing that $\sup_n E[(f^n)^\kappa] < \infty$ for some $\kappa \in (0, 1)$ (Lemma III.4.3). The latter condition can be verified in a particular setting; a sufficient condition is the existence of a martingale measure Q with $(dQ/dP)^{\kappa/(\kappa-1)} \in L^1$. This criterion covers all the specific models (Examples 5.3, 5.5–5.8 and 5.11) discussed in this chapter.

2) In addition to non-concave utilities *and* non-linear expectations described in part 1), some behavioural theories suggest using a stochastic reference point R , with respect to which payoffs at the terminal time T are

evaluated. If the reference point is deterministic, this can be embedded in our analysis by defining a new utility $U_1(x) := U(x - R)$ and slightly modifying the arguments to account for the new domain of the utility. If the reference point is stochastic, one has to think carefully about the correct notion of admissible trading strategies. Jin et al. [63] allow strategies which lead to final positions satisfying $f - R \geq -L$ for a fixed (deterministic) maximal loss L . If we modify U such that it is defined on $(-L, \infty)$, we can use the same notion of admissible trading strategies. For a general reference point, this problem is rather involved. However, Kahneman and Tversky [69] claim that for most decision problems, the reference point R is the status quo or a particular asset position (e.g. all the money on the bank account). In both cases, the agent can then replicate the (possibly stochastic) reference point by some trading strategy with initial capital x^R . In that case, the agent with stochastic reference point R uses the amount x^R to replicate R and invests the remaining part $x - x^R$ as if he had a deterministic reference point x^R ; the problem is ill-posed (for this notion of admissible strategies) if $x - x^R < -L$. In summary, this remark shows that having a stochastic reference point is either very delicate (and essentially an unsolved problem) or rather easy to deal with. \diamond

IV.3 Models on finite probability spaces

The abstract analysis in Section IV.2 gives the existence of a maximizer under the rather abstract Assumption 2.5. The goal of this section is to focus on finite-dimensional models for which Assumption 2.5 is always satisfied (as we shall see in Lemma 3.1 below) and to present some elementary results and examples in order to bring out the intuition and structure, abstracting them from technical complexities.

Formally, consider an \mathbb{R}^{d+1} -valued process $(S_t)_{t=0}^T = (S_t^0, S_t^1, \dots, S_t^d)_{t=0}^T$ with $S_t^0 \equiv 1$, based on and adapted to the finite filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t=0}^T, P)$, which we write as $\Omega = \{\omega_1, \dots, \omega_N\}$. Without loss of generality we assume that \mathcal{F}_0 is trivial, that $\mathcal{F}_T = \mathcal{F}$ is the power set of Ω , and that $P[\{\omega_i\}] > 0$ for all $1 \leq i \leq N$. The setup here is a special case of the one introduced in Section IV.2 in which the price process S is a pure jump process with jumps occurring only at fixed dates.

Lemma 3.1. *Assumption 2.5 is satisfied in the setup described above. Theorem 2.6 can therefore be applied.*

Proof. Let us fix a sequence (f^n) in $\mathcal{C}(x)$ converging weakly to some distribution F . Let us fix some equivalent martingale measure $Q \in \mathcal{M}^e$. The condition $E_Q[f^n] \leq x$ implies $f^n(\omega_i) \leq x/Q[\{\omega_i\}]$ for all $1 \leq i \leq N$ since $f^n \geq 0$. The sequence (f^n) is thus uniformly bounded above by $x/\min_i Q[\{\omega_i\}]$, and we can extract a subsequence (f^{n_k}) converging (a.s.)

to some f^* . This implies that f^* has distribution F . Fatou's lemma gives $E_Q[f^*] \leq \liminf_k E_Q[f^{n_k}] \leq x$ for every $Q \in \mathcal{M}^e$ and it follows from (2.6) that $f^* \in \mathcal{C}(x)$. \square

The existence of a maximizer for the problem $u(x, U)$ in the finite-dimensional case can of course also be proved directly (with the same argument as for Lemma 3.1) and without Assumption 2.2. We have verified Assumption 2.5 in order to give a better understanding for that condition.

For an economic interpretation, a mere existence result is not very satisfying. In the sequel, our goal is therefore to describe the first order condition for optimality in more detail. In the classical case when U is concave, this property of the optimizer is a by-product of convex duality arguments. But in the non-concave case, convex duality cannot be applied. Instead, we use a standard marginal variation argument. In order to avoid dealing with generalized derivatives for non-concave functions, we impose here a stronger assumption on U . Recall that Ω is finite.

Proposition 3.2. *Let U be differentiable and strictly increasing and let f_x^* denote a solution for $u(x, U)$. There exist a martingale measure $\hat{Q} \in \mathcal{M}^a$ and some $\lambda > 0$ such that $U'(f_x^*) = \lambda \frac{d\hat{Q}}{dP}$ on the set $\{f_x^* > 0\}$.*

Proof. The equivalent formulation of $u(x, U)$ in (2.10) shows that $u(x, U)$ can be seen as optimization problem with infinitely many constraints (one for each $Q \in \mathcal{M}^e$). The idea is to rewrite these infinitely many constraints in terms of finitely many constraints by passing to extreme points. This allows us to use the classical tools from finite-dimensional optimization theory. Note first that \mathcal{M}^a can be identified with a bounded, closed, convex set in \mathbb{R}^n and is therefore the convex hull of its finitely many extreme points $\{Q^1, \dots, Q^m\}$.

Set $A := \{f_x^* > 0\}$ and note that $P[A] > 0$ since U is strictly increasing and $x > 0$. We show below that

$$E_{Q^i}[f 1_A] \leq 0 \text{ for } i = 1, \dots, m \implies E[U'(f_x^*) f 1_A] \leq 0 \quad (3.1)$$

holds for every $f \in L^0$. Applying Farkas' lemma (Corollary 22.3.1 in Rockafellar [95]) to (3.1) then gives

$$U'(f_x^*) = \sum_i^m \lambda_i \frac{dQ^i}{dP} \quad \text{on } A \quad (3.2)$$

for some $\lambda_1, \dots, \lambda_m \in \mathbb{R}_+$. We set $\lambda := \sum_i^m \lambda_i$ and note that $\lambda > 0$ by (3.2) since U is strictly increasing and $P[A] > 0$. We now define \hat{Q} via

$$\frac{d\hat{Q}}{dP} := \sum_{i=1}^m \frac{\lambda_i}{\lambda} \frac{dQ^i}{dP}.$$

The measure \hat{Q} is again in \mathcal{M}^a since it is a convex combination of (finitely many) elements of the convex set \mathcal{M}^a . In terms of λ and \hat{Q} , (3.2) reads as $U'(f_x^*) = \lambda \frac{d\hat{Q}}{dP}$ on A which was the claim.

It remains to show (3.1). For this, we fix some $f \in L^0$ with $E_{Q^i}[f1_A] \leq 0$ for all $i = 1, \dots, m$ and define $f(t) := f_x^* + tf1_A$ for $t > 0$. For t small enough, it holds that $f(t) \geq 0$ because $f_x^* > 0$ on A , and by construction, we have $E_{Q^i}[f(t)] \leq x$ for all $i = 1, \dots, m$. Since any $Q \in \mathcal{M}^e$ can be written as a convex combination of the m extreme points $\{Q^1, \dots, Q^m\}$, it follows that $E_Q[f(t)] \leq x$ for any $Q \in \mathcal{M}^a$ which gives by (2.6) that $f(t) \in \mathcal{C}(x)$. Therefore, we find from the optimality of $f_x^* = f(0)$ that

$$\frac{1}{t}(E[U(f(t))] - E[U(f(0))]) \leq 0 \quad (3.3)$$

for t small enough. We now analyze (3.3) in the limit for t going to 0. The expectation is, in the present setting, a sum of a finite number of terms, so that we can interchange limit and expectation (sum). On the set A^c , we have $f(t) = f(0)$, so the term $U(f(t)) - U(f(0))$ vanishes there. On the set A , we use that U is differentiable on $(0, \infty)$ to deduce that

$$\lim_{t \rightarrow 0} \frac{U(f_x^*(\omega) + tf(\omega)) - U(f_x^*(\omega))}{t} = U'(f_x^*(\omega))f(\omega)$$

for each $\omega \in A$. Inequality (3.3) therefore yields $E[U'(f_x^*)f1_A] \leq 0$, which completes the proof of (3.1). \square

If U is concave, the martingale measure \hat{Q} derived from the first order condition for optimality as in Proposition 3.2 has a clear economic meaning as the “least favourable market completion” with “fictitious securities” (see for instance Karatzas et al. [70] for detailed explanations of that concept in a slightly different setting; similar ideas also apply for our present finite-dimensional setup); moreover, it also plays a crucial role for the duality approach. In the next remark, we show that these properties do not carry over to the non-concave case.

Remark 3.3. Let f_x^* be a maximizer for $u(x, U)$ and let $\hat{Q} \in \mathcal{M}^a$ be the martingale measure derived in Proposition 3.2 that satisfies $U'(f_x^*) = \lambda \frac{d\hat{Q}}{dP}$ on $\{f_x^* > 0\}$ for some $\lambda > 0$. Let us compare the problems $u(x, U)$ and

$$u(x, U; \{\hat{Q}\}) := \sup\{E[U(f)] \mid f \in L_+^0(\Omega, \mathcal{F}, P) \text{ with } E_{\hat{Q}}[f] \leq x\} \quad (3.4)$$

defined in (2.11). The maximizer f_x^* for $u(x, U)$ is in $\mathcal{C}(x)$, so it satisfies $E_{\hat{Q}}[f_x^*] \leq x$ and it follows that f_x^* is also a candidate for the problem $u(x, U; \{\hat{Q}\})$. Example 3.4 below illustrates that the maximizer f_x^* for $u(x, U)$ is in general not optimal for the problem $u(x, U; \{\hat{Q}\})$ and that we may have

$$u(x, U) < \inf_{Q \in \mathcal{M}^a} u(x, U; \{Q\}) \leq u(x, U; \{\hat{Q}\}) \quad (3.5)$$

for particular initial capitals x . The economic interpretation is as follows. If we use \hat{Q} to define a pricing rule $f \mapsto E_{\hat{Q}}[f]$ (“pricing by marginal utility”), then an economic agent with initial capital x , *non-concave* utility function U and investing optimally is indifferent (of first order) towards small changes of the final position in a cost-neutral way (i.e., with respect to the pricing rule defined by \hat{Q}). However, if we consider as in (3.4) arbitrary contingent claims (not only those which can be generated by trading) and look only at the price $E_{\hat{Q}}[f]$, then there can be another payoff \tilde{f} with the same “price” $E_{\hat{Q}}[\tilde{f}] = E_{\hat{Q}}[f^*]$ and a higher expected non-concave utility. However, this payoff \tilde{f} cannot be generated by allowed trading. The inequality (3.5) shows that this result even does not change if we choose our pricing rule via another martingale measure. This is in marked contrast to the classical *concave* utility maximization where the solution to the original problem can also be obtained by working with only one (well-chosen) pricing rule defined via the marginal utility. This implies that the idea of introducing “fictitious securities” cannot be used to solve non-concave utility maximization problems. \diamond

It remains to give an explicit example for which one can verify (3.5).

Example 3.4. We set $T = 1$, $d = 1$, $n = 4$, $P[\{\omega_i\}] = 1/4$ and $S_1^1(\omega_i) = a_i$ for $a_1 = 3/2$, $a_2 = 5/4$, $a_3 = 3/4$ and $a_4 = 1/2$. Consider the non-concave utility function on \mathbb{R}_+ defined by

$$U(x) := \begin{cases} (x-1)^{1/3}, & x \geq 1, \\ -(1-x)^{2/3}, & x < 1. \end{cases}$$

The problem $u(1, U)$ can be solved explicitly; Theorem 3.1 in Bernard and Ghossoub [15] shows that the optimal risky holding is

$$\vartheta^* = \frac{1}{2^3} \left(\frac{(1/2)^{1/3} + (1/4)^{1/3}}{(1/2)^{2/3} + (1/4)^{2/3}} \right)^3.$$

The resulting optimal final position $f_1^* = 1 + \vartheta^*(S_1^1 - 1) > 0$ is strictly positive and satisfies $P[f_1^* = 1] = 0$. It follows that the martingale measure defined by $\frac{d\hat{Q}}{dP} := \frac{U'(f_1^*)}{E[U'(f_1^*)]}$ is equivalent to P .

We now show (3.5) for $x = 1$ which, in particular, implies that f_1^* is not optimal for $u(x, U; \{\hat{Q}\})$. The key insight for this is that $f_1^*(\omega_3)$ and $f_1^*(\omega_4)$ are in $\{U < U_c\} := \{x \in \mathbb{R}_+ \mid U(x) < U_c(x)\}$ and we have

$$U(f_1^*(\omega_3)) + U(f_1^*(\omega_4)) < U(0) + U(f_1^*(\omega_3) + f_1^*(\omega_4)). \quad (3.6)$$

We will see below in Theorem V.6.2 that such a candidate payoff f^* (for which two states have payoff values in the part where U is strictly convex)

cannot be optimal for the problem $u(1, U; \{Q\})$ (where one optimizes over payoffs that are not necessarily attainable by trading). In the present setting, all the quantities are explicitly determined and we can also directly show that f_1^* is not optimal for $u(1, U; \{Q\})$. More precisely, we shall show below that we can find for each $Q \in \mathcal{M}^a$ a random variable $f^Q \in L_+^0$ satisfying $E_Q[f^Q] \leq 1$, $f_1^*(\omega_i) = f^Q(\omega_i)$ for $i=1, 2$ and

$$P[f^Q = 0] = P[f^Q = f_1^*(\omega_3) + f_1^*(\omega_4)] = 1/4.$$

This then gives $E[U(f_1^*)] < E[U(f^Q)]$ because of (3.6). Since the distribution of f^Q (under P) is independent of Q , $E[U(f^Q)]$ is also independent of Q and (3.5) follows.

It remains to show how to construct f^Q . Since $f_1^*(\omega_3)$ and $f_1^*(\omega_4)$ are in $\{U < U_c\}$, the idea is to rearrange the payoff in the states ω_3 and ω_4 in such a way that the spread between the values is maximal. In order to also satisfy the budget constraint $E_Q[f^Q] \leq 1$, we choose the higher payoff value in the cheaper state (in the sense that dQ/dP is lower) and the lower payoff value in the more expensive state. Formally, if $Q[\{\omega_3\}] \leq Q[\{\omega_4\}]$, we define a modified payoff by $f^Q(\omega_i) := f_1^*(\omega_i)$ for $i = 1, 2$, $f^Q(\omega_3) := f_1^*(\omega_3) + f_1^*(\omega_4)$ and $f^Q(\omega_4) := 0$. Using that f_1^* and f^Q coincide for ω_1 and ω_2 , plugging in the explicit definition of f^Q and using that $Q[\{\omega_3\}] \leq Q[\{\omega_4\}]$, we obtain that

$$\begin{aligned} E_Q[f_1^*] - E_Q[f^Q] &= Q[\{\omega_3\}](f_1^*(\omega_3) - f^Q(\omega_3)) + Q[\{\omega_4\}](f_1^*(\omega_4) - f^Q(\omega_4)) \\ &= -Q[\{\omega_3\}]f_1^*(\omega_4) + Q[\{\omega_4\}]f_1^*(\omega_4) \\ &= f_1^*(\omega_4)(Q[\{\omega_4\}] - Q[\{\omega_3\}]) \geq 0, \end{aligned} \quad (3.7)$$

which implies $E_Q[f^Q] \leq E_Q[f_1^*] \leq 1$. The conditions $P[f^Q = 0] = 1/4$ and $P[f^Q = f_1^*(\omega_3) + f_1^*(\omega_4)] = 1/4$ are also satisfied since $P[\{\omega_i\}] = 1/4$ for $i=3, 4$. If $Q[\{\omega_3\}] > Q[\{\omega_4\}]$, we define a modified payoff by $f^Q(\omega_i) := f_1^*(\omega_i)$ for $i = 1, 2$, $f^Q(\omega_3) := 0$ and $f^Q(\omega_4) := f_1^*(\omega_3) + f_1^*(\omega_4)$. A calculation similar to that in (3.7) yields $E_Q[f^Q] \leq 1$. This ends the example.

The attentive reader might object that Proposition 3.2 is proved for a differentiable function U , while the function U used in Example 3.4 is not differentiable in 1. But it is obvious that one can slightly change the function U around 1 to “smooth out” the kink in 1 so that f_1^* is still optimal.

IV.4 A sufficient condition for $u(x, U) = u(x, U_c)$

The analysis for models on finite probability spaces in Section IV.3 shows that the problems $u(x, U)$ and $u(x, U_c)$ are not directly linked in general. In the special case $\mathcal{M}^e = \{Q\}$, the problem $u(x, U) = u(x, U; \mathcal{M}^e)$ reduces to $u(x, U; \{Q\})$ defined in (2.11). For that problem, it is known from Section II.5 that there is a large class of models for which $u(x, U; \{Q\})$ and

$u(x, U_c; \{Q\})$ coincide. The goal of this section is to describe a subclass of our general (not necessarily complete) financial market for which $u(x, U; \mathcal{M}^e)$ and $u(x, U_c; \mathcal{M}^e)$ coincide as well.

In the analysis in Section II.5, there are two essential ingredients leading to $u(x, U; \{Q\}) = u(x, U_c; \{Q\})$. The first (explicit) one is that one can generate any payoff *distribution*. A second (implicit) one is that one can generate a given distribution for a reasonable price. In order to make this precise and to motivate an analogous assumption in the present setting, we first introduce some additional notation.

Dybvig [42] considers the setting of a complete market and defines the *distributional price* of a distribution F to be the lowest initial capital such that there is a strategy leading to a final payoff with the distribution F . Note that although completeness allows one to generate by trading any *payoff* on the given underlying space (Ω, \mathcal{F}_T) , it depends on the structure of that space which *distributions* F can actually be obtained from *random variables* on (Ω, \mathcal{F}_T) . The natural analogue in our general setting is to define the distributional price to be

$$\text{DP}(F) := \inf \left\{ \sup_{Q \in \mathcal{M}^e} E_Q[f] \mid f \in L_+^0(\Omega, \mathcal{F}_T, P) \text{ with } f \sim F \right\}. \quad (4.1)$$

Besides the obvious drawback that the set $\{f \in L_+^0(\Omega, \mathcal{F}_T, P), f \sim F\}$ might be empty, it may also happen that $\text{DP}(G) < \text{DP}(F)$ for a distribution G dominating the distribution F in some stochastic order. This indicates that the definition (4.1) has economically unreasonable consequences. One alternative is therefore to consider the price

$$\inf \left\{ \sup_{Q \in \mathcal{M}^e} E_Q[f] \mid f \in L_+^0(\Omega, \mathcal{F}_T, P) \text{ with } f \sim G \text{ and } G \succeq_2 F \right\}, \quad (4.2)$$

where \succeq_2 denotes second order stochastic dominance. As noticed in Jouini and Kallal [66], this is also the minimal amount for which a risk-averse agent can attain the same utility level as with a payoff with distribution F . They therefore refer to the above expression (4.2) as the *utility price*. It is shown in Theorem 4 of Jouini and Porte [68] that the expression (4.2) can also be written equivalently in terms of quantiles as follows. Recall the notation $\mathcal{D}(y)$ from (2.5).

Definition 4.1. For a distribution F on \mathbb{R}_+ , the *utility price* is defined as

$$\text{UP}(F) := \sup_{h \in \mathcal{D}(1)} \left\{ \int_0^1 q_F(s) q_h(1-s) ds \right\}.$$

If $\mathcal{M}^e = \{Q\}$ is a singleton and the probability space $(\Omega, \mathcal{F}_T, P)$ is atomless, we can fix a uniformly distributed random variable \mathcal{U} on $(0, 1)$ satisfying

$q_{\frac{dQ}{dP}}(\mathcal{U}) = \frac{dQ}{dP}$ a.s. For an arbitrary distribution F on \mathbb{R}_+ , there is then a self-financing strategy starting from the initial capital $X_0 = E_Q[q_F(1 - \mathcal{U})]$ and leading to the final position $X_T = q_F(1 - \mathcal{U})$ with distribution F . Since $\frac{dQ}{dP} \in \mathcal{D}(1)$, we find that

$$\text{UP}(F) \geq \int_0^1 q_{\frac{dQ}{dP}}(s) q_F(1 - s) ds = E_Q[q_F(1 - \mathcal{U})] \geq \text{DP}(F)$$

and we even obtain equality since the *utility price* is always smaller than the *distributional price*; see part 1) of Remark 4.4 below. This line of arguments shows that in the case $\mathcal{M}^e = \{Q\}$ with an atomless probability space $(\Omega, \mathcal{F}_T, P)$, the *distributional price* is equal to the *utility price*. The natural analogue in our general setting is as follows.

Assumption 4.2. *For any distribution F on \mathbb{R}_+ with $\text{UP}(F) < \infty$, there exists $f \in \mathcal{C}(\text{UP}(F))$ with $f \sim F$.*

Let us shortly summarize the above motivation.

Remark 4.3. If $\mathcal{M}^e = \{Q\}$ and $(\Omega, \mathcal{F}_T, P)$ is atomless (as in the typical complete Brownian models; see Example II.2.7), then Assumption 4.2 is satisfied. \diamond

Before we discuss Assumption 4.2 in more detail, we briefly comment on the relation between $\text{UP}(F)$ and $\text{DP}(F)$ and show that the *utility price* is lower-semicontinuous with respect to weak convergence. The latter property is used later to compare Assumptions 2.5 and 4.2.

Remark 4.4. 1) We always have $\text{UP}(F) \leq \text{DP}(F)$. To see this, fix $f \sim F$ with $x := \sup_{Q \in \mathcal{M}^e} E_Q[f] < \infty$; by (2.6) and the definition of $\mathcal{C}(x)$, this implies that there is some $X \in \mathcal{X}(x)$ with $f \leq X_T$ P -a.s. The distribution G of X_T therefore dominates F stochastically in the second order and it follows from the Hardy–Littlewood inequality and the definition of $\mathcal{D}(1)$ that

$$\int_0^1 q_G(s) q_h(1 - s) ds \leq E[hX_T] \leq x$$

holds for every $h \in \mathcal{D}(1)$. The result then follows since $f \sim F$ was arbitrary. If one defines the utility price in terms of (4.2), it is sufficient to notice that every candidate for (4.1) is also one for (4.2).

2) Every sequence (F^n) of distributions with $F^n \Rightarrow F$ satisfies

$$\text{UP}(F) \leq \liminf_{n \rightarrow \infty} \text{UP}(F^n). \quad (4.3)$$

Indeed, assume by way of contradiction that there is $\epsilon > 0$ such that $\text{UP}(F) - \epsilon > \liminf_n \text{UP}(F^n)$. By definition, there exists $\tilde{h} \in \mathcal{D}(1)$ with

$\text{UP}(F) - \epsilon \leq \int_0^1 q_{\bar{h}}(s) q_F(1-s) ds$. Recall that $F^n \Rightarrow F$ implies $q_{F^n}(s) \rightarrow q_F(s)$ a.s. Using Fatou's lemma and the definition of the utility price, we obtain

$$\text{UP}(F) - \epsilon \leq \liminf_{n \rightarrow \infty} \int_0^1 q_{\bar{h}}(s) q_{F^n}(1-s) ds \leq \liminf_{n \rightarrow \infty} \text{UP}(F^n),$$

which is the required contradiction. \diamond

Let us briefly discuss direct consequences of Assumption 4.2.

Remark 4.5. 1) Let F be the uniform distribution on $(0, 1)$. It follows from the definition of a quantile function that $q_F(s) \leq 1$ for all $s \in (0, 1)$. This yields

$$\int_0^1 q_F(s) q_h(1-s) ds \leq \int_0^1 q_h(1-s) ds = E[h] \leq 1$$

for every $h \in \mathcal{D}(1)$, where we used in the last step the definition of $\mathcal{D}(1)$. Taking the supremum over all $h \in \mathcal{D}(1)$ gives $\text{UP}(F) \leq 1 < \infty$. We thus can apply Assumption 4.2 for F to obtain $f \in \mathcal{C}(\text{UP}(F))$ with $f \sim F$. This shows that the probability space $(\Omega, \mathcal{F}_T, P)$ supports a continuous distribution which implies that $(\Omega, \mathcal{F}_T, P)$ is atomless.

Remark 4.3 shows that under the additional assumption $\mathcal{M}^e = \{Q\}$, we also have the converse direction in the sense that an atomless underlying probability space implies Assumption 4.2. However, this does not hold for general models as we see below in Example 5.11.

2) *Assumption 4.2 implies Assumption 2.5:* Let us fix a sequence (f^n) in $\mathcal{C}(x)$ with $f^n \sim F^n$ and $F^n \Rightarrow F$. For fixed $n \in \mathbb{N}$ and fixed $h \in \mathcal{D}(1)$, the Hardy–Littlewood inequality and $f^n \in \mathcal{C}(x)$ imply

$$\int_0^1 q_{F^n}(s) q_h(1-s) ds \leq E[h f^n] \leq x,$$

and taking the supremum over $h \in \mathcal{D}(1)$ gives $\text{UP}(F^n) \leq x$. But the *utility price* is lower-semicontinuous with respect to weak convergence by part 2) of Remark 4.4, so that we obtain $\text{UP}(F) \leq \liminf_n \text{UP}(F^n) \leq x$. Assumption 4.2 then yields $f^* \in \mathcal{C}(x)$ with distribution F . \diamond

Our main result of this section now relates $u(x, U)$ and $u(x, U_c)$. In particular, it implies that the maximizer for $u(x, U)$ also maximizes $u(x, U_c)$. Note that this allows one to describe the maximizer for $u(x, U)$ in the same way as one can describe the maximizers for a concave utility maximization problem; see for instance Theorem 3.2 of Bouchard et al. [22].

Theorem 4.6. *Suppose that Assumptions 2.2, 2.3 and 4.2 are satisfied. Then the non-concave problem $u(x, U)$ admits a maximizer and it holds that*

$$u(x, U) = u(x, U_c) \quad \text{for all } x > 0. \quad (4.4)$$

The proof is based on (a modification of) Proposition II.5.3. For completeness, we include the modified version we use here. As in the previous chapters, we use the notation $\{U < U_c\} := \{x \in \mathbb{R}_+ \mid U(x) < U_c(x)\}$; it is shown in Lemma II.2.11 that $\{U < U_c\}$ can be written as a countable union of finite disjoint open intervals $\{U < U_c\} = \bigcup_i (a_i, b_i)$ for some a_i and b_i .

Proposition 4.7. *Let $(\Omega, \mathcal{F}_T, P)$ be atomless and let $f \in \mathcal{C}(x)$ be fixed. Then there is a distribution F^* such that for each $h \in \mathcal{D}(1)$, there exists $f_h^* \in L_+^0(\Omega, \mathcal{F}_T, P)$ with $f_h^* \sim F^*$ such that $E[hf_h^*] \leq E[hf]$ and*

$$E[U(f_h^*)] = E[U_c(f)].$$

Proof. This result is a slight modification of Proposition II.5.3. For the convenience of the reader, we include the argument. We first consider the case that $E[h] > 0$. We apply¹ Proposition II.5.3 to f for the pricing density $\varphi := h/E[h]$ to obtain a modified payoff $f_h^* \in L_+^0(\Omega, \mathcal{F}_T, P)$ satisfying

$$E[hf_h^*] \leq E[hf], \quad (4.5)$$

$$E[U_c(f_h^*)] = E[U_c(f)] \quad (4.6)$$

and $P[f_h^* \in \{U < U_c\}] = 0$. We show below that the distribution F_h of f_h^* constructed in Proposition II.5.3 is independent of the particular choice of $h \in \mathcal{D}(1)$ which then yields the desired result by choosing $F^* = F_h$ for the case $E[h] > 0$. If $E[h] = 0$, we choose an arbitrary $f \sim F^*$ which is possible since $(\Omega, \mathcal{F}_T, P)$ is atomless. Condition (4.5) is satisfied since $h = 0$ a.s.

In order to show that F_h is independent of h , we express F_h in terms of the distribution F of f . A closer inspection of the construction of f_h^* in part 2) of the proof of Proposition II.5.3 shows that $f_h^* := f$ on $\{f \in \{U = U_c\}\}$ and that

$$\{f_h^* \in [a_i, b_i]\} = \{f \in [a_i, b_i]\} \quad (4.7)$$

for all i . For $s \notin \bigcup_i [a_i, b_i]$, we now show $\{f_h^* \leq s\} = \{f \leq s\}$ as follows. For $\omega \in \{f \notin \bigcup_i [a_i, b_i]\}$, we have $U(f(\omega)) = U_c(f(\omega))$ so that $f(\omega) = f_h^*(\omega)$. This implies

$$\{f \leq s\} \cap \left\{f \notin \bigcup_i [a_i, b_i]\right\} = \{f_h^* \leq s\} \cap \left\{f \notin \bigcup_i [a_i, b_i]\right\}. \quad (4.8)$$

For the part $\{f \leq s\} \cap \{f \in \bigcup_i [a_i, b_i]\}$, recall first that $s \notin \bigcup_i [a_i, b_i]$. Therefore there is a subset $I \in \mathbb{N}$ (depending on s) such that

$$\{g \leq s\} \cap \left\{g \in \bigcup_i [a_i, b_i]\right\} = \left\{g \in \bigcup_{i \in I} [a_i, b_i]\right\} \quad (4.9)$$

¹In Proposition II.5.3, it is assumed that $h > 0$. However, this property is not used in the proof there.

holds for any random variable g . Applying (4.9) for f , using (4.7), applying (4.9) for f_h^* and using again (4.7) give

$$\begin{aligned}
\{f \leq s\} \cap \left\{ f \in \bigcup_i [a_i, b_i] \right\} &= \left\{ f \in \bigcup_{i \in I} [a_i, b_i] \right\} \\
&= \left\{ f_h^* \in \bigcup_{i \in I} [a_i, b_i] \right\} \\
&= \{f_h^* \leq s\} \cap \left\{ f_h^* \in \bigcup_i [a_i, b_i] \right\} \\
&= \{f_h^* \leq s\} \cap \left\{ f \in \bigcup_i [a_i, b_i] \right\}.
\end{aligned} \tag{4.10}$$

Combining (4.8) and (4.10) yields $\{f \leq s\} = \{f_h^* \leq s\}$ and we deduce $F_h(s) = F(s)$ for all $s \notin \bigcup_i [a_i, b_i]$. For $s \in \bigcup_i [a_i, b_i]$, note that $F_h(s) = F_h(a_i)$ for all $s \in (a_i, b_i)$ since $f_h^* \notin \{U < U_c\}$. The proof is therefore completed by showing that $F_h(a_i)$ is independent of h . In part 3) of the proof of Proposition II.5.3, it is shown that $E[U_c(f)1_{\{f \in (a_i, b_i)\}}] = E[U_c(f_h^*)1_{\{f \in (a_i, b_i)\}}]$ which together with $f_h^* = f$ on $\{f \in \{U = U_c\}\}$ gives

$$E[U_c(f)1_{\{f \in [a_i, b_i]\}}] = E[U_c(f_h^*)1_{\{f \in [a_i, b_i]\}}]. \tag{4.11}$$

Because of $f_h^* \notin \{U < U_c\}$, equality (4.11) can be rewritten as

$$E[U_c(f)1_{\{f \in [a_i, b_i]\}}] = P[f_h^* = a_i]U_c(a_i) + P[f_h^* = b_i]U_c(b_i),$$

where $P[f_h^* = a_i] + P[f_h^* = b_i] = P[f \in [a_i, b_i]]$. Solving for $P[f_h^* = a_i]$ gives

$$P[f_h^* = a_i] = \frac{U_c(b_i)P[f \in [a_i, b_i]] - E[U_c(f)1_{\{f \in [a_i, b_i]\}}]}{U_c(b_i) - U_c(a_i)}.$$

Note that the right-hand side is independent of the particular choice $h \in \mathcal{D}(1)$. Hence also

$$F_h(a_i) = \lim_{x \nearrow a_i} F_h(x) + P[f_h^* = a_i] = \lim_{x \nearrow a_i} F(x) + P[f_h^* = a_i]$$

is independent of the particular choice $h \in \mathcal{D}(1)$. □

Theorem 4.6 follows now from Proposition 4.7, as follows.

Proof of Theorem 4.6. It is argued in part 2) of Remark 4.5 that Assumption 4.2 implies Assumption 2.5; the existence part of Theorem 4.6 therefore follows from Theorem 2.6. The inequality “ \leq ” for (4.4) follows from $U \leq U_c$. For “ \geq ”, we start with some $f \in \mathcal{C}(x)$. The payoff f satisfies $E[hf] \leq x$ for all $h \in \mathcal{D}(1)$ by the definition of $\mathcal{D}(1)$ and $\mathcal{Y}(1)$. Moreover, it is argued in part 1) of Remark 4.5 that the probability space $(\Omega, \mathcal{F}_T, P)$ is atomless.

This allows us to apply Proposition 4.7 to f for an arbitrary $h \in \mathcal{D}(1)$. This gives a modified payoff $f_h^* \in L_+^0(\Omega, \mathcal{F}_T, P)$ satisfying

$$E[hf_h^*] \leq E[hf], \quad (4.12)$$

$$E[U(f_h^*)] = E[U_c(f)], \quad (4.13)$$

and such that the distribution of f_h^* is independent of the choice of $h \in \mathcal{D}(1)$. We denote this distribution by F^* and infer that $q_{f_h^*}(s) = q_{F^*}(s)$ a.s. for each $h \in \mathcal{D}(1)$. Using the Hardy–Littlewood inequality (Theorem A.24 in Föllmer and Schied [43]), (4.12) and $f \in \mathcal{C}(x)$, we obtain

$$\int_0^1 q_{F^*}(s)q_h(1-s)ds = \int_0^1 q_{f_h^*}(s)q_h(1-s)ds \leq E[hf_h^*] \leq E[hf] \leq x. \quad (4.14)$$

Taking the supremum over $h \in \mathcal{D}(1)$ in (4.14) therefore gives $\text{UP}(F^*) \leq x$. Finally, we use Assumption 4.2 to obtain $f^* \in \mathcal{C}(\text{UP}(F^*)) \subseteq \mathcal{C}(x)$ with $f^* \sim F^*$. The payoff f^* has by construction the same distribution as f_h^* (for an arbitrary $h \in \mathcal{D}(1)$), and together with (4.13), we therefore arrive at

$$E[U_c(f)] = E[U(f_h^*)] = E[U(f^*)] \leq u(x, U).$$

The inequality “ \geq ” in (4.4) follows since $f \in \mathcal{C}(x)$ is arbitrary. \square

At this point, it seems appropriate to comment on other assumptions in the literature leading to $u(x, U) = u(x, U_c)$.

Remark 4.8. 1) It is known from the literature on non-smooth (concave) utility maximization that the problem $u(x, U_c) = u(x, U_c; \mathcal{M}^e)$ has (under the Assumptions 2.2 and 2.3) a maximizer $f^* \in \partial J(Y_T^*)$ for some (dual) element $Y^* \in \mathcal{Y}(y)$ which can be determined via a dual problem of minimizing $E[J(Y_T^*)]$ over the set $Y \in \mathcal{Y}(y)$ and a suitable $y > 0$ (see for instance Theorem 3.2 of Bouchard et al. [22]). If Y_T^* has a continuous distribution, it follows from subdifferential calculus that $P[f^* \in \{U < U_c\}] = 0$, and this implies in turn that f^* is also a maximizer for $u(x, U_c)$ and that $u(x, U) = u(x, U_c)$; see for instance Lemma II.5.7 for detailed arguments. This line of argument shows that the continuity of the dual minimizer is another sufficient criterion for $u(x, U) = u(x, U_c)$. Bichuch and Sturm [17] explore this more thoroughly and show that some classes of incomplete market models such as the lognormal mixture models of Brigo and Mercurio [24] satisfy this criterion. Like our Assumption 4.2, their criterion is relatively difficult to verify. In contrast to our Assumption 4.2, however, their criterion cannot be used to show the existence of a solution for the behavioural portfolio selection problem (2.15) with distorted beliefs as discussed in part 1) of Remark 2.7. In Example 5.8 below, we present a model in which $u(x, U)$ and $u(x, U_c)$ coincide on $(0, \infty)$, our Assumption 4.2 is satisfied and the dual minimizer *always* has a *discontinuous* distribution.

2) Rásonyi and Rodrigues [90] study the problem (2.15) for functionals on the whole real line (similar to the one discussed in part 1) of Remark 2.7) and allowing for a stochastic reference point R . They assume that

- (i) there exists $Q^* \in \mathcal{M}^e$ such that $\rho := dQ^*/dP$ has a continuous distribution and satisfies $E[\rho^p] < \infty$ for all (positive and negative) p ,
- (ii) there exists an \mathcal{F}_T -measurable random variable U^* uniformly distributed on $(0, 1)$ and independent of ρ , and
- (iii) every $\sigma(\rho, R, U^*)$ -measurable random variable f (integrable with respect to Q^*) is *hedgeable* in the sense that there exists a self-financing strategy ϑ^* such that $X_T^{\vartheta^*} = f$ holds and $(X_t^{\vartheta^*})_{t \in [0, T]}$ is a martingale under Q^* .

Under these assumptions (and some growth conditions imposed on T and U), they prove the existence of a maximizer for their optimization problem. In order to relate their assumptions to our Assumption 4.2 (and therefore also Assumption 2.5), fix a distribution F on \mathbb{R}_+ with $\text{UP}(F) < \infty$ and suppose that (i)–(iii) above are satisfied. The distribution of the random variable defined by $f := q_F(1 - F_\rho(\rho))$ is F since $F_\rho(\rho)$ is uniformly distributed on $(0, 1)$. Moreover, f is hedgeable by (iii) since it is a function of ρ . This means that f is the final position $X_T^{\vartheta^*}$ of a self-financing strategy ϑ^* such that $(X_t^{\vartheta^*})_{t \in [0, T]}$ is a martingale under Q^* . We deduce that $X^{\vartheta^*} \in \mathcal{X}(X_0^{\vartheta^*})$ since $f = X_T^{\vartheta^*}$ is nonnegative. We infer from the definition of $\text{UP}(F)$, the Hardy–Littlewood inequality and the martingale property of $(X_t^{\vartheta^*})_{t \in [0, T]}$ that

$$\text{UP}(F) \geq \int_0^1 q_\rho(s) q_F(1 - s) ds = E_{Q^*}[X_T^{\vartheta^*}] = X_0^{\vartheta^*}.$$

We conclude that $f \in \mathcal{C}(X_0^{\vartheta^*}) \subseteq \mathcal{C}(\text{UP}(F))$ which means that Assumption 4.2 (and therefore also Assumption 2.5) are satisfied. This implies that $u(x, U) = u(x, U_c)$ holds for the model class discussed in Rásonyi and Rodrigues [90]. While they focus on general functionals allowing for distorted beliefs and show the existence of a maximizer, our result shows that by restricting oneself to standard beliefs, one additionally gets a description of the function $u(x, U)$ and the optimal final position associated to $u(x, U)$. \diamond

While Assumption 4.2 as well as the assumptions in Bichuch and Sturm [17] and Rásonyi and Rodrigues [90] are sufficient for $u(x, U) = u(x, U_c)$, they are not necessary. This is illustrated below in Example 5.3. It would be interesting to see a necessary and sufficient condition for $u(x, U) = u(x, U_c)$ in the general case.

IV.5 Examples

In Section IV.3, we have discussed models on finite probability spaces. In that case, Assumption 2.5 is satisfied and the underlying probability space is atomic which implies, by part 1) of Remark 4.5, that Assumption 4.2 always fails to hold. The goal of this section is to discuss several model classes in continuous time.

IV.5.1 Special case $\mathcal{M}^e = \{Q\}$

In the special case $\mathcal{M}^e = \{Q\}$, the problem $u(x, U)$ can be rewritten as optimization problem (2.11) with only one constraint. For that problem, the existence of a maximizer for $u(x, U)$ is proved in Theorem II.3.4. The proof there uses *Fatou's lemma in several dimensions* in a version proved in Balder [6]. The drawback of the proof of Theorem II.3.4 is that it only works for preference functionals with standard beliefs and cannot be applied for the preference functionals with distortion of the form (2.14). In this section, we verify that Assumption 2.5 is satisfied. This gives, via Theorem 2.6, an alternative proof for the existence of a maximizer for $u(x, U)$ with the advantage that one can apply the same proof for preference functionals with distorted beliefs, as discussed in part 1) of Remark 2.7.

Proposition 5.1. *Let the financial market satisfy $\mathcal{M}^e(S) = \{Q\}$. Then Assumption 2.5 is satisfied.*

For later use, we formulate (and prove) the following abstract version of Proposition 5.1.

Lemma 5.2. *Let (Ω, \mathcal{F}, P) and $Q \approx P$ be fixed. Let (f^n) be a sequence of nonnegative random variables with $\limsup_n E_Q[f^n] < \infty$. Let F^n denote the distribution of f^n and assume that $F^n \Rightarrow F$ for some distribution function F . Then there is \bar{f} with distribution F and $E_Q[\bar{f}] \leq \liminf_{n \rightarrow \infty} E_Q[f^n]$.*

Proof. The probability space (Ω, \mathcal{F}, P) can be decomposed into an atomic part Ω^{pa} which consists of (at most countably many) P -atoms and an atomless part $\Omega^{\text{na}} = \Omega \setminus \Omega^{\text{pa}}$ which contains no atoms. The main idea is to prove the statement on both of the parts separately. We work directly with a (re-labelled) sequence realizing the \liminf of $(E_Q[f^n])$ and we assume, without loss of generality, that $P[\Omega^{\text{na}}] > 0$ and $P[\Omega^{\text{pa}}] > 0$.

1) We start with the atomic part Ω^{pa} which consists of at most countable many atoms $\omega_1, \omega_2, \dots$, and we use the *Arzelà–Ascoli diagonalization argument* to extract a subsequence (n_k) and a limit element f^{pa} such that $f^{n_k} \rightarrow f^{\text{pa}}$ a.s. on Ω^{pa} , as follows. Let us first consider the sequence $(f^n(\omega_1))$. Since P -atoms are Q -atoms and vice versa, we find that the sequence $(f^n(\omega_1))$ is bounded above by $x/Q[\{\omega_1\}]$. Hence we can extract a converging subsequence $(f^{n_{1,\ell}}(\omega_1))$ converging to some $f^{\text{pa}}(\omega_1) \geq 0$ as

$\ell \rightarrow \infty$. Now let us do the same with the sequence $(f^{n_1, \ell}(\omega_2))_{\ell \in \mathbb{N}}$. This gives a further subsequence $(f^{n_2, \ell})_{\ell \in \mathbb{N}}$. We proceed in this way and consider the diagonal sequence $(f^{n_{\ell}, \ell})_{\ell \in \mathbb{N}}$. One easily verifies that $(f^{n_{\ell}, \ell}(\omega_i))$ converges to $f^{\text{pa}}(\omega_i)$ as $\ell \rightarrow \infty$ for every $i = 1, \dots$, which means that on Ω^{pa} the sequence $(f^{n_{\ell}, \ell})$ converges pointwise to f^{pa} as $\ell \rightarrow \infty$. In particular, this gives

$$\lim_{\ell \rightarrow \infty} P[f^{n_{\ell}, \ell} \leq a \mid \Omega^{\text{pa}}] = P[f^{\text{pa}} \leq a \mid \Omega^{\text{pa}}]. \quad (5.1)$$

2) We now consider the subsequence $(f^{n_{\ell}, \ell})_{\ell \in \mathbb{N}}$ constructed in part 1) as a sequence on the atomless probability space $(\Omega^{\text{na}}, \mathcal{F} \cap \Omega^{\text{na}}, P[\cdot \mid \Omega^{\text{na}}])$. Note that the random variable

$$\varphi := \frac{dQ}{dP} \frac{1}{E\left[\frac{dQ}{dP} \mid \Omega^{\text{na}}\right]}$$

is a strictly positive random variable on $(\Omega^{\text{na}}, \mathcal{F} \cap \Omega^{\text{na}}, P[\cdot \mid \Omega^{\text{na}}])$ with expectation 1. Moreover, we have that

$$E[\varphi f^{n_{\ell}, \ell} \mid \Omega^{\text{na}}] \leq \frac{E\left[\frac{dQ}{dP} f^{n_{\ell}, \ell} \mid \Omega^{\text{na}}\right]}{E\left[\frac{dQ}{dP} \mid \Omega^{\text{na}}\right]} \leq \frac{x}{P[\Omega^{\text{na}}] E\left[\frac{dQ}{dP} \mid \Omega^{\text{na}}\right]} =: \tilde{x},$$

where we used in the second inequality that $E_Q[f^{n_{\ell}, \ell}] \leq x$ holds for all ℓ since $f^{n_{\ell}, \ell} \in \mathcal{C}(x)$. It then follows from Lemma III.4.2 that the sequence $(f^{n_{\ell}, \ell})_{\ell \in \mathbb{N}}$ is tight (with respect to $P[\cdot \mid \Omega^{\text{na}}]$). This allows us to apply Helly's selection theorem (Billingsley [18, Theorem 6.1 and p.227]) to get a further subsequence (which by a slight abuse of notation is denoted by) $(n_{\ell})_{\ell \in \mathbb{N}}$ and a distribution function \bar{F} such that $\lim_{\ell \rightarrow \infty} P[f^{n_{\ell}} \leq a \mid \Omega^{\text{na}}] = \bar{F}(a)$ holds for all continuity points a of \bar{F} .

Since $(\Omega^{\text{na}}, \mathcal{F} \cap \Omega^{\text{na}}, P[\cdot \mid \Omega^{\text{na}}])$ is atomless, it is possible to find on it a random variable \mathcal{U} uniformly distributed on $(0, 1)$ such that $\varphi = q_{\varphi|\Omega^{\text{na}}}(\mathcal{U})$ $P[\cdot \mid \Omega^{\text{na}}]$ -a.s. (Lemma A.28 in Föllmer and Schied [43]), where we use $q_{\varphi|\Omega^{\text{na}}}$ to denote the quantile function of the distribution function $P[\varphi \leq a \mid \Omega^{\text{na}}]$. Define the random variable $f^{\text{na}} := q_{\bar{F}|\Omega^{\text{na}}}(1 - \mathcal{U})$. Since $1 - \mathcal{U}$ is uniformly distributed on $(0, 1)$, the distribution of f^{na} on $(\Omega^{\text{na}}, \mathcal{F} \cap \Omega^{\text{na}}, P[\cdot \mid \Omega^{\text{na}}])$ is \bar{F} (Lemma A.19 in [43]). With the arguments so far, we have shown that

$$\lim_{\ell \rightarrow \infty} P[f^{n_{\ell}} \leq a \mid \Omega^{\text{na}}] = P[f^{\text{na}} \leq a \mid \Omega^{\text{na}}] \quad (5.2)$$

holds for all continuity points a of \bar{F} .

We now show that

$$E_Q[f^{\text{na}} 1_{\Omega^{\text{na}}}] \leq \liminf_{\ell \rightarrow \infty} E_Q[f^{n_{\ell}} 1_{\Omega^{\text{na}}}] \quad (5.3)$$

For this, we rewrite the first expectation in terms of φ and $E[\cdot \mid \Omega^{\text{na}}]$, which allows us to rewrite it in terms of \mathcal{U} . We then combine Fatou's lemma and the

fact that weak convergence implies convergence of any quantile functions to get a first inequality. A second one follows by applying the Hardy–Littlewood inequality (Theorem A.24 in [43]). These steps together give

$$\begin{aligned}
E_Q[f^{\text{na}} 1_{\Omega^{\text{na}}}] &= E[\varphi f^{\text{na}} \mid \Omega^{\text{na}}] E\left[\frac{dQ}{dP} \mid \Omega^{\text{na}}\right] P[\Omega^{\text{na}}] \\
&= E[q_{\varphi|\Omega^{\text{na}}}(\mathcal{U}) q_{\bar{F}|\Omega^{\text{na}}}(1 - \mathcal{U}) \mid \Omega^{\text{na}}] E\left[\frac{dQ}{dP} \mid \Omega^{\text{na}}\right] P[\Omega^{\text{na}}] \\
&= \int_0^1 q_{\varphi|\Omega^{\text{na}}}(s) q_{\bar{F}|\Omega^{\text{na}}}(1 - s) ds E\left[\frac{dQ}{dP} \mid \Omega^{\text{na}}\right] P[\Omega^{\text{na}}] \\
&\leq \liminf_{\ell \rightarrow \infty} \int_0^1 q_{\varphi|\Omega^{\text{na}}}(s) q_{F^{\text{na}}|\Omega^{\text{na}}}(1 - s) ds E\left[\frac{dQ}{dP} \mid \Omega^{\text{na}}\right] P[\Omega^{\text{na}}] \\
&\leq \liminf_{\ell \rightarrow \infty} E[\varphi f^{\text{na}} \mid \Omega^{\text{na}}] E\left[\frac{dQ}{dP} \mid \Omega^{\text{na}}\right] P[\Omega^{\text{na}}] \\
&\leq \liminf_{\ell \rightarrow \infty} E_Q[f^{\text{na}} 1_{\Omega^{\text{na}}}],
\end{aligned}$$

which finishes the proof of (5.3).

3) We define a random variable f^* on $(\Omega, \mathcal{F}_T, P)$ by

$$f^* = f^{\text{na}} 1_{\Omega^{\text{na}}} + f^{\text{pa}} 1_{\Omega^{\text{pa}}}.$$

In order to determine the distribution of f^* , we decompose the distribution function of f^* into the one on the atomic part and the one on the atomless part and use (5.1) and (5.2). This gives

$$\begin{aligned}
P[f^* \leq a] &= P[f^* \leq a \mid \Omega^{\text{pa}}] P[\Omega^{\text{pa}}] + P[f^* \leq a \mid \Omega^{\text{na}}] P[\Omega^{\text{na}}] \\
&= \lim_{\ell \rightarrow \infty} P[f^{\text{na}} \leq a \mid \Omega^{\text{pa}}] P[\Omega^{\text{pa}}] + \lim_{\ell \rightarrow \infty} P[f^{\text{na}} \leq a \mid \Omega^{\text{na}}] P[\Omega^{\text{na}}] \\
&= \lim_{\ell \rightarrow \infty} P[f^{\text{na}} \leq a] \\
&= F(a)
\end{aligned}$$

for all continuity points a of F .

4) The proof is completed by showing that $f^* \in \mathcal{C}(x)$. Recall that on Ω^{pa} , the sequence $(f^{\text{na}})_{\ell \in \mathbb{N}}$ converges pointwise to f^* . Fatou's lemma therefore gives $E_Q[f^* 1_{\Omega^{\text{pa}}}] \leq \liminf_{\ell} E_Q[f^{\text{na}} 1_{\Omega^{\text{pa}}}]$. Together with (5.3), we arrive at

$$\begin{aligned}
E_Q[f^*] &= E_Q[f^* 1_{\Omega^{\text{na}}}] + E_Q[f^* 1_{\Omega^{\text{pa}}}] \\
&\leq \liminf_{\ell \rightarrow \infty} E_Q[f^{\text{na}} 1_{\Omega^{\text{na}}}] + \liminf_{\ell \rightarrow \infty} E_Q[f^{\text{na}} 1_{\Omega^{\text{pa}}}] \\
&\leq \liminf_{\ell \rightarrow \infty} E_Q[f^{\text{na}}] \leq x,
\end{aligned}$$

where we used $f^{\text{na}} \in \mathcal{C}(x)$ for all $\ell \in \mathbb{N}$ in the last step. \square

For completeness, we show how Proposition 5.1 follows from Lemma 5.2.

Proof of Proposition 5.1. Let us fix a sequence (f^n) in $\mathcal{C}(x)$, let F^n denote the distribution of f^n and assume that $F^n \Rightarrow F$ for some distribution function F . It follows from (2.6) that each f^n satisfies $E_Q[f^n] \leq x$. Lemma 5.2

therefore yields some \bar{f} with distribution F and $E_Q[\bar{f}] \leq x$. This implies, again by (2.6) because $\mathcal{M}^e = \{Q\}$, that $\bar{f} \in \mathcal{C}(x)$. \square

We conclude this section with a model where the underlying probability space $(\Omega, \mathcal{F}_T, P)$ is neither purely atomic nor atomless. We then choose the model's parameters in such a way that the relation $u(x, U) = u(x, U_c)$ is satisfied. This shows that Assumption 4.2 is not necessary for $u(x, U) = u(x, U_c)$.

Example 5.3. Consider some time horizon $T \in (0, \infty)$ and a probability space (Ω, \mathcal{F}, P) on which there is a Poisson process $(N_t)_{t \geq 0}$ with intensity γ . Let $(\mathcal{F}_t)_{0 \leq t \leq T}$ be the filtration generated by the Poisson process. We consider a (discounted) market consisting of a savings account $B \equiv 1$ and one risky stock S described by

$$dS_t = \alpha S_t dt + \sigma S_{t-} d\tilde{N}_t, \quad S_0 = 1 > 0, \sigma > -1, \sigma \neq 0, \alpha/\sigma < \gamma$$

where $\tilde{N}_t := N_t - \gamma t$ is the compensated Poisson process. The unique martingale measure is defined by $dQ/dP = e^{(\gamma - \tilde{\gamma})T} \left(\frac{\tilde{\gamma}}{\gamma}\right)^{N_T}$ for $\tilde{\gamma} := \gamma - \alpha/\sigma$. Note also that dQ/dP satisfies $(dQ/dP)^{\kappa/(\kappa-1)} \in L^1$ for $\kappa \in (0, 1)$, so the moment condition discussed in part 1) of Remark 2.7 is satisfied. This model is not atomless; it thus follows from part 1) of Remark 4.5 that Assumption 4.2 is not satisfied. However, we now show that $u(x, U) = u(x, U_c)$ is satisfied for $\alpha = 0$ and $T\gamma \geq \ln(2)$. Note first that $\alpha = 0$ implies $\tilde{\gamma} = \gamma$ and hence $dQ/dP = 1$. The statement is therefore obvious for $x \notin \{U < U_c\} := \{x \in \mathbb{R}_+ \mid U(x) < U_c(x)\}$ since $U \leq U_c$, Jensen's inequality and $P = Q$ give

$$E[U(f)] \leq E[U_c(f)] \leq U_c(E[f]) \leq U_c(x) = U(x) \quad (5.4)$$

for all $f \in \mathcal{C}(x)$. For $x \in \{U < U_c\}$, there are constants a and b with $a < x < b$, $U(a) = U_c(a)$ and $U(b) = U_c(b)$ such that U is affine on (a, b) (Lemma II.2.11). The idea is now to choose $f^* := a1_A + b1_{A^c}$ for $A \in \mathcal{F}_T$ with $P[A] = \frac{b-x}{b-a}$. Note that the choice of A is possible since \mathcal{F}_T contains only one atom $\{N_T = 0\}$ whose probability is $P[N_T = 0] = \exp(-\gamma T) \leq 0.5$; the last inequality uses the assumption $T\gamma \geq \ln(2)$. It follows from the definition of f^* that $E_Q[f^*] = E[f^*] = x$ and a slight modification of (5.4) yields optimality of f^* . This ends the example.

IV.5.2 Weakly complete models

We now consider a class of models for which the set \mathcal{M}^e , not necessarily a singleton, contains a martingale measure $Q \in \mathcal{M}^e$ whose density dQ/dP dominates all the other martingale densities stochastically in the second order (see Definition 2.55 and Proposition 2.57 in Föllmer and Schied [43] for a definition and equivalent formulations). This condition appeared in Kramkov and Sîrbu [75] and Schachermayer et al. [99]; the latter referred to it as *weak completeness assumption*.

Proposition 5.4. *Suppose that there is a measure $\hat{Q} \in \mathcal{M}^e$ such that for all $Q \in \mathcal{M}^e$, $d\hat{Q}/dP$ dominates dQ/dP stochastically in the second order, i.e.*

$$\int_0^t P\left[\frac{d\hat{Q}}{dP} \leq u\right] du \leq \int_0^t P\left[\frac{dQ}{dP} \leq u\right] du \quad \text{for each } t \geq 0. \quad (5.5)$$

Moreover, assume that the distribution of $d\hat{Q}/dP$ is continuous. Then Assumptions 2.5 and 4.2 are satisfied.

Proof. We only verify Assumption 4.2. Assumption 2.5 then follows from part 2) of Remark 4.5. So we fix a distribution F on \mathbb{R}_+ with $\text{UP}(F) < \infty$. We denote by \hat{G} the distribution of $d\hat{Q}/dP$ and define $f := q_F(1 - \hat{G}(\frac{d\hat{Q}}{dP}))$. By assumption, the distribution of $d\hat{Q}/dP$ is continuous. Hence $\hat{G}(\frac{d\hat{Q}}{dP})$ (and then also $1 - \hat{G}(\frac{d\hat{Q}}{dP})$) are uniformly distributed on $(0, 1)$ and it follows that $f \sim F$. It remains to show that $f \in \mathcal{C}(\text{UP}(F))$. For this, we note first that $d\hat{Q}/dP \in \mathcal{D}(1)$; therefore

$$\begin{aligned} \text{UP}(F) &\geq \int_0^1 q_F(s) q_{\hat{G}}(1-s) ds \\ &= \int_0^1 q_F(1-s) q_{\hat{G}}(s) ds \\ &= E_{\hat{Q}} \left[q_F \left(1 - \hat{G} \left(\frac{d\hat{Q}}{dP} \right) \right) \right] =: x, \end{aligned}$$

where in the last step we used that $q_{\hat{G}}(\hat{G}(\frac{d\hat{Q}}{dP})) = \frac{d\hat{Q}}{dP}$ a.s. The proof is therefore completed by showing that $f \in \mathcal{C}(x)$ which is, by (2.6), equivalent to showing that $E_Q[f] \leq x$ for all $Q \in \mathcal{M}^e$. We define a negative, convex and decreasing function $g : (0, \infty) \rightarrow \mathbb{R}$ by $g(x) := -\int_0^x q_F(1 - F_{\hat{G}}(y)) dy$. The stochastic dominance of \hat{Q} then gives (see, for instance, Theorem 2.57 of Föllmer and Schied [43])

$$E\left[g\left(\frac{d\hat{Q}}{dP}\right)\right] \leq E\left[g\left(\frac{dQ}{dP}\right)\right] \quad \text{for all } Q \in \mathcal{M}^e. \quad (5.6)$$

The inequality (5.6) finally implies (Theorem 5 of Rüschendorf [98])

$$E\left[g'\left(\frac{d\hat{Q}}{dP}\right)\left(\frac{dQ}{dP} - \frac{d\hat{Q}}{dP}\right)\right] \geq 0 \quad \text{for all } Q \in \mathcal{M}^e, \quad (5.7)$$

which gives $E_Q[f] \leq x$ for all $Q \in \mathcal{M}^e$ since $-g'(d\hat{Q}/dP) = f$. For the convenience of the reader, we briefly show how to deduce (5.7) from (5.6). We fix an arbitrary $Q \in \mathcal{M}^e$ and define $Q_\epsilon := \epsilon Q + (1 - \epsilon)\hat{Q}$. The function $\epsilon \mapsto g(\frac{dQ_\epsilon}{dP})$ is convex on $[0, 1]$, and so

$$Z_\epsilon := \frac{g\left(\frac{dQ_\epsilon}{dP}\right) - g\left(\frac{d\hat{Q}}{dP}\right)}{\epsilon}$$

is increasing in ϵ and decreasing to $Z_0 = g'(\frac{d\hat{Q}}{dP})(\frac{dQ}{dP} - \frac{d\hat{Q}}{dP})$ as $\epsilon \searrow 0$. Since $g \leq 0$ is convex and $dQ_\epsilon/dP \in L^1$, it follows that $g(\frac{dQ_\epsilon}{dP}) \in L^1$ for any $\epsilon \in [0, 1]$. This implies $Z_\epsilon \in L^1$ and (5.6) gives $E[Z_\epsilon] \geq 0$. Monotone convergence finally gives $Z_0 \in L^1$ and $E[Z_0] \geq 0$, which is equivalent to (5.7). \square

One example for the setup described in Proposition 5.4 is an Itô process model with deterministic coefficients as described in the next example. See also Section 4 of He and Zhou [56].

Example 5.5. We consider a probability space (Ω, \mathcal{F}, P) on which there is an n -dimensional Brownian motion $W = (W_t^1, \dots, W_t^n)_{t \geq 0}$ in the augmented filtration generated by W . We consider the asset prices $S = (S^1, \dots, S^d)$ given by

$$\frac{dS_t^i}{S_t^i} = \mu_t^i dt + \sum_{j=1}^n \sigma_t^{ij} dW_t^j, \quad S_0^i = s^i > 0$$

for $i = 1, \dots, d$ and $d \leq n$. The processes μ and σ describe the appreciation rates and volatilities of the d discounted stocks and are assumed to be progressively measurable processes valued respectively in \mathbb{R}^d and $\mathbb{R}^{d \times n}$ with $\int_0^T |\mu_t^i| dt + \int_0^T \sum_{j=1}^n |\sigma_t^{ij}|^2 dt < \infty$ a.s. for all $i = 1, \dots, d$. We assume that for all $t \in [0, T]$, the matrix σ_t is of full rank equal to d . The square $d \times d$ matrix $\sigma_t \sigma_t^{\text{tr}}$ is thus invertible, and we define the progressively measurable process γ valued in \mathbb{R}^n by

$$\gamma_t = \sigma_t^{\text{tr}} (\sigma_t \sigma_t^{\text{tr}})^{-1} \mu_t, \quad 0 \leq t \leq T.$$

Moreover, we assume that $0 < \int_0^T \gamma_t^2 dt < \infty$ a.s. and that γ is *deterministic*.

This market satisfies condition (2.3) and it can be incomplete (since $d < n$ is allowed). It is shown in the second part of the proof of Theorem 6.6.4 of Karatzas and Shreve [71] (starting from equation (6.34) there) that $d\hat{Q}/dP := \exp(-\int_0^T \gamma dW - \int_0^T \gamma^2 dt)$ satisfies $E[g(d\hat{Q}/dP)] \leq E[g(dQ/dP)]$ for every convex and decreasing function $g : (0, 1) \rightarrow \mathbb{R}$ and every $Q \in \mathcal{M}^e$. This is equivalent to (5.5) by Theorem 2.57 in Föllmer and Schied [43]. The distribution of $d\hat{Q}/dP$ is continuous since the distribution of $\int \gamma dW$ is continuous; in fact, the latter is normal with mean 0 and variance $\int_0^T \gamma_t^2 dt$. Note also that $d\hat{Q}/dP$ satisfies $(d\hat{Q}/dP)^{\kappa/(\kappa-1)} \in L^1$ for $\kappa \in (0, 1)$, so the moment condition discussed in part 1) of Remark 2.7 is satisfied.

Note that in Example 5.5, only the market price of risk γ is deterministic; drift and volatility may be stochastic. The following special case of Example 5.5 with stochastic drift and volatility appears in German [47].

Example 5.6. Set $d = 1$ and $n = 2$ and fix a process V defined by $dV_t = \eta(t, V_t) + \nu(t, V_t) d\tilde{W}_t$ for $\tilde{W}_t := \rho W_t^1 + \sqrt{1 - \rho^2} W_t^2$, $0 < \rho < 1$ and $\eta(t, V_t)$ and $\nu(t, V_t)$ such that V is strictly positive and well defined. We now define $\mu_t^1 := \gamma \sqrt{V_t}$, $\sigma_t^1 = \sqrt{V_t}$ and $\sigma^2 = 0$ for $\gamma \in \mathbb{R} \setminus \{0\}$.

The next example is discussed in detail in Kramkov and Sîrbu [75].

Example 5.7. Let $W = (W_t)_{0 \leq t \leq T}$ and $B = (B_t)_{0 \leq t \leq T}$ be two independent Brownian motions on a filtered probability space $(\Omega, \mathcal{F}, \mathcal{F}_{0 \leq t \leq T}, P)$, where the filtration is generated by W and B . There is one traded asset S^1 described by

$$dS_t^1 = S_t^1(\mu dt + \sigma dW_t)$$

and one nontraded, but observable asset S^2 described by

$$dS_t^2 = S_t^2(v dt + \eta(\rho dW_t + \sqrt{1 - \rho^2} dB_t))$$

for constants $v \in \mathbb{R}$, $\mu \in \mathbb{R} \setminus \{0\}$, $\sigma > 0$, $\eta > 0$ and $0 < \rho < 1$. It is shown in Section 7 of Kramkov and Sîrbu [75] that for all $Q \in \mathcal{M}^e$, the martingale density $d\hat{Q}/dP := \mathcal{E}(-\gamma W)_T$ for $\gamma = \mu/\sigma$ dominates dQ/dP stochastically in the second order. The distribution of $d\hat{Q}/dP$ is continuous.

The three examples mentioned in this section so far explicitly exclude the case $\gamma = 0$, in which the measure P is already a martingale measure. In this case, the fluctuation of the price already follows a martingale which can be seen as a strong version of the efficient market hypothesis in the sense that there is no self-financing strategy with a positive expected gain; see Remark 5.15 of Föllmer and Schied [43]. In that case, the density of the measure P (i.e. the constant 1) satisfies $E[g(dQ/dP)] \geq g(1)$ for any convex function $g : (0, 1) \rightarrow \mathbb{R}$ and any other martingale measure $Q \in \mathcal{M}^e$ by Jensen's inequality. The density of the measure P (i.e. the constant 1) therefore satisfies (5.5) by Theorem 2.57 in [43]. However, its distribution is obviously not continuous. In the next example, we show that Assumption 4.2 (and therefore also Assumption 2.5) are also satisfied in the case $\gamma = 0$. We focus on the one-dimensional case and allow the volatility coefficient to be more general. The extension to the multidimensional case is straightforward.

Example 5.8. We consider the stochastic volatility model

$$\begin{aligned} dS_t &= \sigma(t, S_t, V_t) S_t dW_t^1, \\ dV_t &= \eta(t, S_t, V_t) + \nu(t, S_t, V_t) dW_t^2, \end{aligned} \tag{5.8}$$

where W^1 and W^2 are standard Brownian motions and σ , η and ν are supposed to be such that (5.8) has a unique strong solution and that σ is strictly positive. As filtration, we take the P -augmentation of the filtration generated by W^1 and W^2 . One explicit example is to choose $\sigma(t, S_t, V_t) = \sqrt{V_t}$, $\eta(t, S_t, V_t) = \kappa(\delta - V_t)$ and $\nu(t, S_t, V_t) = \xi\sqrt{V_t}$ for $2\kappa\delta \geq \xi^2$, which describes the Heston model satisfying the Feller condition.

In order to verify Assumption 4.2 (and therefore also Assumption 2.5), we fix a distribution F on \mathbb{R}_+ with $\text{UP}(F) < \infty$. We denote by G the (normal) distribution of W_T^1 . The function G is continuous, hence $G(W_T^1)$ is

uniformly distributed on $(0, 1)$, and it follows that $f := q_F(G(W_T^1)) \sim F$. It remains to show that $f \in \mathcal{C}(\text{UP}(F))$. Since $P \in \mathcal{M}^e$, we find that

$$x := E[q_F(G(W_T^1))] = \int_0^1 q_F(s) ds = \int_0^1 q_F(s) q_1(1-s) ds \leq \text{UP}(F) < \infty.$$

The martingale representation property therefore gives a process $\tilde{\vartheta}$ such that $0 \leq q_F(G(W_T^1)) = E[q_F(G(W_T^1))] + \int_0^T \tilde{\vartheta}_t dW_t^1$. Defining $\vartheta_t := \tilde{\vartheta}_t / (S_t^1 \sigma_t)$, we finally obtain $f = x + \int_0^T \vartheta_t dS_t^1$ which shows that $f \in \mathcal{C}(x) \subseteq \mathcal{C}(\text{UP}(F))$.

Remark 5.9. 1) He and Zhou [56] solve directly the portfolio optimization problem for the preference functional (2.14) with distortion in the model described in Example 5.5. In contrast to their proof, we use the properties of the model to show Assumption 4.2 which, by part 2) of Remark 4.5, implies Assumption 2.5 and which eventually gives the existence of a maximizer via Theorem 2.6 and part 1) of Remark 2.7.

2) Examples 5.5–5.8 all belong to the class of *weakly complete* financial markets introduced in Schachermayer et al. [99]. Although all examples come to the same conclusion that Assumption 4.2 is satisfied, the argument to verify Assumption 4.2 in the Examples 5.5–5.7 differs from the argument in Example 5.8. It would be interesting to see a unified argument for all *weakly complete* financial markets. \diamond

IV.5.3 Black–Scholes model in a random environment

We finally discuss another class of models where Assumption 2.5 is satisfied but Assumption 4.2 need not be. We consider a probability space (Ω, \mathcal{F}, P) on which there are a Brownian motion $W = (W_t)_{0 \leq t \leq T}$ and a random variable η , independent of W , taking the finitely many values $0, \dots, N$ with probabilities $p_i \geq 0$ and $\sum_{i=0}^N p_i = 1$. The filtration will be the one generated by W and augmented at time 0 by the complete knowledge of η . We define the price process S by

$$\frac{dS_t}{S_t} = \sigma_t(\eta)(\gamma_t(\eta)dt + dW_t), \quad S_0 = s_0 > 0, \quad (5.9)$$

where the market price of risk $\gamma_t(i)$, $i = 0, \dots, N$ and the volatility $\sigma_t(i) > 0$, $i = 1, \dots, N$ are deterministic and bounded functions. Moreover, we assume that $\sigma_t(0) = 0$. This market is arbitrage-free and there is more than one martingale measure.

The process S models the following situation: At time 0, we roll a dice with $N+1$ possible outcomes and according to the result, we start a geometric Brownian motion with mean $\gamma_t(i)\sigma_t(i)$ and volatility $\sigma_t(i)$. With probability p_0 , the price process remains constant on $[0, T]$.

Proposition 5.10. *Suppose that S is defined by (5.9). Then Assumption 2.5 is satisfied.*

Proof. Let us fix a sequence (f^n) in $\mathcal{C}(x)$ such that the sequence (F^n) of distribution functions of (f^n) converges weakly to \bar{F} for some distribution function \bar{F} . The goal is to construct a payoff $\bar{f} \in \mathcal{C}(x)$ with distribution \bar{F} . The method to achieve this is to decompose the probability space $(\Omega, \mathcal{F}_T, P)$ into $N+1$ parts $(\Omega \cap \{\eta = i\}, \mathcal{F}_T \cap \{\eta = i\}, P[\cdot | \eta = i])$ and to argue on each part separately. Conditionally on $\{\eta = i\}$ and $i = 1, \dots, N$, the market is complete. Thus for every $Q \in \mathcal{M}^e$, the conditional expectation $E_Q[f | \eta = i]$ is independent of $Q \in \mathcal{M}^e$ and given by

$$E_Q[f | \eta = i] = E[Z_T^i f | \eta = i], \quad (5.10)$$

for $Z_t^i := \mathcal{E}(-\int_0^t \gamma_s(i) dW_s)_t$ and $i = 1, \dots, N$. On each part, we then apply Lemma 5.2.

1) We first assume that f^n is the final value X_T^n of a nonnegative process $X^n \in \mathcal{X}(x)$ and consider the sequence $(f^n)_{n \in \mathbb{N}}$ as a sequence on the probability space $(\Omega \cap \{\eta = 1\}, \mathcal{F}_T \cap \{\eta = 1\}, P[\cdot | \eta = 1])$. We define

$$\varphi := \frac{Z_T^1}{E[Z_T^1 | \eta = 1]}$$

and observe that φ is a strictly positive random variable with expectation 1 on $(\Omega \cap \{\eta = 1\}, \mathcal{F}_T \cap \{\eta = 1\}, P[\cdot | \eta = 1])$. Moreover, we have that

$$E[\varphi f^n | \eta = 1] = \frac{E_Q[f | \eta = 1]}{E[Z_T^1 | \eta = 1]} \leq \frac{x}{P[\eta = 1]E[Z_T^1 | \eta = 1]} =: \tilde{x},$$

where we used that $E_Q[f^n] \leq x$ holds for all n since $f^n \in \mathcal{C}(x)$. It then follows from Lemma III.4.2 that the sequence (f^n) is tight (with respect to the measure $P[\cdot | \eta = 1]$). This allows us to apply Helly's selection theorem (Billingsley [18, Theorem 6.1 and p.227]) to get a subsequence $(f^{n_{i,\ell}})_{\ell \in \mathbb{N}}$ and a distribution function \bar{F}^1 such that $\lim_{\ell \rightarrow \infty} P[f^{n_{i,\ell}} \leq a | \eta = 1] = \bar{F}^1(a)$ holds for all continuity points a of \bar{F}^1 . We now apply Lemma 5.2 (with the measure defined by φ) to construct \bar{f}^1 on $(\Omega \cap \{\eta = 1\}, \mathcal{F}_T \cap \{\eta = 1\}, P[\cdot | \eta = 1])$ satisfying $\bar{f}^1 \sim \bar{F}^1$ and

$$E[\varphi \bar{f}^1 | \eta = 1] \leq \liminf_{\ell \rightarrow \infty} E[\varphi f^{n_{i,\ell}} | \eta = 1].$$

The latter inequality implies, via equality (5.10) and the definition of φ , that $E_Q[\bar{f}^1 | \eta = 1] \leq \liminf_{\ell} E_Q[f^{n_{i,\ell}} | \eta = 1]$ holds for each $Q \in \mathcal{M}^e$.

2) We now iteratively apply the same arguments again for $(f^{n_{i,\ell}})_{\ell \in \mathbb{N}}$ as a sequence on $(\Omega \cap \{\eta = i+1\}, \mathcal{F}_T \cap \{\eta = i+1\}, P[\cdot | \eta = i+1])$ and $i = 1, \dots, N-1$. This allows us for all $i = 1, \dots, N$ to construct \bar{f}^i on $(\Omega \cap \{\eta = i\}, \mathcal{F}_T \cap \{\eta = i\}, P[\cdot | \eta = i])$ satisfying

$$\lim_{\ell \rightarrow \infty} P[f^{n_{i,\ell}} \leq a | \eta = i] = P[\bar{f}^i \leq a | \eta = i] \quad (5.11)$$

and

$$E_Q[\bar{f}^i \mid \eta = i] \leq \liminf_{\ell \rightarrow \infty} E_Q[f^{n_{i,\ell}} \mid \eta = i]. \quad (5.12)$$

For $\eta = 0$, recall that f^n is by the assumption in 1) the final position X_T^n of a nonnegative value process $X^n \in \mathcal{X}(x)$. Since S remains constant on $\{\eta = 0\}$, it follows that $f^n = x$ holds there. This is the only place where $f^n = X_T^n$ is being used. We thus define $\bar{f}^0 = x$ and observe that \bar{f}^0 trivially satisfies (5.11) and (5.12) for $\eta = 0$ and the sequence $(n_{0,\ell})_{\ell \in \mathbb{N}}$ defined by $n_{0,\ell} := n_{N,\ell}$ for all $\ell \in \mathbb{N}$.

3) The candidate \bar{f} on $(\Omega, \mathcal{F}_T, P)$ can now be defined by

$$\bar{f} = \sum_{i=0}^N \bar{f}^i 1_{\{\eta=i\}}.$$

In order to verify that the distribution of \bar{f} is \bar{F} , we decompose the distribution function \bar{F} into $N+1$ conditional distribution functions, use (5.12) and recall that the sequence $(n_{N,\ell})_{\ell \in \mathbb{N}}$ is a subsequence of any other sequence $(n_{i,\ell})_{\ell \in \mathbb{N}}$ for $i = 0, \dots, N$. This gives

$$\begin{aligned} P[\bar{f} \leq a] &= \sum_{i=0}^N P[\bar{f}^i \leq a \mid \eta = i] P[\eta = i] \\ &= \sum_{i=0}^N \lim_{\ell \rightarrow \infty} P[f^{n_{i,\ell}} \leq a \mid \eta = i] P[\eta = i] \\ &= \lim_{\ell \rightarrow \infty} P[f^{n_{N,\ell}} \leq a] \\ &= \bar{F}(a) \end{aligned}$$

for all points a that are continuity points of $P[\bar{f}^i \leq a \mid \eta = i]$ for each $i = 0, \dots, N$ as well as of \bar{F} . But this means that the right-continuous functions $P[\bar{f} \leq a]$ and $\bar{F}(a)$ coincide Lebesgue-a.e., so they are in fact identical.

4) In order to show $\bar{f} \in \mathcal{C}(x)$, we fix some $Q \in \mathcal{M}^e$, rewrite $E_Q[\bar{f}]$ in conditional terms and use (5.12) to obtain

$$\begin{aligned} E_Q[\bar{f}] &= \sum_{i=0}^N E_Q[\bar{f} \mid \eta = i] Q[\eta = i] \\ &\leq \sum_{i=0}^N \liminf_{\ell \rightarrow \infty} E_Q[f^{n_{i,\ell}} \mid \eta = i] Q[\eta = i] \\ &\leq \liminf_{\ell \rightarrow \infty} E_Q[f^{n_{N,\ell}}] \leq x, \end{aligned}$$

where we used $f^{n_{N,\ell}} \in \mathcal{C}(x)$ for all $\ell \in \mathbb{N}$.

5) So far, we have shown the claim for a sequence (f^n) where f^n is the final position $f^n = X_T^n$ of a nonnegative value process $X^n \in \mathcal{X}(x)$. For an arbitrary sequence (f^n) with limit distribution \bar{F} and $f^n \in \mathcal{C}(x)$, there is, by the definition of $\mathcal{C}(x)$, a value process $X^n \in \mathcal{X}(x)$ satisfying $f^n \leq X_T^n$. Applying the above argument for the sequence (X_T^n) yields a random variable $\tilde{f} \in \mathcal{C}(x)$ with distribution \tilde{F} such that we have $X_T^n \Rightarrow \tilde{f}$ along a subsequence. Weak convergence implies a.e. convergence of the quantiles, so we have $q_{X_T^n}(t) \rightarrow q_{\tilde{f}}(t)$ for a.e. $t \in (0, 1)$ and $q_{f^n}(t) \rightarrow q_{\bar{F}}(t)$ for a.e. $t \in (0, 1)$. The relation $f^n \leq X_T^n$ gives $q_{f^n}(t) \leq q_{X_T^n}(t)$ and this yields then $q_{\bar{F}}(t) \leq q_{\tilde{f}}(t)$. Finally, we choose a uniformly distributed random variable \mathcal{U} satisfying $\tilde{f} = q_{\tilde{f}}(\mathcal{U})$ which is possible since $(\Omega, \mathcal{F}_T, P)$ is atomless. We define $f := q_{\bar{F}}(\mathcal{U})$ and conclude that $f \in \mathcal{C}(x)$ since $f = q_{\bar{F}}(\mathcal{U}) \leq q_{\tilde{f}}(\mathcal{U}) = \tilde{f} \in \mathcal{C}(x)$. \square

We conclude this section with an example which shows that $u(x, U)$ may be strictly smaller than $u(x, U_c)$, even if the underlying probability space $(\Omega, \mathcal{F}_T, P)$ is atomless.

Example 5.11. Let $N = 1$, $p_0 = p_1 = 1/2$, $\gamma_t(i) = 1$ for $i = 0, 1$, $\sigma_t(1) = 1$, and $s_0 = 1$. The asset S is therefore defined by

$$S_t = 1_{\{\eta=0\}} + \exp(W_t + t/2)1_{\{\eta=1\}}.$$

On the set $\{\eta = 0\}$, the price process S remains constant. So any self-financing trading strategy ends up with the initial capital there. This allows us to write $E[U(X_T)] = P[\eta = 0]U(x) + E[U(X_T)1_{\{\eta=1\}}]$ for any value process $X \in \mathcal{X}(x)$. Taking the supremum over $f \in \mathcal{C}(x)$ yields

$$u(x, U) = \frac{1}{2}U(x) + \sup_{f \in \mathcal{C}(x)} E[U(f)1_{\{\eta=1\}}].$$

The analogous argument for U_c gives

$$u(x, U_c) = \frac{1}{2}U_c(x) + \sup_{f \in \mathcal{C}(x)} E[U_c(f)1_{\{\eta=1\}}].$$

An initial capital $x \in \{U < U_c\}$ satisfies $U(x) < U_c(x)$ and this finally gives $u(x, U) < u(x, U_c)$ because of $U \leq U_c$. Theorem 4.6 therefore implies that Assumption 4.2 is not satisfied in this example.

Remark 5.12. A similar example (with $p_0 = 0$) appears in Bichuch and Sturm [17]. They impose additional assumptions on (γ_t) such that the distributions of all martingale densities dQ/dP have a continuous distribution and then proceed as described in part 1) of Remark 4.8 above. \diamond

IV.6 Conclusion

In this chapter, we study the problem of non-concave utility maximization from terminal wealth in a general financial market. Our general formulation allows us to analyze the problem in a unified way and enables us to work out more systematically the impact of the non-concavity in the utility function on the portfolio choice. This chapter is clearly just a first step, and there are many open and interesting problems. To name a few, we could mention:

- (i) Are there general and easily verifiable structural assumptions on the price process S such that Assumption 2.5 is satisfied?
- (ii) Are there a model and a non-concave utility U such that there is a maximizer for $u(x, U_c)$ but none for $u(x, U)$?
- (iii) Does the first order condition for optimality (proved for the case of finite probability spaces in Proposition 3.2) also hold for the general setting?
- (iv) Are there (easily verifiable) necessary *and* sufficient conditions such that $u(x, U) = u(x, U_c)$ holds for any $x > 0$?

The present chapter gives answers to the above questions for a particular choice of the model. It will be interesting to see more answers.

Chapter V

Implications for a financial market equilibrium: The pricing kernel puzzle

This chapter is a modified version of the article [58]. We study implications of behavioural portfolio selection for a general financial market equilibrium. We focus on the relation between the pricing kernel(s)¹ (or pricing densities) and the aggregate endowment.

V.1 Introduction

A pricing kernel is the Radon–Nikodým derivative of a risk-neutral probability (or martingale) measure with respect to the objective probability measure. In equilibrium models, this is an important quantity since it provides the connection between asset prices and fundamental economic principles such as the scarcity of endowment and the decreasing marginal utility of wealth. In standard models of financial economics, risk-averse agents with correct beliefs trade in a complete market. In these models, the pricing kernel is a positive and decreasing function of the aggregate endowment. However, some empirical studies claim to show that this function is positive and generally decreasing² but also has increasing parts. The latter is known as the *pricing kernel puzzle*.

In order to explain the pricing kernel puzzle (in the sense that one can generate examples with a (locally) increasing relation between the pricing kernel and the aggregate endowment), one needs to relax at least one of the assumptions of the standard model. This, in principle, leads to (at least)

¹In this chapter, it is more convenient to refer to dQ/dP as the pricing kernel rather than pricing density in order to be in line with the literature.

²A sufficiently smooth function h on \mathbb{R}_+ is said to be *generally decreasing* if $\lim_{x \searrow 0} h(x) \geq \lim_{x \nearrow \infty} h(x)$.

three routes for explaining the pricing kernel puzzle. The first task is to understand the resulting economic mechanisms that lead to an increasing relation between the pricing kernel and the aggregate endowment. The second task is then to understand whether the relaxations of the standard model lead to pricing kernels that are consistent with the empirical findings.

While the main branch of the pricing kernel literature is focusing on the second task, we are also interested in the first one. For a good understanding of the pricing kernel, it is necessary to understand the mechanism of how different deviations from the standard assumptions influence the pricing kernel. For this, we consider a simple and unifying setup where possible deviations from the standard model can be analyzed and compared. Given the three main assumptions of the standard model, we explore three possible reasons for the pricing kernel puzzle: non-concave utility functions, distorted (and incorrect) beliefs, and incomplete markets.

Non-concave utility functions refer to risk-seeking behaviour. While risk aversion is a standard assumption in finance, there is considerable empirical evidence that agents might show risk aversion for some ranges of returns and risk-seeking behaviour for others (examples can be found in Kahneman and Tversky [69]). Formally, risk-seeking behaviour is described by a partially convex utility function, and the convexity in the utility can, in principle, lead to a non-decreasing relation between the pricing kernel and the aggregate endowment. However, we show in Theorem 6.2 that the agent chooses a final position which for at most one state takes values in the area where his utility is strictly convex. This allows us to conclude that non-concave utility functions as an explanation is rather a pathological phenomenon of theoretical interest and not a reasonable explanation for the empirical findings.

Incorrect beliefs refer to situations where the subjective belief(s) of the agent(s) do not coincide with the objective probability measure. Reasons for such settings are misestimation (by the agent or the statistician), heterogeneous beliefs, ambiguity aversion or behavioural biases such as distortion. In such models, a non-decreasing relation between the pricing kernel and the aggregate endowment can be viewed as a difference between the belief that determines the equilibrium and the estimated probability. Taking a specific model for beliefs or distortions (depending on some parameters), the pricing kernel is a parametric function of the aggregate endowment (depending on the model parameters) and one can check whether reasonable choices of the parameters lead to reasonable pricing kernels. We find that even with unrealistic parameter values, distorted and misestimated beliefs in isolation are not sufficient to explain the empirical results. A combination of misestimated and distorted beliefs can reproduce the empirical results, but the implied parameters are not realistic.

A third relaxation of the standard model is to consider an *incomplete market*. While the market spanned by an index (e.g. S&P 500) and its options is usually assumed to be complete, background risk (in a sense to be

made precise later) and market frictions can lead to situations where the agents face a portfolio optimization problem in incomplete markets. In such a setting, the pricing kernel that is relevant for the agents is not necessarily unique and the final positions chosen by different agents correspond to different pricing kernels. The aggregation becomes sometimes impossible, the representative agent may fail to exist and the decreasing relation between the pricing kernels and the aggregate endowment may be violated. In particular the (heterogeneous) background risk (such as labour risk or housing risk not captured by an index) leads to very flexible pricing kernels. The challenge is then to give plausible restrictions such that the pricing kernel is a generally decreasing function of the aggregate endowment with increasing areas.

In the literature, one finds many (specific) models providing many possible explanations for the pricing kernel puzzle; see Section V.2 for a detailed review. One goal of this chapter is to analyze and classify these explanations more systematically and to explain the economic mechanism leading to a non-decreasing relation between the aggregate endowment and the pricing kernel(s). In particular, this includes a detailed analysis of the pricing kernel(s) in incomplete markets as suggested in the conclusion of Ziegler [111].

The chapter is organized as follows. In Section V.2, we review the literature on the pricing kernel puzzle. In Section V.3, we introduce the model and we define our notion of a financial market equilibrium. Section V.4 considers the case of risk-averse agents having common beliefs in a complete market. The boundary behaviour of the pricing kernel is analyzed in Section V.5. Section V.6 is devoted to the study of the case of partially risk-seeking agents. Section V.7 provides a detailed exposition of the case where risk-averse agents have incorrect beliefs. In Section V.8, we look more closely at the pricing kernels in incomplete markets. Finally, Section V.9 summarizes the main results. In an effort to keep clear the main lines of the argument, some of the more technical mathematical calculations are placed in appendices. For standard results in financial economics, the corresponding results in Magill and Quinzii [82] are cited as one possible reference.

V.2 Related literature

In the literature, the pricing kernel(s) is mainly analyzed from the econometric viewpoint. The main part of the literature is focusing on a setting where the aggregate endowment is equal to (a multiple of) an index (such as the S&P 500 or the DAX; plus a deterministic constant) and where the market spanned by the index and the options written on the index forms a complete market (we will explain this type of setting in detail in Remarks 3.2 and 4.5). For this type of setting, researchers have taken great interest in estimating the pricing kernel. One often-used approach relies on a model of a representative agent in which the (unique) pricing kernel is a paramet-

ric function of the aggregate endowment. Market data are then used to estimate the parameters. Two among numerous examples are Brown and Gibbons [25] and Hansen and Singleton [52]. Both use a pricing kernel implied by a power utility. Due to the parametric form, the pricing kernel is necessarily a decreasing function of the aggregate endowment. Another approach is based on the no-arbitrage principle. While the techniques of this method have become more and more sophisticated, the basic approach has remained the same. Along the lines of Breeden and Litzenberger [23], option data are used to estimate the martingale measure (or risk-neutral distribution). Other methods (and historical data) is used to determine an estimated probability measure which is seen as a proxy for the objective probability measure. Some examples are Jackwerth and Rubinstein [60], Aït-Sahalia and Lo [1], Jackwerth [59], Aït-Sahalia and Lo [2], Brown and Jackwerth [26], Rosenberg and Engle [97], Yatchew and Härdle [110] and Barone-Adesi et al. [8]. The most robust observation in that part of the literature is that the pricing kernel is a generally decreasing function of the underlying index. Often, but not always, there is an interval, usually in the area of zero return of the index, where the pricing kernel is an increasing function of the index. Note that this area is highly relevant since most of the monthly returns of the indices are between -4% and $+2\%$.

In the papers mentioned above, one finds many hypotheses which are invoked to explain the pricing kernel puzzle. Many empirical studies (see, for instance, Rosenberg and Engle [97], Detlefsen et al. [40] and Golubev et al. [49]) consider the pricing kernel as marginal utility of the representative agent and state that the non-decreasing relation between the pricing kernel and the aggregate endowment is evidence for *risk-seeking behaviour* of the representative agent. Assuming complete markets, this argument is questionable since an agent with such a utility chooses a final position which is different from the index (see Theorem 6.2).

Chabi-Yo et al. [33] and Benzoni et al. [12] consider a representative agent having a state-dependent utility in an *incomplete* market specified via a latent state variable. In both cases, the models are also calibrated to the empirical data and can capture the stylized facts of the data. Grith et al. [51] consider agents with a state-dependent utility in a complete market setup. The state-dependence is specified in such a way that the risk aversion depends on the wealth level. In all these articles, the mechanism leading to a non-decreasing relation between the pricing kernel and the aggregate endowment is (somehow) hidden in advanced models. The present chapter works out explicitly the fundamental economic principle leading to this non-decreasing relation.

Explanations regarding *incorrect beliefs* appear in various forms in the literature. Shefrin [101] explains the puzzle with heterogeneous beliefs. Jack-

werth [59] examines whether a *Peso problem*³ could explain the puzzle and concludes that this is not the case. Ziegler [111] considers risk-averse agents in complete markets and compares different explanations regarding heterogeneous and wrong beliefs. He comes to the conclusion that the degree of pessimism needed is implausibly high. More recently, Dierkes [41] and Polkovnichenko and Zhao [86] analyze the pricing kernel in a setup with distorted beliefs but do not connect this systematically to the pricing kernel puzzle. Gollier [48] uses ambiguity aversion to explain the puzzle, but he does not calibrate the model to the empirical data. On an abstract level, all these explanations are related and our analysis provides a unified perspective.

Given the huge variety of all these efforts, we thought that it is time to give a unifying and simple framework of a financial market in which all of these hypotheses can be analyzed and compared. This simple setup allows us to describe in a transparent way the fundamental economic problem as well as the principles leading to a non-decreasing relation between the pricing kernel and the aggregate endowment. But of course, there is a price to pay; there are also some explanations which cannot be classified in our framework. Most important seem to be statistical problems and challenges in the estimation procedures; see, for instance, Jackwerth [59] for several useful comments.

V.3 Setup

We consider a one-period exchange economy. Let $\Omega = \{\omega_1, \dots, \omega_N\}$ denote the states of nature at time 1. The set $\mathcal{F} = 2^\Omega$ is the power set on Ω , i.e., the set of all possible events arising from Ω . Uncertainty is modeled by the probability space (Ω, \mathcal{F}, P) , where the objective probability P on Ω satisfies $P[\{\omega_n\}] > 0$ for $n = 1, \dots, N$, i.e., every state of the world has a strictly positive probability to occur.

There are $d + 1$ assets, whose payoffs at date $t = 1$ are described by $S_1^i \in \mathbb{R}^N$. The asset 0 is the risk-free asset with payoff $S_1^0 = 1$. The price of the i -th asset at date $t = 0$ is denoted by S_0^i . The risk-free asset supply is unlimited and the price S_0^0 is exogenously given by 1. The latter assumption does not restrict the generality of the model as we always may choose the bond as numeraire. In other words, the payoffs are already discounted. The prices of the other assets are endogenously derived by demand and supply. The marketed subspace \mathcal{X} is the span of $(S_1^i)_{i=0,1,\dots,d}$. Without loss of generality, we assume that no asset is redundant, i.e., $\dim(\mathcal{X}) = d + 1$, and that $d + 1 \leq N$ holds. The market is called complete if $d + 1 = N$ holds.

We consider a finite set \mathcal{I} of agents. Agent j has a stochastic income $W_j \in \mathbb{R}_+^N$ at date 1. This summarizes the initial capital and the value

³A *Peso problem* arises when the estimated probability is not an accurate estimate of the objective probability (for instance, since the sample is too short or some events are unlikely to be observed) but the agents are cognizant of the objective probability.

of the initial holdings in stocks at date 1. The sum over W_j is called the *aggregate endowment*. The variable $\vartheta_j = (\vartheta_j^0, \dots, \vartheta_j^d) \in \mathbb{R}^{d+1}$ denotes the j -th agent's portfolio giving the number of units of each of the $d+1$ securities purchased (if $\vartheta_j^i > 0$) or sold (if $\vartheta_j^i < 0$) at time 0 by agent j . Buying and selling these $d+1$ securities is the only trading opportunity available to agent j . Thus, given the available securities, agent j can generate any payoff $f = W_j + \sum_{i=0}^d S_1^i \vartheta_j^i$, where ϑ_j satisfies the *budget restriction* $\sum_{i=0}^d S_0^i \vartheta_j^i \leq 0$. Moreover, we assume that the resulting income must be non-negative in all states of nature, i.e., $f \geq 0$. The subset of payoffs in \mathcal{X} that are positive and can be generated by trading for agent j is denoted by $\mathcal{C}_j(S_0)$, i.e.,

$$\mathcal{C}_j(S_0) := \left\{ f \in \mathbb{R}_+^N \mid f \leq W_j + \sum_{i=0}^d S_1^i \vartheta_j^i \text{ for some } \vartheta_j \in \mathbb{R}^{d+1} \text{ s.t. } \sum_{i=0}^d S_0^i \vartheta_j^i \leq 0 \right\}.$$

Every agent has his own (subjective) *belief* about the future; the belief of agent j is represented by a probability measure $P_j \approx P$. The preferences of agent j are described by a strictly increasing⁴ functional $V_j : \mathcal{X} \rightarrow \mathbb{R}$. This functional summarizes the utility function U_j as well as the belief P_j of the agent j . We often use the expected utility functional $V_j(f) := E_{P_j}[U_j(f)]$, where E_{P_j} denotes the expectation with respect to P_j . We will explicitly define functionals in the next sections. In order to optimize the preference functional, agents may want to buy and sell assets. An allocation $(f_j)_{j \in \mathcal{I}}$ is called *feasible* if the resulting total demand matches the overall supply (i.e., the sum over the initial holdings in the stocks). Formally, this means that the *market clearing condition*

$$\sum_{j \in \mathcal{I}} \vartheta_j = 0 \tag{3.1}$$

has to be satisfied. Note that the market clearing conditions for the financial contracts imply that the final positions $(f_j)_{j \in \mathcal{I}}$ satisfy

$$\sum_{j \in \mathcal{I}} f_j = \sum_{j \in \mathcal{I}} W_j. \tag{3.2}$$

In a financial market equilibrium, the prices of assets are derived in such a way that the resulting total demand matches the overall supply.

Definition 3.1. A price vector $S_0 = (1, S_0^1, \dots, S_0^d)$ together with a feasible allocation $(f_j)_{j \in \mathcal{I}}$ is called a *financial market equilibrium* if each f_j maximizes the functional V_j over all $f \in \mathcal{C}_j(S_0)$.

Since the preference functional is strictly increasing, the agents would exploit arbitrage opportunities in the sense of a sure gain without any risk.

⁴A functional V is called *strictly increasing* if $f_1(\omega) \geq f_2(\omega)$ for all $\omega \in \Omega$ and $f_1(\omega_n) > f_2(\omega_n)$ for at least one $\omega_n \in \Omega$ implies $V(f_1) > V(f_2)$.

This means that if there were such an opportunity, every agent would rush to exploit it and so competition will make it disappear very quickly. Thus, we conclude that the condition

$$\left\{ f \in \mathbb{R}_+^N \mid f \leq \sum_{i=0}^d S_1^i \vartheta_j^i \text{ for } \vartheta_j \in \mathbb{R}^{d+1} \text{ s.t. } \sum_{i=0}^d S_0^i \vartheta_j^i \leq 0 \right\} = \{0\}$$

is satisfied in equilibrium. This implies (Theorem 9.3 in Magill and Quinzii [82]) the existence of a martingale measure $Q \approx P$ such that $S_0^i = E_Q[S_1^i]$ holds for all assets i . The Q -probabilities $Q[\{\omega_n\}]$, $n = 1, \dots, N$ are also called *state prices* or *risk-neutral probabilities*. A *pricing kernel* is the Radon–Nikodým derivative of an equivalent martingale measure Q with respect to P , denoted by dQ/dP . Note that in the present setting of this chapter, it follows that

$$\frac{dQ}{dP}(\omega_n) = \frac{Q[\{\omega_n\}]}{P[\{\omega_n\}]} \quad \text{for } n = 1, \dots, N.$$

In the next sections, we sometimes consider other probability measures $R \approx P$. One example is the belief $P_j \approx P$ (of agent j). Another is the estimated probability measure $\hat{P} \approx P$ (of the statistician); this object will be defined and explained in detail in the next section. In order to distinguish different cases, we use the following notion. For a probability measure $R \approx P$, a *pricing kernel with respect to R* is the Radon–Nikodým derivative of a martingale measure Q with respect to R , denoted by dQ/dR . With this notion, a pricing kernel can equivalently be defined as a pricing kernel with respect to the objective probability measure P . Note that a pricing kernel (with respect to any R) is not unique if the market is incomplete; each martingale measure Q defines one.

Remark 3.2. 1) In order to illustrate our abstract setting, we finally explain one specification of our abstract setting commonly used in the empirical literature. One first fixes some initial date t_0 and some maturity T . The set Ω is given by possible values of an index (e.g. S&P 500 or Dax) at time $t_0 + T$. The first asset S^1 then describes the index; S_0^1 describes the value of the index at time t_0 and S_1^1 describes the index at time $t_0 + T$. The other assets S^2, \dots, S^d describe call options (or put options) on the index with maturity T for different strike prices K . The maturity T is usually chosen to be 1 or 3 months since the options for these maturities are traded most frequently.

In practice, for each seller of a call option there is also a buyer and vice versa. Therefore, the call options are usually assumed to be in zero net supply and it follows that the call options do not have any impact on the aggregate endowment. The aggregate endowment is therefore equal to (a multiple of) the index S^1 (plus some deterministic constant). One tacit motivation for this particular setting is that (changes in) an index is a good

proxy for (changes in) the aggregate endowment in the economy represented by the index.

2) In the setting described in part 1), completeness of the financial market is ensured if there are sufficiently many options. More precisely, it is sufficient to have $N - 2$ call options (for well-chosen strike prices). In a more general infinite-dimensional (but still one-period) setting, one essentially needs to have call options for all strike prices (see for instance Lemma 7.23 in Föllmer and Schied [43] for a precise result in this direction).

3) In practice, the number of different values of an index (equal to the number of states in the present setting) may be assumed to be finite (but very large). However, the number of strike prices for which the call options are traded is much smaller. This leads to an incomplete market. In the empirical literature, one then often interpolates the (traded) call option prices (for different strike prices) to obtain call option prices for all strike prices. This means that one singles out one particular martingale measure and considers a complete market with this particular martingale measure. \diamond

V.4 The pricing kernel puzzle

In this section, we assume that $P_j = P$ for all j and that the preference functional V_j of agent j is given by $V_j(f) := E_{P_j}[U_j(f)] = E[U_j(f)]$ for a strictly increasing, strictly concave and differentiable utility function $U_j : \mathbb{R}_+ \rightarrow \mathbb{R}$ satisfying the *Inada conditions*

$$U'_j(0) := \lim_{x \searrow 0} U'_j(x) = +\infty, \quad (4.1)$$

$$U'_j(\infty) := \lim_{x \rightarrow \infty} U'_j(x) = 0. \quad (4.2)$$

This means that all agents have the same belief which coincides with the objective probability measure P and that each agent is risk-averse in the sense that his utility function is concave. These preference functionals are well known; see for instance Section 2.3 of Föllmer and Schied [43] for a discussion and an axiomatic foundation.

Moreover, we assume that the financial market is complete. In particular, this implies that the pricing kernel is unique. Under these assumptions, there is a decreasing relation between the (unique) pricing kernel and aggregate endowment.

Theorem 4.1. *Consider a financial market satisfying $\dim(\mathcal{X}) = N$ and let the preference functionals V_j be given as above. If $(S_0, (f_j)_{j \in \mathcal{I}})$ is a financial market equilibrium with pricing kernel dQ/dP , then there exists a strictly decreasing function $g : \mathbb{R}_+ \rightarrow \mathbb{R}_+$, such that*

$$\frac{dQ}{dP}(\omega_n) = \frac{Q[\{\omega_n\}]}{P[\{\omega_n\}]} = g(W(\omega_n)), \quad n = 1, \dots, N.$$

A formal proof is given in Theorem 16.7 of Magill and Quinzii [82]. Intuitively, every agent forms his portfolio according to the first order conditions for optimality. Since the financial market is complete, there is a unique martingale measure Q and it follows that the requested final position has the form

$$f_j(\omega_n) = (U'_j)^{-1} \left(\lambda_j \frac{Q[\{\omega_n\}]}{P[\{\omega_n\}]} \right)$$

for a suitable Lagrange parameter λ_j . Because of the decreasing marginal rate of substitution (or in other words, because U'_j is decreasing since U_j is concave), this final position is a decreasing function of the pricing kernel. The same holds true for the sum of all final positions of the agents. Due to the market clearing condition (3.1) and its consequence (3.2), this sum is equal to the aggregate endowment. This implies that in equilibrium the aggregate endowment is a decreasing function of the pricing kernel.

Remark 4.2. An equivalent way of demonstrating Theorem 4.1 goes via aggregation. Since markets are complete, equilibrium allocations are Pareto-efficient and can therefore be supported by the maximization of an aggregate utility, which as a positive weighted sum of the individual utilities inherits concavity. Thus the pricing kernel, being proportional to the gradient of the aggregate utility, is a decreasing function of aggregate endowment. For details, see Chapter 16 in Magill and Quinzii [82]. \diamond

Remark 4.3. The assumptions of Theorem 4.1 can be relaxed. It is enough to assume that the utility functions U_j are increasing and concave (i.e., not necessarily strictly concave and not necessarily satisfying the Inada conditions). Indeed, Theorem 1 of Dybvig [42] and its generalization in Appendix A there show that the final position f_j of agent j and the pricing kernel are anti-comonotonic.⁵ Hence, this also holds for the sum over all agents. Using the market clearing condition, it follows that the sum $W = \sum_{j \in \mathcal{I}} f_j$ and the pricing kernel are anti-comonotonic. \diamond

Remark 4.4. If we restrict ourselves to mean-variance type preferences, we end up in the CAPM which is the traditional example in finance. There, the pricing kernel is an affine decreasing function of the aggregate endowment (see Theorem 17.3 in Magill and Quinzii [82]). \diamond

Remark 4.5. For the specific setting described in Remark 3.2, the probability space is given by possible values of some index and the aggregate

⁵Comonotonicity of two random variables intuitively means that their realizations have the same rank order. In our setup, two random variables X^1 and X^2 are called comonotonic if $(X^1(\omega_n) - X^1(\omega_{n'}))(X^2(\omega_n) - X^2(\omega_{n'})) \geq 0$ for all $n, n' \in \{1, \dots, N\}$. Random variables X^1 and X^2 are called anti-comonotonic if X^1 and $-X^2$ are comonotonic. See Föllmer and Schied [43] for a general definition (Definition 4.82) and equivalent formulations (Lemma 4.89).




Figure V.1: The pricing kernel with respect to the estimated probability measure in Rosenberg and Engle [97].⁷

endowment is given by the index. Moreover, we explained there how additional assumptions on non-traded call options lead to a complete market. Following this path, Theorem 4.1 then tells us that the (unique) pricing kernel is (under the assumptions in Theorem 4.1 imposed on the preferences) a decreasing function of the index. It is tempting to verify this decreasing relation in the empirical data. The (unique) martingale measure can be determined via the second derivative of the prices of the call options with respect to the strike price (see Breeden and Litzenberger [23] for details). In contrast, the objective probability measure P cannot be determined from the data. Instead, a statistician determines an *estimated probability measure* $\hat{P} \approx P$ as a proxy for the objective probability measure P (based on his favourite statistical model/data set). One can think of \hat{P} as a subjective belief \hat{P} (of the statistician) about the future. In this setting, one can therefore only determine the pricing kernel with respect to the estimated probability measure \hat{P} , but not the pricing kernel (with respect to the objective probability measure P). In the empirical literature, one often tacitly assumes that $\hat{P} = P$. However, for a good understanding, a distinction between these objects is necessary.

In the empirical literature, the exercise to determine Q and \hat{P} depending on an index is executed for several indices, several different dates and several specific methods (both for Q and for \hat{P}). The results are always similar. One typical graph is shown in Figure V.1. The x -axis is given by the S&P 500 return (which is up to normalization equal to the S&P 500). The y -axis describes the pricing kernel with respect to \hat{P} . The figure shows two functions (together with confidence bands), each of which corresponds to a different (statistical) method. In particular the non-monotonic line is exem-

⁷Reprinted from *Journal of Financial Economics*, Vol. 64, Rosenberg, J. and Engle R.F., Empirical pricing kernels, p. 361, ©(2002) with permission from Elsevier.

plary for the findings in the empirical literature; the left end of the function is much higher than the right end of the line, but it also has some increasing parts. Assuming (as it is usually done in the empirical literature) that the interpolation of the call option prices is harmless and that $P = \hat{P}$, this means that the empirical observation is not consistent with the theoretical statement in Theorem 4.1 and some of the more fundamental assumptions (imposed on the preferences or on the general setting) seem to be violated. This empirical observation as well as the associated line of arguments goes back to Jackwerth [59] and Brown and Jackwerth [26]. They refer to it as the *pricing kernel puzzle*. \diamond

Motivated by the findings and arguments in the empirical literature regarding the *pricing kernel puzzle*, we find it important to understand the relation between the pricing kernel and the aggregate endowment in our slightly more general setting. In the following section, we first explain the global behaviour of the pricing kernel as a function of the aggregate endowment. We then alternately omit one of the three main assumptions of complete markets, risk aversion, and correct beliefs (i.e. $P_j = P$ for all j) and try to understand how the freedom gained may generate a non-decreasing relation between the pricing kernel(s) and the aggregate endowment.

V.5 Boundary behaviour of the pricing kernel

The preceding section showed that in a complete market, the (unique) pricing kernel is a monotonically decreasing function of the aggregate endowment if the agents are risk-averse and have correct beliefs in the sense that $P_j = P$ for each j . As we show later, as soon as we drop one of the assumptions, the statement does not hold anymore. The goal of this section is to show that even without those assumptions, the values of the pricing kernel are higher for the states with very low aggregate endowment compared to the values of the pricing kernel for the states with very high aggregate endowment. In between, the decreasing relation may be violated.

Formally, the preference functionals are defined by $V_j(f) := E_{P_j}[U_j(f)]$. The probability measure P_j represents the belief of agent j . We assume the following boundary behaviour of the utility functions: For sufficiently small and large arguments, the utility function $U_j : \mathbb{R}_+ \rightarrow \mathbb{R}$ of agent j is strictly increasing, concave and continuously differentiable, and it satisfies the Inada conditions (4.1) and (4.2). Observe that we do not assume that U_j is concave. Moreover, we do not impose any further assumptions on the financial market, i.e. the market can be complete or incomplete which means that there may be multiple pricing kernels. In order to be able to apply our statement on the empirical application explained in Remarks 3.2 and 4.5, we formulate the statement for the pricing kernel with respect to the estimated probability measure $\hat{P} \approx P$. Recall from Remark 4.5 that \hat{P} should be

interpreted as a (subjective) belief of a statistician used as a proxy for the (possibly unobservable) objective probability measure P . Note that except for the assumption $\hat{P} \approx P$, we do not impose any assumption on \hat{P} ; this means that the statement is also valid for the pricing kernel (with respect to the probability measure P) if we (formally) set $\hat{P} = P$. Recall that the aggregate endowment is denoted by

$$W = \sum_{j \in \mathcal{I}} W_j.$$

Without loss of generality, we assume that $0 < W(\omega_1) \leq W(\omega_2) \leq \dots \leq W(\omega_N)$. The next theorem states that if $W(\omega_N) - W(\omega_1)$ is sufficiently large, then the value of the pricing kernel (with respect to \hat{P}) in state ω_1 is much larger than the one in state ω_N .

Theorem 5.1. *Let V_j be defined as above. For $\hat{P} \approx P$ and $\Delta \geq 1$, there is a constant ρ (depending on \hat{P} , $(P_j, U_j)_{j \in \mathcal{I}}$, $W(\omega_1)$ and Δ) with the following property: If $W(\omega_N) - W(\omega_1) > \rho$, then for every equilibrium there is a pricing kernel with respect to \hat{P} satisfying*

$$\frac{Q[\{\omega_1\}]}{\hat{P}[\{\omega_1\}]} \geq \frac{Q[\{\omega_N\}]}{\hat{P}[\{\omega_N\}]} \Delta. \quad (5.1)$$

In particular, the measure Q can be chosen to be proportional to some agent j 's marginal utility, i.e., $Q[\{\omega_n\}] = \lambda_j U_j'(f_j(\omega_n)) P_j[\{\omega_n\}]$, $n = 1, \dots, N$.

Note that ρ does not depend on $W(\omega_N)$; we can therefore choose $W(\omega_N)$ in such a way that $W(\omega_N) - W(\omega_1) > \rho$ is satisfied. The proof of Theorem 5.1 can be found in Appendix V.10.1. Less formally, the statement of Theorem 5.1 can be explained as follows: We fix the preferences of the agents, a belief \hat{P} and W satisfying $W(\omega_N) - W(\omega_1) > \rho$. Every equilibrium in that economy supports at least one martingale measure for which the value of the pricing kernel (with respect to \hat{P}) in the state with low values of aggregate endowment is much higher than the value of the pricing kernel (with respect to \hat{P}) in the state with high values of aggregate endowment. Moreover, the associated martingale measure is a reasonable one; there is at least one agent such that his utility gradient is proportional to the pricing kernel. If the market is complete, there is only one pricing kernel (with respect to \hat{P}) and (5.1) consequently holds for this unique one.

Let us finally discuss two illustrative special cases. In the case of a single agent with $P_1 = \hat{P} = P$, the first order condition for optimality and the Inada conditions guarantee that the values of the pricing kernel are high for the states with very low aggregate endowment and low for those with very high aggregate endowment. If there are two or more agents with $P_j = \hat{P} = P$, the final positions of the agents are small in the states with low aggregate endowment since all agents have positive final positions. Moreover, there is

at least one agent who has a high final position in a state where there is high aggregate endowment (due to the market clearing condition). For this agent, we can again look at the first order condition for optimality and use the Inada condition to deduce (5.1). In the case $\hat{P} \neq P \neq P_j$, the arguments are slightly more technical, but conceptually very similar.

V.6 Non-concave utility

In this section, we consider agents with non-concave utility functions and true beliefs (in the sense that $P_j = P$ for all j) in a complete market. Formally, the preference functional V_j is described by $V_j(f) = E[U_j(f)]$, where $U_j : \mathbb{R}_+ \rightarrow \mathbb{R}$ is strictly increasing. Hence, the main difference to the situation of Theorem 4.1 is that U_j is not necessarily concave. In the literature, the most prominent examples of non-concave utilities are the one suggested by Friedman and Savage [45] and the one arising in prospect theory suggested in Kahneman and Tversky [69].

Before we analyze whether risk-seeking behaviour (in the sense that U_j is not necessarily concave) is a possible reason for the findings in the empirical literature, we want to ensure that risk-seeking behaviour can, in principle, induce an increasing relation between the pricing kernels and the aggregate endowment. In Section V.4, we argued that if the final positions of all agents are decreasing functions of the pricing kernel, the aggregate endowment is one as well. So, in order to have a non-decreasing relation between the pricing kernel(s) and the aggregate endowment, it is necessary that at least one agent has a final position which is not a decreasing function of the pricing kernel. Therefore, the main point is to show that for some agent, the decreasing relation between the pricing kernel and the optimal final position may be violated.

Example 6.1. We consider an economy with two states, two assets and a single (representative) agent. The underlying probabilities are defined by $P[\{\omega_1\}] = 2/3$ and $P[\{\omega_2\}] = 1/3$. The asset prices at time 1 are given by

$$S_1 = (S_1^0, S_1^1) = \begin{pmatrix} 1 & \frac{2}{3} \\ 1 & 2 \end{pmatrix}$$

where the first (second) column describes S_1^0 (S_1^1). The utility function of the agent is given by

$$U(x) = \begin{cases} (x-1)^{\frac{1}{3}}, & x \geq 1, \\ -\lambda(1-x)^{\frac{2}{3}}, & x < 1. \end{cases}$$

The parameter λ describes the loss aversion of the agent and 1 can be interpreted as his reference point. The agent is risk-seeking on the interval $(0, 1)$ and risk-averse on $(1, \infty)$. The stochastic income W_1 of the agent is

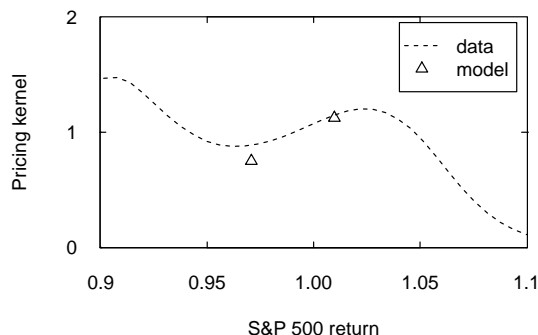


Figure V.2: The pricing kernel for Example 6.1 for a representative agent having a non-concave utility in an economy with two states. The dashed line shows the pricing kernel implied in the data in Figure 2 of Jackwerth [59].

$W_1 = (1.0098, 0.9707)$. It is shown in Appendix V.10.2 that the price vector $S_0 = (S_0^0, S_0^1) = (1, 1)$ together with the allocation $f = W$ forms an equilibrium. In order to analyze the pricing kernel, note first that the martingale measure Q defined by the equation $(1, 1) = (E_Q[S_1^0], E_Q[S_1^1])$ is given by $Q[\{\omega_1\}] = 3/4$ and $Q[\{\omega_2\}] = 1/4$. Hence, the resulting unique pricing kernel is given by $\frac{dQ}{dP}(\omega_1) = 9/8$ and $\frac{dQ}{dP}(\omega_2) = 3/4$. The pricing kernel of this example is shown in Figure V.2. We conclude that there is no decreasing relation between the pricing kernel and the final position of the single representative investor.

Example 6.1 has implications for related questions in portfolio optimization. Dybvig [42] shows that the optimal final position and the pricing kernel are anti-comonotonic if all states have the same probability. This result is then generalized to more general setups and/or more general preferences; see, for instance, Carlier and Dana [31] and He and Zhou [56]. Example 6.1 shows that one cannot drop all assumptions. For general probabilistic structures and a general utility function, the optimal final position and the pricing kernel are not necessarily anti-comonotonic.

While the above example and its implications are of theoretical interest, risk-seeking behaviour is not a reasonable explanation for the empirical findings. In order to justify this doubt, we next show that risk-seeking behaviour on the aggregate level (i.e. a representative agent with a non-concave utility) can be excluded as an explanation for the pattern found in the economic literature. For this we use the next theorem proved in Appendix V.10.3.

Theorem 6.2. *Suppose that $\dim(\mathcal{X}) = N$ and let U be an increasing, differentiable non-concave utility function. Let C denote the interior of the interval where U is strictly convex. Moreover, let f^* be the optimal final po-*

sition for the pricing kernel dQ/dP and stochastic income W . Then, there exists at most one state $n \in \{1, \dots, N\}$ with $f^*(\omega_n) \in C$.

Before combining the theorem and the puzzle, we explain this theorem in a broader context. In Corollary II.5.6, we have seen that on an atomless probability space, the optimal final position f^* and the pricing kernel are anti-comonotonic and that f^* satisfies $P[f^* \in C] = P[f^* \in \{U < U_c\}] = 0$. This implies that on an atomless probability space, non-concave utility functions cannot generate a non-decreasing relation between the pricing kernel and the aggregate endowment. For probability spaces which are not atomless (as in the present setting), this may happen as we have seen in Example 6.1. Theorem 6.2 then provides the insight that even in a general model, it is optimal to choose a final position in such a way that at most one state lies in the set $\{f^* \in C\}$. In the case that there are a lot of states, the influence of a single state becomes small.

Let us now analyze the implications of Theorem 6.2 in the context of Remarks 3.2 and 4.5. Recall that this means that the aggregate endowment is equal to S^1 , that the market is complete, that $\hat{P} = P$ and that the relation between the (unique) pricing kernel and the aggregate endowment $W = S^1$ is defined by the non-monotonic function shown in Figure V.1. We now argue that this situation cannot be explained by a representative agent having a non-concave utility function. The utility function implied in the pricing kernel (formally by integrating the pricing kernel) in Figure V.1 is concave for low and high values and convex in the middle part. In particular, the utility function is strictly convex in the area between -4% and 2% return of the aggregate endowment. However, a representative agent with such a utility function would, as shown in Theorem 6.2, choose a final position f^* in such a way that at most one state lies in the set where the aggregate endowment has a return between -4% and 2% . But claims to payoffs in these states have to be held by someone. Hence, prices need to adjust such that the pricing kernel is a decreasing function of the aggregate endowment and it becomes more attractive again to hold the assets which provide payoffs in those states. This shows that the pricing kernel shown in Figure V.1 is not consistent with a representative agent having a non-concave utility.

We conclude that non-concave utilities can, in principle, be seen as a possible argument for a non-decreasing relation between the pricing kernel and the aggregate endowment. However, it only works for pathological examples with few states and, in our setup, it is not an explanation for the pattern found in the empirical literature.

One other way of generating a non-decreasing relation between the pricing kernel and the aggregate endowment is to generalize the utility function to be state-dependent, i.e., $V(X) := \sum_{n=1}^N P[\{\omega_n\}]U(f(\omega_n), \omega_n)$ where the utility $U(\cdot, \omega_n)$ is concave for every $n = 1, \dots, N$. Then for any pricing kernel dQ/dP there exist a state-dependent utility function $U(\cdot, \omega_n)$ such that

the first order condition for optimality

$$\frac{Q[\{\omega_n\}]}{P[\{\omega_n\}]} = \lambda U'(f(\omega_n), \omega_n), n = 1, \dots, N$$

holds. In particular, no robustness problem arises. Assuming for example that the degree of risk aversion is state-dependent, i.e.,

$$U(x, \omega_n) = \frac{x^{1-\alpha_n}}{1-\alpha_n}, n = 1, \dots, N,$$

one can generate the typical form of the pricing kernel by assuming that for small losses the investors are less risk-averse than for large gains or losses. Since state-dependent utilities are so flexible, the challenge is to give plausible restrictions such that the pricing kernel is a generally decreasing function of the aggregate endowment with increasing areas. Attempts in this direction are done in Chabi-Yo et al. [33] and Benzoni et al. [12] by generalizing the model to multiple periods.

V.7 Incorrect beliefs

We now analyze settings where the belief P_j of agent j may differ from the objective probability P (and also from the beliefs $P_{j'}$ of the other agents). There are several different motivations for such a setting. One is to account for the fact that each agent in an economy may have a different view about the future. But even in a setting where there is only one agent knowing P (e.g. in an experiment) it might be that the agent distorts probabilities. This leads to $P \neq P_j$ again, but refers to a bias in decision making (we explain this in more detail in Section V.7.1). We first analyze the utility maximization problem for general incorrect beliefs, then analyze the phenomena independently, and finally combine them in the last part of this section.

In order to formally describe the different phenomena in a unified way, we assume that the preferences of agent $j \in \mathcal{I}$ are described by the functional

$$V_j(f) := E_{P_j}[U_j(f)] = \sum_{n=1}^N P_j[\{\omega_n\}]U_j(f(\omega_n)),$$

where P_j is a set function on \mathcal{F} .⁸ The set function P_j represents the (subjective) belief of agent j about the future. In order to isolate the effect of incorrect beliefs, we consider again the case of a complete market and strictly concave utility functions satisfying the Inada conditions. In equilibrium, each agent j maximizes his preference functional V_j over the set $\mathcal{C}_j(S_0)$

⁸A set function P_j on \mathcal{F} is a function $P_j : \mathcal{F} \rightarrow [0, 1]$ satisfying $P_j[\emptyset] = 0$ and $P_j[\Omega] = 1$; see Chapter 2 in Denneberg [39] for a detailed discussion. In particular, a probability measure as well as a distorted probability measure are set functions.

and the final position f_j of agent j solves the maximization problem

$$\text{maximize } E_{P_j}[U_j(f)] \text{ over } f \in \mathcal{C}_j(S_0). \quad (7.1)$$

By the definition of $\mathcal{C}_j(S_0)$, each final position f is of the form $W_j + \sum_{i=0}^d S_1^i \vartheta_j^i$, where ϑ_j satisfies the budget restriction $\sum_{i=0}^d S_0^i \vartheta_j^i \leq 0$. Since the prices $S_0^i = E_Q[S_1^i]$ can be written in terms of the martingale measure Q , the constraint $\sum_{i=0}^d S_0^i \vartheta_j^i \leq 0$ can be rewritten as $E_Q[f] \leq E_Q[W_j]$ (for details see Magill and Quinzii [82, Page 83ff]). This modification allows the Lagrange method to be used and we conclude that the final position of agent j has the form

$$f_j(\omega_n) = (U_j')^{-1} \left(\lambda_j \frac{Q[\{\omega_n\}]}{P_j[\{\omega_n\}]} \right)$$

for a suitable Lagrange parameter λ_j . Note that the objective probability P does not appear in the final position of agent j . The final position f_j of agent j is a decreasing function of the pricing kernel with respect to P_j . However, it is not necessarily a decreasing function neither of the pricing kernel (with respect to P) nor of the pricing kernel with respect to another $P_{j'}$. Thus, in a model with a single (representative) agent, a non-decreasing relation between the pricing kernel (with respect to P) and the aggregate endowment can be viewed as a difference between the objective probability P and the belief P_1 of the representative agent. In order to relate incorrect beliefs and the pattern found in the empirical literature, we now analyze different sources of incorrect beliefs in the context of Remarks 3.2 and 4.5 in the following sense: We assume that the aggregate endowment is equal to S^1 , that the market is complete and that there is a representative agent. Moreover, we fix an estimated probability measure \hat{P} .

V.7.1 Distorted beliefs

In this section, we consider distortions (in isolation) as a special case for incorrect beliefs, i.e., we assume, in addition, that $P = \hat{P}$. Kahneman and Tversky [69] show that agents tend to overweight extreme events which occur with small probability. In order to make this idea precise, we fix a payoff f which is ordered in an increasing way and we fix an increasing and continuous function $T : [0, 1] \rightarrow [0, 1]$ with $T(0) = 0$ and $T(1) = 1$. The preferences are defined by

$$V(f) := \sum_{n=1}^N U(f(\omega_n)) \left(T(P[\bigcup_{k=1}^n \{\omega_k\}]) - T(P[\bigcup_{k=1}^{n-1} \{\omega_k\}]) \right), n = 1, \dots, N. \quad (7.2)$$

Note that T is applied to the cumulative distribution rather than the decumulative distribution function.⁹ In the classical case without distortion

⁹We follow here the approach of Polkovnichenko and Zhao [86] and Dierkes [41]. This has the advantage that one can compare our results with theirs.




Figure V.3: Distortion implied in the data in Figure 2 of Jackwerth [59] and the parametric form of Prelec [87] for $\gamma = 0.7$.

function (i.e. $T(p) = p$), the term $T(P[\bigcup_{k=1}^n \{\omega_k\}]) - T(P[\bigcup_{k=1}^{n-1} \{\omega_k\}])$ reduces to $P[\{\omega_n\}]$ and we recover the classical expected utility functional. The term $T(P[\bigcup_{k=1}^n \{\omega_k\}]) - T(P[\bigcup_{k=1}^{n-1} \{\omega_k\}])$ is bigger than $P[\{\omega_n\}]$ if $T' > 1$ between $P[\bigcup_{k=1}^n \{\omega_k\}]$ and $P[\bigcup_{k=1}^{n-1} \{\omega_k\}]$ and it is smaller than $P[\{\omega_n\}]$ if $T' < 1$ between $P[\bigcup_{k=1}^n \{\omega_k\}]$ and $P[\bigcup_{k=1}^{n-1} \{\omega_k\}]$. If T has a concave-convex form as the dashed line in Figure V.3, the distortion T satisfies $T' > 1$ close to 0 and 1. Consequently, the functional given in (7.2) overweights the influence of the extreme events on the payoff f as observed in the experimental literature. One widely used parametric specification of the weighting function, due to Kahneman and Tversky [69], is

$$T(p) = \frac{p^\gamma}{(p^\gamma + (1-p)^\gamma)^{\frac{1}{\gamma}}}. \quad (7.3)$$

Another widely used specification, due to Prelec [87], is

$$T(p) = \exp(-(-\log(p))^\gamma). \quad (7.4)$$

Experimental studies find $\gamma \in [0.5, 0.7]$ for (7.3) as well as for (7.4). The function T described in (7.4) is shown in Figure V.3 for $\gamma = 0.7$ (dashed line).

If the different values of f occur with small probabilities (as in the case when f describes (the return of) an index such as the S&P 500), we have

$$T(P[\bigcup_{k=1}^n \{\omega_k\}]) - T(P[\bigcup_{k=1}^{n-1} \{\omega_k\}]) \approx P[\{\omega_n\}]T'(P[\bigcup_{k=1}^n \{\omega_k\}]).$$

Denoting the distribution function of f by F , we can therefore write the resulting first order condition for optimality approximately as

$$U'(f(\omega_n))P[\{\omega_n\}]T'(F(f(\omega_n))) \approx \lambda Q[\{\omega_n\}]; \quad (7.5)$$

for details, see Polkovnichenko and Zhao [86]. If there is no distortion (i.e. $T(p) = p$), it follows that $T'(p) = 1$ and (7.5) coincides with the first order condition for optimality for the classical expected utility functional. Note that the first order conditions are necessary, but not sufficient for optimality in the present case. By writing down (7.5), we thus implicitly assume that there exists a representative agent with preferences described by (7.2) and that the payoff f is his optimal final position for (7.1). These assumptions are restrictive.

In the literature, the relation (7.5) is used to deduce the probability distortions implied by the prices of the S&P 500 index and the associated options. Dierkes [41] assumes a parametric form for the utility and the distortion and calibrates (7.5) yielding parameters that are comparable to the ones suggested in the experimental studies. In addition, Polkovnichenko and Zhao [86] use (7.5) to estimate the distortion non-parametrically. They find that the distortion functions are time-varying. For some years, the form is concave-convex as suggested in the experimental literature; for other years, the form is slightly different. For comparison, we apply a similar approach here. We set $U(x) := \log(x)$, fix $P = \hat{P}$, Q and f as given in Figure 2 of Jackwerth [59] and determine T in such a way that (7.5) is satisfied. Note that the (discretized) lognormal distribution fitted to \hat{P} has parameters $\mu = 0.0029$ and $\sigma = 0.0388$. The implied distortion T is shown in Figure V.3 (solid line). While the form of the implied distortion is similar to the one suggested by experimental studies (dashed line) for low values, the form is substantially different for high values. So even if the restrictive assumptions held true, the empirical implications would not be reasonable. We conclude that a distortion, in isolation, is not a good explanation for the pattern found in the empirical literature. One reason for this, we believe, is that \hat{P} is a bad proxy for the probabilities P perceived by the agent. In the next subsection, we therefore analyze the case $P \neq \hat{P}$ in more detail and we then combine distortions and $P \neq \hat{P}$ in Subsection V.7.3.

Remark 7.1. Recently, Gollier [48] has used ambiguity aversion to explain the pricing kernel puzzle. In his model, the agent faces a finite set of beliefs and he assigns a probability to each of these beliefs. The pricing kernel is

then defined with respect to a weighted belief. It is then argued that the aversion to ambiguity of the representative agent affects the equilibrium in a way that is observationally equivalent to a distortion in our framework. \diamond

V.7.2 Misestimated beliefs

In the analysis in Section V.7.1, we have assumed that $P = P_1 = \hat{P}$. This means that a statistician may observe the objective probability measure P as well as the belief P_1 of the representative agent. However, these assumptions are questionable. In practice, the procedure for the estimation uses past data, whereas the belief(s) of the agent(s) are forward-looking. In this section, we therefore analyze the case $P = P_1 \neq \hat{P}$ in more detail.¹⁰ We refer to this situation as *misestimated beliefs* since the estimated belief \hat{P} is a bad proxy for the belief P_1 .

In the present representative agent economy, equilibrium asset prices reflect the agent's preferences and beliefs. More precisely, any two of the following imply the third: the martingale measure, the belief of the representative agent and the preferences of the representative agent. One approach to explore misestimated beliefs is therefore to set up a model for the belief and the preferences of the representative agents and to verify whether or not such a model can explain the estimates for the martingale measure in the data. This is done in detail in Ziegler [111]. In the simplest case, this boils down to fixing a parametric utility function (e.g., log-utility, power utility or exponential utility) and specifying the belief P_1 . Recall that in the present setting the probability space is given by the values of S^1 . The belief P_1 can thus be specified via the distribution of S_1^1 . One tractable way is to assume that the distribution of S_1^1 is given by a discretized (and renormalized) lognormal distribution with parameters μ and σ . This determines the martingale measure via the first order condition

$$U'(S_1^1(\omega_n))P_1[\{\omega_n\}] = \lambda Q[\{\omega_n\}], \text{ for } n = 1, \dots, N. \quad (7.6)$$

We refer to Q defined via (7.6) as implied martingale measure. A calibration analysis then shows whether or not this implied martingale measure Q can reproduce the (empirical) martingale measure \hat{Q} estimated in the data. As explained in Ziegler [111], the advantage of this approach (compared to the calibration of the pricing kernel) is that the tails are less important. As an illustration for the calibration, we set $U(x) = \log(x)$ and choose the empirical martingale measure \hat{Q} to be the one given in Figure 2 of Jackwerth [59]. Note that the implied martingale measure now depends on three parameters λ , μ and σ . We now fit this implied martingale measure to \hat{Q} using nonlinear least-squares. This procedure yields the parameters $\lambda = 1.0398$, $\mu = 0.0061$

¹⁰On the first point of view, the assumption $P = P_1$ may still appear restrictive. However, note that in a setting with $P \neq P_1$, the role of P is somewhat artificial since P is not relevant for the equilibrium. We therefore directly start with $P = P_1$.

and $\sigma = 0.0341$. Even though these parameters seem to be reasonable, the comparison of the implied and the empirical martingale measure (the probabilities $Q[\{\omega_i\}]$ and $\hat{Q}[\{\omega_i\}]$ of the implied as well as of the empirical martingale measure are depicted in Figure V.4) shows that the probability weights of the implied martingale measure (dotted line) misses some essential features of the probability weights of the estimated martingale measure (solid line). This is in line with the results of Ziegler [111]. He also checks whether mixtures of lognormal beliefs or the extension to several agents give better results and concludes that the degree of pessimism required to explain the puzzle is implausibly high. This leads to the conclusion that misestimated beliefs, in isolation, are not sufficient to explain the puzzle.

Remark 7.2. The approach used here is conceptually similar to Bliss and Panigirtzoglou [21]. They also form a subjective distribution based on the implied martingale measure via the first order condition for optimality. In contrast to our calibration, they use power utilities and fit the risk aversion coefficient in order to explain the observations. As in our approach in this subsection, this implicitly assumes a decreasing relation between the pricing kernel and the aggregate endowment. \diamond

V.7.3 Combination of distortion and misestimation

We now combine the arguments of the preceding subsections. As in Section V.7.2, we assume that the belief of the representative agent is specified via a discretized (and renormalized) lognormal distribution about (the return of) the index S_1^1 with parameters μ and σ . As in Section V.7.1, the agent distorts the perceived distribution in the way described in Polkovnichenko and Zhao [86]. Following them, we also make the (tricky) assumption that the index is the optimal final position for the representative agent. As described in Section V.7.1, the latter leads to the first order condition

$$U'(S_1^1(\omega_n))P_1[\{\omega_n\}]T'(F_{P_1}(f(\omega_n))) \approx \lambda Q[\{\omega_n\}], \text{ for } n = 1, \dots, N, \quad (7.7)$$

where $P_1[\{\omega_n\}]$ and F_{P_1} are the weights and the cumulative distribution function of the belief. For the distortion T , we fix the parametric form given in (7.4) which has some computational advantages compared to (7.3). For the utility, we set $U(x) = \log(x)$. For this representative agent, the implied martingale measure Q (determined via (7.7)) depends on the four parameters λ, μ, σ and γ . We then calibrate (7.7) to the estimate of the martingale measure given in Figure 2 of Jackwerth [59] using nonlinear least-squares yielding the parameters $\lambda = 1.0177$, $\mu = -0.0015$, $\sigma = 0.0207$ and $\gamma = 0.5401$. The implied martingale measure for this representative agent (and the estimated values for λ, μ, σ and γ) is also depicted in Figure V.4. On the one hand, we see that the model with distortion (dashed line)—in




Figure V.4: The solid line shows the probability density (weights) of the martingale measure estimated in Figure 2 of Jackwerth [59]. The probability density (weights) of the fitted martingale measure with distorted beliefs (dashed line) closely reproduces the solid line. The probability density (weights) of the fitted martingale measure in the model without distortion (dotted line), on the other hand, misses some essential features of the data.

contrast to the model without distortion (dotted line)—closely reproduces the probability density (weights) of the empirical martingale measure (solid line). It captures well the essential features well, in particular its thick left tail, and the parameter γ for the distortion is comparable to experimental studies. But on the other hand, we also observe that μ is negative which is not plausible.

To highlight the main insight of our analysis, let us recall the findings of Ziegler [111]. The most promising attempt to explain the puzzle there is to consider two groups of agents with heterogeneous beliefs. He finds that the pessimism required to explain the puzzle is implausibly high. A closer inspection of the result shows that the pessimism is necessary to generate the fat left tail of the probability density function of the empirical martingale measure. In our case with distortion, the agents overweight the extreme events and this leads to a fat left tail in the probability density of the implied martingale measure. This allows us to closely reproduce the empirical martingale measure density with only one (group of) agent(s).

V.8 Incomplete markets

Up to now, we have restricted ourselves to the case of a complete market economy, i.e. $\dim(\mathcal{X}) = N$. In this section, we want to analyze the case $\dim(\mathcal{X}) < N$. To understand the impact of market incompleteness on the shape of the pricing kernel, we isolate this extension in the sense that we keep the other two assumptions that agents are risk-averse and that all the agents have common beliefs equal to the objective probability measure P , i.e. $P_j = P$ for all j .

In equilibrium, every agent maximizes his preference functional subject to his budget set. This means that each agent solves a problem of the form

$$\max E[U_j(f)] \text{ over } f \in \mathcal{C}_j(S_0) \quad (8.1)$$

for the equilibrium asset prices S_0 . In incomplete markets, there are infinitely many probability measures $Q \approx P$ satisfying the equation $S_0^i = E_Q[S_1^i]$ for each $i = 0, \dots, d$ and hence there are also infinitely many pricing kernels. The constraint $\sum_{i=0}^d S_0^i \vartheta_j^i \leq 0$ can be rewritten using the pricing kernels, and writing down the first order conditions for optimality for that optimization problem, it turns out (Theorem 10.4 in Magill and Quinzii [82]) that every solution for (8.1) has the same form as in the complete market case for a particular pricing kernel. More precisely, for every agent j with optimal final position f_j , there is a martingale measure Q_j such that his final position is of the form

$$f_j = (U_j')^{-1} \left(\lambda_j \frac{dQ_j}{dP} \right). \quad (8.2)$$

This shows that every final position is a decreasing function of **some** pricing kernel. Put differently, if f_j is the optimal final position of an agent with a concave utility U_j , then $U_j'(f_j)/E[U_j'(f_j)]$ is a pricing kernel. Hence, if there is a single representative agent, there exists some pricing kernel such that the final position of the representative agent (which is equal to the aggregate endowment by the market clearing condition) is a decreasing function of that pricing kernel. However, the assumption of a representative agent in an incomplete market is a delicate one. To substantiate this, consider an economy with heterogeneous agents. For every agent j , $U_j'(f_j)/E[U_j'(f_j)]$ is a pricing kernel. However, due to the incompleteness, the different pricing kernels are not necessarily the same. This can be used to give an example in which no pricing kernel is a decreasing function of the aggregate endowment. That is to say, in our example we have a unique equilibrium allocation and whatever martingale measure we select out of the continuum of martingale measures supporting it, the resulting pricing kernel is not a decreasing function of the aggregate endowment. If one now assumes that there is a risk-averse representative agent, one cannot explain the asset prices even though they result from a (classical) financial market equilibrium with risk-averse agents having correct beliefs $P_j = P$ for all j . This is illustrated in the next example.

Example 8.1. We consider an economy with three states, two assets and two agents. The underlying probabilities are defined by $P[\{\omega_n\}] = 1/3$ for $n = 1, 2, 3$. The prices of the assets at time 1 are given by

$$S_1 = (S_1^0, S_1^1) = \begin{pmatrix} 1 & \frac{13}{15} \\ 1 & \frac{4}{3} \\ 1 & \frac{13}{15} \end{pmatrix}.$$

There are two agents. Both of them have utility $U_1(x) = U_2(x) = \log(x)$ and they have common and true beliefs, i.e., they evaluate utilities according to the probabilities $P[\{\omega_n\}] = 1/3$ for $n = 1, 2, 3$. The stochastic incomes are given by $W_1 = (\frac{76}{45}, \frac{19}{6}, \frac{212}{15})$ and $W_2 = (\frac{418}{15}, \frac{23}{6}, \frac{382}{285})$. It is shown in Appendix V.10.4 that $S_0 = (S_0^0, S_0^1) = (1, 1)$, $f_1 = (\frac{14}{9}, \frac{7}{2}, 14)$ and $f_2 = (28, \frac{7}{2}, \frac{28}{19})$ is the unique financial market equilibrium. In order to analyze the pricing kernels, note first that every martingale measure Q satisfies the equation $Q[\{\omega_1\}]\frac{13}{15} + Q[\{\omega_2\}]\frac{4}{3} + Q[\{\omega_3\}]\frac{13}{15} = 1$. It follows that all these probability measures can be written as a convex combination of the two extreme points $(0, \frac{2}{7}, \frac{5}{7})$ and $(\frac{5}{7}, \frac{2}{7}, 0)$. We infer that $Q[\{\omega_2\}] < \max(Q[\{\omega_1\}], Q[\{\omega_3\}])$ holds for every martingale measure Q . Because of $P[\{\omega_n\}] = 1/3$ for $n = 1, 2, 3$, the same holds true for all pricing kernels, i.e.,

$$\frac{Q[\{\omega_2\}]}{P[\{\omega_2\}]} \leq \max\left(\frac{Q[\{\omega_1\}]}{P[\{\omega_1\}]}, \frac{Q[\{\omega_3\}]}{P[\{\omega_3\}]}\right)$$

for all martingale measures Q . We infer that **no** pricing kernel is a decreasing function of the aggregate endowment since the aggregate endowment $W = f_1 + f_2 = (\frac{266}{9}, 7, \frac{294}{19})$ has the lowest value in state 2.

In order to decide whether incomplete markets provide an explanation for the pattern found in the empirical literature, we have to specify the reason(s) for the incompleteness. In the remainder of this section, we separately consider the effect on the pricing kernel of illiquid options for extreme strikes as well as of heterogeneous background risk.

V.8.1 Incompleteness due to a lack of options

Recall from Remark 3.2 that in the empirical literature, one usually assumes that there are sufficiently many call options ensuring completeness of the market. This assumption is questionable; in particular for extreme strike prices, there are not many liquid call options. We therefore shortly discuss one related special case of our setup. As in Remark 3.2, we assume that the probability space is given by the values of “an index” S^1 and that the aggregate endowment is equal to this index (plus a deterministic constant). We assume that there are sufficiently many call options with strike prices between -5% and 5% return of the index such that the martingale probabilities for these states are uniquely determined. For the “extreme” states (below

−5% and above 5% return of the index), there are not sufficiently many call options and the martingale probabilities for these states are not uniquely determined. Formally, this means that for any two martingale measures Q and Q' , we have $Q[\{\omega\}] = Q'[\{\omega\}]$ for the “middle” states ω where the index is between −5% and 5% return of the index. For the other states ω , $Q[\{\omega\}]$ and $Q'[\{\omega\}]$ might differ. It then follows from (8.2) that for the “middle” states, the final position f_j of agent j is a decreasing function of dQ/dP (which is unique for these states). Thus also the aggregate endowment (which is by (3.2) the sum over f_j) is for these “middle” states a decreasing function of dQ/dP . However, recall from Figure V.1 that the violation of the decreasing relation between dQ/dP and the aggregate endowment is in this middle part. We conclude that a lack of call options for extreme strike prices (in the sense specified above) is not a reasonable explanation for the pattern found in the empirical literature.

V.8.2 Incompleteness due to background risk

A second main assumption in the setting described in Remark 3.2 is that the probability space is defined by the values of the index. This means that the different values of the index reflect all the uncertainty in the economy. In reality, however, there can be other risks such as labour risk or housing risk which is not necessarily captured by the index (see Franke et al. [44] and references therein for detailed explanations). So even if the market spanned by the index and its options is complete, there can be other risks which cannot be insured completely by the agents. The goal of this section is to explain this form of incompleteness in detail and relate it to the empirical pattern in the literature. We consider four states and three assets. For simplicity, we assume that $P[\{\omega_2\}] = P[\{\omega_3\}]$. The matrix of the assets at time 1 is given by

$$S_1 = (S_1^0, S_1^1, S_1^2) = \begin{pmatrix} 1 & a & 0 \\ 1 & b & d \\ 1 & b & d \\ 1 & c & e \end{pmatrix}$$

for values $0 < a < 1 < b < c$ and $0 < d < e$ such that arbitrage is excluded. We assume that asset S^2 is in zero net supply. For the prices $S_0 = (S_0^0, S_0^1, S_0^2) := (1, 1, 1)$, this market is incomplete and there are infinitely many martingale measures $Q \approx P$, each of which can be written as a convex combination of two extreme points $\underline{Q} = (q_1, q_{23}, 0, q_4)$ and $\overline{Q} = (q_1, 0, q_{23}, q_4)$. However, given the information of the three assets, we cannot separate state 2 and state 3. In this sense, the market spanned by the bank account, the “index” (asset 1) and the call option written on the index (asset 2) has only 3 observable states and the market which can be observed from the prices is complete. Hence there is a unique set of state prices for

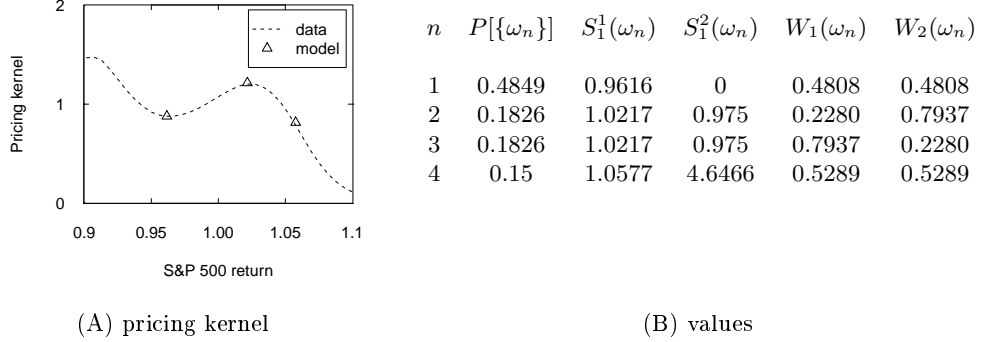


Figure V.5: Panel (A) shows the (unique) pricing kernel for the market with respect to the filtration generated by the assets. For comparison, the dashed line shows the pricing kernel implied in the data in Figure 2 of Jackwerth [59]. Panel (B) shows the corresponding values for risky probabilities, the payoffs for assets S_1^1 and S_1^2 and the stochastic incomes W_1 and W_2 . The stochastic income W_j and the parameter e are chosen as described in Appendix V.10.5.

this marketed subspace which means that $Q[\{\omega_1\}]$, $Q[\{\omega_2\}] + Q[\{\omega_3\}]$ and $Q[\{\omega_4\}]$ are uniquely defined. But, on the other hand, the individual agents may face risks which are not captured by the index. In our example, this means that the stochastic income does not necessarily lie in the marketed subspace. For illustrative purposes, we consider the (extreme) case that the background risk only matters on the individual level, i.e., the aggregate endowment lies again in the marketed subspace. In this way, we can construct a situation where the index is equal to the aggregate endowment (plus a deterministic constant); the market with respect to the filtration generated by the assets is complete and admits a unique pricing kernel. However, the agents face a portfolio selection problem in the (incomplete) market with respect to another (larger) filtration. It is shown in Appendix V.10.5 that for all parameters values a, b, c, d , $P[\{\omega_1\}]$, $P[\{\omega_2\}] = P[\{\omega_3\}]$ and $P[\{\omega_4\}]$, we can choose a parameter e , and stochastic incomes W_1 and W_2 such that the aggregate endowment $W_1 + W_2$ is equal to the first asset S^1 and that $(W_j)_{j=1,2}$ together with the prices $S_0 = (S_0^0, S_0^1, S_0^2) = (1, 1, 1)$ form an equilibrium for two agents having logarithmic utility functions. We choose the parameters a, b, c, d , $P[\{\omega_1\}]$, $P[\{\omega_2\}] = P[\{\omega_3\}]$ and $P[\{\omega_4\}]$ in such a way that we can capture the essential features of the pricing kernel of Figure 2 in Jackwerth [59]. The corresponding values and the pricing kernel are shown in Figure V.5. The constructed example has only three observable states, but the same idea can easily be extended to arbitrarily many states. We also made the rather restrictive assumption that background risk only matters on the individual level. Relaxing this assumption as well gives additional

freedom. Everything being so flexible, the challenge is then to give plausible restrictions such that the pricing kernel is a generally decreasing function of the aggregate endowment which may have increasing areas. One frequently used method to specify background risk on the aggregate level in combination with state-dependent utilities is to introduce a second state variable (Garcia et al. [46], Chabi-Yo et al. [33] and Benzoni et al. [12]). In those models, the pricing kernel is not only a function of the aggregate endowment but also of other state variables. Considering the pricing kernel only as a function of aggregate endowment, the other state variables can be seen as background risk which is not captured by the aggregate endowment.

V.9 Conclusion

In the present chapter, we finally study the implications of behavioural effects such as risk-seeking behaviour and distorted beliefs on the financial market equilibrium. We focus on the relation between the pricing kernel (or pricing density) dQ/dP and the aggregate endowment in the economy.

In an economy with complete markets and risk-averse investors having correct beliefs, the pricing kernel is a monotonically decreasing function of aggregate endowment. As soon as we relax at least one assumption, one can construct examples where the decreasing relation between the pricing kernel and the aggregate endowment is violated. Non-concave utility functions and distorted beliefs can therefore lead to an increasing relation between the pricing kernel and the aggregate endowment. In this sense, behavioural effects can explain the pricing kernel puzzle in empirical finance. In order to check whether the behavioural effects can explain the empirical findings, we compare the resulting pricing kernel with the empirical estimates from the S&P 500. We conclude that risk-seeking behaviour as well as distorted beliefs are, in isolation, not satisfying explanations for the empirical estimates.

V.10 Appendix: Proofs

V.10.1 Proof of Theorem 5.1

Proof. We fix some probability measure $\hat{P} \approx P$ and $\Delta \geq 1$. We want to show that we can find ρ and some martingale measure $Q \approx P$ such that

$$W(\omega_N) - W(\omega_1) > \rho \implies \frac{Q[\{\omega_1\}]}{\hat{P}[\{\omega_1\}]} \geq \frac{Q[\{\omega_N\}]}{\hat{P}[\{\omega_N\}]} \Delta$$

holds in equilibrium. The idea is as follows: In equilibrium, every agent j is optimizing his expected non-concave utility. Hence his final position satisfies the first order condition for optimality

$$U'_j(f_j(\omega_n)) = \lambda_j \frac{Q_j[\{\omega_n\}]}{P_j[\{\omega_n\}]} \quad \text{for all } n = 1, \dots, N,$$

for a particular martingale measure Q_j as described in Proposition IV.3.2. Solving for λ_j and setting the term for state 1 equal to the term for state N gives

$$\frac{P_j[\{\omega_1\}]U'_j(f_j(\omega_1))}{P_j[\{\omega_N\}]U'_j(f_j(\omega_N))} = \frac{Q_j[\{\omega_1\}]}{Q_j[\{\omega_N\}]} \quad (10.1)$$

Since the final positions have to be positive, it follows from the market clearing condition that $f_j(\omega_1) \leq W(\omega_1)$ holds for all agents j . The term $U'_j(f_j(\omega_1))$ is therefore large as $W(\omega_1)$ is small because of the Inada condition (4.1) at 0. On the other hand, if $W(\omega_N) \rightarrow \infty$, the market clearing condition $W(\omega_N) = \sum_{j \in I} f_j(\omega_N)$ implies that at least one agent (say i) has an arbitrarily large final position in state N . It follows then from the Inada condition (4.2) at ∞ , that $U'_i(f_i(\omega_N))$ is arbitrarily small. The term $U'_j(f_j(\omega_1))/U'_j(f_j(\omega_N))$ is therefore arbitrarily large if $W(\omega_N) - W(\omega_1)$ is sufficiently large.

Let us now be precise. Recall that $W(\omega_1) > 0$. The strict monotonicity of U_j , continuous differentiability of U_j , and the Inada condition (4.1) at 0 imply that $z := \inf_{\tilde{x} \in [0, W(\omega_1)]} U'_j(\tilde{x}) > 0$ is attained. Because of the Inada condition (4.2) at ∞ , there is x_j such that

$$U'_j(x) \leq \frac{z}{\Delta} \frac{P_j[\{\omega_1\}]}{P_j[\{\omega_N\}]} \frac{\hat{P}[\{\omega_N\}]}{\hat{P}[\{\omega_1\}]}$$

for all $x \geq x_j$. This gives

$$\frac{P_j[\{\omega_1\}] \inf_{\tilde{x} \in [0, W(\omega_1)]} U'_j(\tilde{x})}{P_j[\{\omega_N\}] U'_j(x)} = \frac{P_j[\{\omega_1\}] z}{P_j[\{\omega_N\}] U'_j(x)} \geq \frac{\hat{P}[\{\omega_1\}]}{\hat{P}[\{\omega_N\}]} \Delta \quad (10.2)$$

for all $x \geq x_j$. We fix $\rho := |\mathcal{I}| \max_j x_j$, where $|\mathcal{I}|$ denotes the number of agents in the economy. The condition $W(\omega_N) - W(\omega_1) > \rho$ then gives $W(\omega_N) > W(\omega_1) + \rho \geq \rho = |\mathcal{I}| \max_j x_j$. In the equilibrium, the market clearing condition $W(\omega_N) = \sum_{j \in I} f_j(\omega_N)$ implies that there is at least one agent (say i) whose final position in state N satisfies $f_i(\omega_N) \geq W(\omega_N)/|\mathcal{I}|$. Putting together (10.1), the inequalities $f_i(\omega_N) \geq W(\omega_N)/|\mathcal{I}| \geq \max_j x_j$ and $U'_i(f_i(\omega_1)) \geq z$, as well as (10.2), we obtain

$$\frac{Q_i[\{\omega_1\}]}{Q_i[\{\omega_N\}]} = \frac{P_i[\{\omega_1\}] U'_i(f_i(\omega_1))}{P_i[\{\omega_N\}] U'_i(f_i(\omega_N))} \geq \frac{P_i[\{\omega_1\}] z}{P_i[\{\omega_N\}] U'_i(f_i(\omega_N))} \geq \frac{\hat{P}[\{\omega_1\}]}{\hat{P}[\{\omega_N\}]} \Delta,$$

which is equivalent to $\frac{Q_i[\{\omega_1\}]}{\hat{P}[\{\omega_1\}]} \geq \frac{Q_i[\{\omega_N\}]}{\hat{P}[\{\omega_N\}]} \Delta$. \square

V.10.2 Example 6.1

Before formally verifying the equilibrium conditions, let us shortly explain how we have constructed this example. We first fixed the prices $S_0 = (1, 1)$

and the initial capital 1. We then solved the non-concave utility maximization problem for fixed prices. Setting the stochastic income W_1 to be equal to the optimal final position, it then follows by construction that W_1 and $S_0 = (1, 1)$ form an equilibrium.

Let us now formally verify that the above approach indeed gives an equilibrium. The market clearing condition is trivially satisfied; we only need to verify that $f_1 = W_1$ maximizes V_1 over $\mathcal{C}_1(S_0)$. The martingale measure for $S_0 = (S_0^0, S_0^1) = (1, 1)$ is $Q[\{\omega_1\}] = 3/4$ and $Q[\{\omega_2\}] = 1/4$ and it follows that $E_Q[W_1] = 1$. It is therefore sufficient to show that $(\hat{f}(\omega_1), \hat{f}(\omega_2)) \simeq (1.0098, 0.9707)$ solves the problem

$$\max E[U(f)] \text{ over } f \geq 0, E_Q[f] \leq 1,$$

where the inequality constraint $E_Q[f] \leq 1$ can be replaced by an equality constraint $E_Q[f] = 1$ since U is strictly increasing. In order to prove optimality of \hat{f} , we consider the three cases $f(\omega_1) = f(\omega_2) = 1$, $f(\omega_1) > 1 > f(\omega_2)$ and $f(\omega_1) < 1 < f(\omega_2)$ independently. In the case $f(\omega_1) > 1 > f(\omega_2)$, plugging in the constraint $E_Q[f] = 1$, differentiating

$$E[U(f)] = P[\{\omega_1\}]U(f(\omega_1)) + P[\{\omega_2\}]U\left(\frac{1 - Q[\{\omega_1\}]f(\omega_1)}{Q[\{\omega_2\}]}\right)$$

with respect to $f(\omega_1)$ and setting the result equal to 0 gives

$$\begin{aligned} \hat{f}(\omega_1) &= 1 + \left(\frac{P[\{\omega_1\}]}{2\lambda P[\{\omega_2\}]}\right)^3 \left(\frac{Q[\{\omega_2\}]}{Q[\{\omega_1\}]}\right)^2, \\ \hat{f}(\omega_2) &= \frac{1 - Q[\{\omega_1\}]\hat{f}(\omega_1)}{Q[\{\omega_2\}]} = 1 - \left(\frac{P[\{\omega_1\}]}{2\lambda P[\{\omega_2\}]}\right)^3 \left(\frac{Q[\{\omega_2\}]}{Q[\{\omega_1\}]}\right). \end{aligned}$$

Plugging the candidate $\hat{f}(\omega_1)$ into the second derivatives shows that the final position $(\hat{f}(\omega_1), \hat{f}(\omega_2))$ is a local maximum. The expected utility is

$$\frac{P[\{\omega_1\}]^2 Q[\{\omega_2\}]^{\frac{2}{3}}}{4\lambda P[\{\omega_2\}] Q[\{\omega_1\}]^{\frac{2}{3}}}.$$

The same procedure for the case $f(\omega_1) > 1 > f(\omega_1)$ shows that

$$\begin{aligned} \tilde{f}(\omega_1) &= 1 - \left(\frac{Q[\{\omega_1\}]}{Q[\{\omega_2\}]}\right) \left(\frac{P[\{\omega_2\}]}{2\lambda P[\{\omega_1\}]}\right)^3, \\ \tilde{f}(\omega_2) &= \frac{1 - Q[\{\omega_1\}]\tilde{f}(\omega_1)}{Q[\{\omega_2\}]} = 1 + \left(\frac{Q[\{\omega_1\}]}{Q[\{\omega_2\}]}\right)^2 \left(\frac{P[\{\omega_2\}]}{2\lambda P[\{\omega_1\}]}\right)^3 \end{aligned}$$

is a local maximum. The expected utility is

$$\frac{P[\{\omega_2\}]^2 Q[\{\omega_1\}]^{\frac{2}{3}}}{4\lambda P[\{\omega_1\}] Q[\{\omega_2\}]^{\frac{2}{3}}}.$$

Comparing the above two local maxima and $U(1)$ resulting from the final position $f(\omega_1) = f(\omega_2) = 1$ shows that $(\hat{f}(\omega_1), \hat{f}(\omega_2)) = (1.0098, 0.9707)$ for $P[\{\omega_1\}] = 2/3$ and $Q[\{\omega_1\}] = 3/4$ is optimal.

V.10.3 Proof of Theorem 6.2

Proof. The idea for this proof is as follows. By way of contradiction, we assume that there is an optimal final position f with two states ω_n and $\omega_{n'}$ having values in C . We now rearrange the payoff in the states ω_n and $\omega_{n'}$ in a cost-efficient way (in the sense that the Q -expectation remains constant) to (strictly) increase the expected non-concave utility. In Example IV.3.4, we explicitly specified the rearrangement for a particular setting. In the present more general case, we use the strict convexity of U on the open set C to obtain the same result.

We define $a := Q[\{\omega_n\}]f(\omega_n) + Q[\{\omega_{n'}\}]f(\omega_{n'})$ and consider the function

$$h(x) := P[\{\omega_n\}]U(x) + P[\{\omega_{n'}\}]U\left(\frac{a - Q[\{\omega_n\}]x}{Q[\{\omega_{n'}\}]}\right)$$

on the open set

$$C^* := \left\{x \in C \mid \frac{a - Q[\{\omega_n\}]x}{Q[\{\omega_{n'}\}]} \in C\right\}.$$

Since U is strictly convex on C , the same holds true for h on C^* . Maximizing the convex function h on the closure of C^* gives a solution on the boundary. Thus, there is $\tilde{x} \in C^*$ such that $h(\tilde{x}) > h(f(\omega_n))$ is satisfied. We define \tilde{f} by

$$\begin{aligned}\tilde{f}(\omega_n) &= \tilde{x}, \\ \tilde{f}(\omega_{n'}) &= \frac{a - Q[\{\omega_n\}]\tilde{x}}{Q[\{\omega_{n'}\}]}, \\ \tilde{f}(\omega_{n''}) &= f(\omega_{n''}) \quad \text{for } n'' \in \{1, \dots, N\} \text{ and } n \neq n', n \neq n''.\end{aligned}$$

By construction, we have that $E_Q[\tilde{f}] = E_Q[f]$, so \tilde{f} can be generated with the same stochastic income. Moreover, we have $E[U(\tilde{f})] > E[U(f)]$ since $h(\tilde{f}(\omega_n)) > h(f(\omega_n))$. This gives a contradiction to the optimality of f . \square

V.10.4 Example 8.1

In order to show that the final position $f_1 = (\frac{14}{9}, \frac{7}{2}, 14)$ and $f_2 = (28, \frac{7}{2}, \frac{28}{19})$ together with the prices $S_0 = (S_0^0, S_0^1) = (1, 1)$ form a financial market equilibrium, we have to check feasibility of the allocation $(f_j)_{j \in \mathcal{I}}$ and optimality of f_j for the utility maximization problem of agent j . This can be easily done by verifying that the first order conditions for optimality for the agents' optimization problems are satisfied. However, since we also want to show the uniqueness of the equilibrium, we have to determine all possible equilibria.

We first solve the utility maximization problem of the agents for a general price S_0^1 and we then solve the market clearing condition to determine S_0^1 . We finally check that the optimal allocation $(f_j)_{j \in \mathcal{I}}$ forms a feasible allocation.

In order to maximize the expected utility, the agent j chooses a strategy $\vartheta_j = (\vartheta_j^0, \vartheta_j^1)$, i.e., agent j buys ϑ_j^i units of asset i subject to his budget condition. Formally, this can be described by the optimization problem

$$\begin{aligned} & \text{maximize} \quad \sum_{n=1}^3 P[\{\omega_n\}] \log(W_j(\omega_n) + \vartheta_j^0 + \vartheta_j^1 S_1^1(\omega_n)) \\ & \text{subject to} \quad \vartheta_j^0 + \vartheta_j^1 S_0^1 \leq 0 \text{ and } W_j + \vartheta_j^0 + \vartheta_j^1 S_1^1 > 0 \end{aligned}$$

for agent j . In this optimization problem, the stochastic income W_j , the price $S_0 = (S_0^0, S_0^1)$ and the probabilities are fixed. Due to the monotonicity of the log function and the positivity of asset payoffs, we can replace the inequality in the constraint $\vartheta_j^0 + \vartheta_j^1 S_0^1 \leq 0$ by equality and hence replace ϑ_j^0 by $-\vartheta_j^1 \cdot S_0^1$. This simplifies the problem to a maximization of a function depending on ϑ_j^1 . The boundary restriction $W_j + \vartheta_j^0 + \vartheta_j^1 S_1^1 = W_j + \vartheta_j^1 (S_1^1 - S_0^1) > 0$ has to be satisfied for every state n . The price S_0^1 is determined in such a way that arbitrage is excluded; it follows that $S_1^1 - S_0^1$ is both positive and negative for at least one state. We conclude that the boundary condition defines a bounded interval of possible values for ϑ_j^1 . The property $\log(0) = -\infty$ implies that a candidate that satisfies $W_j(\omega_n) + \vartheta_j^0 + \vartheta_j^1 S_1^1(\omega_n) = 0$ in at least one coordinate cannot be optimal. Hence, a solution exists and satisfies the first order conditions for optimality. Differentiating the function $\sum_{n=1}^3 P[\{\omega_n\}] \log(W_j(\omega_n) + \vartheta_j^1 (S_1^1(\omega_n) - S_0^1))$ with respect to ϑ_j^1 and setting the resulting term equal to 0 gives

$$\begin{aligned} 0 = & \frac{S_1^1(\omega_1) - S_0^1}{W_j(\omega_1) - S_0^1 \vartheta_j^1 + \vartheta_j^1 S_1^1(\omega_1)} + \frac{S_1^1(\omega_2) - S_0^1}{W_j(\omega_2) - S_0^1 \vartheta_j^1 + \vartheta_j^1 S_1^1(\omega_2)} \\ & + \frac{S_1^1(\omega_3) - S_0^1}{W_j(\omega_3) - S_0^1 \vartheta_j^1 + \vartheta_j^1 S_1^1(\omega_3)} \end{aligned} \quad (10.3)$$

for agent $j = 1, 2$. Plugging in the explicit numbers for the payoffs and the stochastic income and solving the equations for ϑ_j^1 gives multiple solutions ϑ_1^{1+} and ϑ_1^{1-} for $j = 1$ and ϑ_2^{1+} and ϑ_2^{1-} for $j = 2$ (depending on S_0^1). Thus there are four possible combinations and every combination determines an equilibrium price S_0^1 via the market clearing condition $\sum_{j \in \mathcal{I}} \vartheta_j^1 = 0$:

- *Case +/+*: The market clearing condition gives the price $S_0^1 = 1$. It follows that $\vartheta_1^1 = 1$ and $\vartheta_2^1 = -1$. We see that the boundary condition $W_j + \vartheta_j^1 (S_1^1 - S_0^1) > 0$ is satisfied. The market clearing condition holds by construction.

- *Case -/+*: The market clearing condition gives the price $S_0^1 \approx 1.2461$. It follows that $\vartheta_1^1 \approx 23.07$, and $\vartheta_1^0 = -S_0^1 \vartheta_1^1 \approx -28.7527$. This implies $W_1 + \vartheta_1^0 + \vartheta_1^1 S_1^1 < 0$ in state 1, i.e., the boundary condition is violated. Hence, this cannot be an equilibrium.
- *Case +/-*: The market clearing condition gives the price $S_0^1 \approx 0.3412$. It follows that $\vartheta_1^1 \approx -18.99$, $\vartheta_2^1 = -\vartheta_1^1$ and $\vartheta_1^0 = -S_0^1 \vartheta_1^1 \approx 6.4834$. This implies $W_1 + \vartheta_1^0 + \vartheta_1^1 S_1^1 < 0$ in state 1, i.e., the boundary condition is violated. Hence, this cannot be an equilibrium.
- *Case -/-*: The market clearing condition has no solution.

We conclude that the +/+-combination leads to the unique equilibrium.

V.10.5 Calculations for the equilibrium in Section V.8.2

We consider a setting with four states, three assets and two agents with logarithmic utility functions. Recall from Section V.8.2 that $P[\{\omega_2\}] = P[\{\omega_3\}]$. We define $p_{23} := 2P[\{\omega_2\}] = 2P[\{\omega_3\}]$. The payoff matrix of the assets is given by

$$S_1 = (S_1^0, S_1^1, S_1^2) = \begin{pmatrix} 1 & a & 0 \\ 1 & b & d \\ 1 & b & d \\ 1 & c & e \end{pmatrix}$$

for $a < 1 < b < c$ and $0 < d < e$ such that arbitrage is excluded. We now show that for each a, b, c, d, p_1, p_{23} and p_4 , we can choose e, W_1 and W_2 such that the aggregate endowment $\sum_{j=1}^2 W_j$ is equal to asset S_1^1 and $(W_j)_{j=1,2}$ and $S_0 = (S_0^0, S_0^1, S_0^2) = (1, 1, 1)$ form an equilibrium for the two agents.

Before we formally prove this claim, let us shortly explain the (heuristic) arguments for the construction of e, W_1 and W_2 . If we fix the prices of the assets to be $S_0 := (S_0^0, S_0^1, S_0^2) = (1, 1, 1)$, we can explicitly parametrize the martingale measures as convex combinations of the two extreme points

$$\underline{Q} = (q_1, q_{23}, 0, q_4) \text{ and } \overline{Q} = (q_1, 0, q_{23}, q_4)$$

where

$$\begin{aligned} q_1 &= \frac{be - cd - e + c - b + d}{ad - ae + be - cd}, \\ q_{23} &= \frac{e - ae + a - c}{ad - ae + be - cd}, \\ q_4 &= \frac{-a + b - d + ad}{ad - ae + be - cd}. \end{aligned}$$

An agent with utility U_j and stochastic income $W_j = (U_j')^{-1}(\lambda_j dQ_j/dP)$ for $Q_j = \mu_j \underline{Q} + (1 - \mu_j) \overline{Q}$ will then not trade at all since the stochastic income

is already optimal. This type of stochastic income together with the prices $S_0 = (1, 1, 1)$ thus form an equilibrium. By setting $\mu := \mu_1 := 1 - \mu_2$ and $\lambda := \lambda_1 = \lambda_2$, choosing the stochastic incomes W_1 and W_2 (of the above type) reduces to choosing the two values λ and μ . On the other hand, the required condition $W_1 + W_2 = S_1^1$ leads to three equations since the equation in state ω_2 and ω_3 are identical. Since we can also choose the parameter e , we finally have a system of three equations and three unknowns e , λ and μ , which can be solved explicitly.

These arguments and the associated calculations motivate to define

$$e := \frac{cP[\{\omega_1\}](-a + b - d + ad) - aP[\{\omega_4\}](-cd + c - b + d)}{aP[\{\omega_4\}](b - 1)},$$

$$\lambda := \frac{2P[\{\omega_4\}]}{cq_4},$$

$$\mu := \frac{1}{2} + \sqrt{\frac{1}{4} - \frac{p_{23}cq_4}{b4P[\{\omega_4\}]q_{23}}}$$

and

$$W_j(\omega_n) := (U'_j)^{-1} \left(\lambda \frac{Q_j[\{\omega_n\}]}{P[\{\omega_n\}]} \right)$$

for $Q_1 := \mu Q + (1 - \mu)\bar{Q}$ and $Q_2 := (1 - \mu)Q + \mu\bar{Q}$. Let us now formally verify that $(W_j)_{j=1,2}$ and $S_0 = (S_0^0, S_0^1, S_0^2) = (1, 1, 1)$ form an equilibrium and that $W_1 + W_2 = S_1^1$ holds. By construction, the allocation $(W_j)_{j \in \mathcal{I}}$ is feasible and we claim that W_j is optimal for (8.1) and $S_0 = (S_0^0, S_0^1, S_0^2) = (1, 1, 1)$. This can be seen as follows. If we introduce an additional (fictitious) asset in such a way that Q_j is the unique martingale measure, then the final position W_j generated by no trading at all satisfies the first order conditions for optimality for the problem

$$\max E[U_j(f)] \text{ over } f \geq 0, E_{Q_j}[f] \leq E_{Q_j}[W_j].$$

Recall that U_j is concave and satisfies the Inada conditions (4.1) and (4.2). The first order conditions for optimality are thus also sufficient for optimality. But since W_j can also be generated by (no) trading in the incomplete market, it follows that W_j is also optimal for the utility maximization problem in the incomplete market.

It remains to show that $W_1 + W_2$ is equal to S_1^1 . We start with state ω_4 . It follows from the definition of U_j and λ that

$$(U'_1)^{-1} \left(\lambda \frac{Q_1[\{\omega_4\}]}{P[\{\omega_4\}]} \right) + (U'_2)^{-1} \left(\lambda \frac{Q_2[\{\omega_4\}]}{P[\{\omega_4\}]} \right) = \frac{2P[\{\omega_4\}]}{\lambda q_4} = c. \quad (10.4)$$

For state 1, note that definition of e implies

$$\begin{aligned} cP[\{\omega_1\}](-a + b - d + ad) &= aP[\{\omega_4\}](be - cd - e + c - b + d) \\ &= aP[\{\omega_4\}](-cd + c - b + d) + eaP[\{\omega_4\}](b - 1) \end{aligned}$$

or, putting it differently,

$$c = \frac{aP[\{\omega_4\}](be - cd - e + c - b + d)}{P[\{\omega_1\}](-a + b - d + ad)} = \frac{aq_1P[\{\omega_4\}]}{P[\{\omega_1\}]q_4}. \quad (10.5)$$

Using the definition of W_j , (10.4) and (10.5) yields

$$W_1(\omega_1) + W_2(\omega_1) = \frac{2P[\{\omega_1\}]}{\lambda q_1} = \frac{2P[\{\omega_1\}]cq_4}{2P[\{\omega_4\}]q_1} = \frac{2P[\{\omega_1\}]aq_1P[\{\omega_4\}]q_4}{2P[\{\omega_4\}]q_1P[\{\omega_1\}]q_4} = a.$$

For states 2 and 3, note first that for $n = 2, 3$ we have

$$\begin{aligned} & W_1(\omega_n) + W_2(\omega_n) \\ &= (U'_1)^{-1} \left(\lambda \frac{Q_1[\{\omega_n\}]}{P[\{\omega_n\}]} \right) + (U'_2)^{-1} \left(\lambda \frac{Q_2[\{\omega_n\}]}{P[\{\omega_n\}]} \right) \\ &= \frac{\frac{p_{23}}{2\lambda}}{\mu \underline{Q}[\{\omega_n\}] + (1 - \mu) \overline{Q}[\{\omega_n\}]} + \frac{\frac{p_{23}}{2\lambda}}{(1 - \mu) \underline{Q}[\{\omega_n\}] + \mu \overline{Q}[\{\omega_n\}]} \\ &= \frac{\frac{p_{23}}{2}}{\lambda \mu q_{23}} + \frac{\frac{p_{23}}{2}}{\lambda(1 - \mu)q_{23}} \end{aligned} \quad (10.6)$$

since $\underline{Q}[\{\omega_3\}] = 0$ and $\overline{Q}[\{\omega_2\}] = 0$. It follows from the definition of μ that

$$\mu - \mu^2 = \frac{p_{23}cq_4}{b4P[\{\omega_4\}]q_{23}} = \frac{p_{23}}{2bq_{23}\lambda}.$$

This gives

$$\frac{2b\lambda q_{23}}{p_{23}} = \frac{1}{\mu - \mu^2} = \frac{1}{\mu} + \frac{1}{1 - \mu} \iff \frac{\frac{p_{23}}{2}}{\lambda \mu q_{23}} + \frac{\frac{p_{23}}{2}}{\lambda(1 - \mu)q_{23}} = b,$$

which, together with (10.6), finally yield $W_1(\omega_n) + W_2(\omega_n) = b$ for $n = 2, 3$.

Bibliography

- [1] Y. Aït-Sahalia and A.W. Lo. Nonparametric estimation of state-price densities implicit in financial asset prices. *Journal of Finance*, 53:499–547, 1998.
- [2] Y. Aït-Sahalia and A.W. Lo. Nonparametric risk management and implied risk aversion. *Journal of Econometrics*, 94:9–51, 2000.
- [3] C.D. Aliprantis and K.C. Border. *Infinite Dimensional Analysis: A Hitchhiker's Guide*. Springer, Berlin, 3rd edition, 2006.
- [4] R.J. Aumann and M. Perles. A variational problem arising in economics. *Journal of Mathematical Analysis and Applications*, 11:488–503, 1965.
- [5] M.J. Bailey, M. Olson, and P. Wonnacott. The marginal utility of income does not increase: Borrowing, lending, and Friedman–Savage gambles. *American Economic Review*, 70:372–379, 1980.
- [6] E.J. Balder. A general approach to lower semicontinuity and lower closure in optimal control theory. *SIAM Journal on Control and Optimization*, 22:570–598, 1984.
- [7] E.J. Balder and M.R. Pistorius. On an optimal consumption problem for p -integrable consumption plans. *Economic Theory*, 17:721–737, 2001.
- [8] G. Barone-Adesi, R.F. Engle, and L. Mancini. A GARCH option pricing model with filtered historical simulation. *Review of Financial Studies*, 21:1223–1258, 2008.
- [9] S. Basak and D. Makarov. Strategic asset allocation in money management. *Preprint, London Business School*. To appear in *Journal of Finance*, 2011. Available at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1336546.
- [10] B.K. Beare. Measure preserving derivatives and the pricing kernel puzzle. *Journal of Mathematical Economics*, 47:689–697, 2011.
- [11] S. Benartzi and R.H. Thaler. Myopic loss aversion and the equity premium puzzle. *Quarterly Journal of Economics*, 110:73–92, 1995.

- [12] L. Benzoni, P. Collin-Dufresne, and R.S. Goldstein. Explaining asset pricing puzzles associated with the 1987 market crash. *Journal of Financial Economics*, 101:552–573, 2011.
- [13] A.B. Berkelaar, R. Kouwenberg, and T. Post. Optimal portfolio choice under loss aversion. *Review of Economics and Statistics*, 86:973–987, 2004.
- [14] C. Bernard and P.P. Boyle. Explicit representation of cost-efficient strategies. *Preprint, University of Waterloo*, 2010. Available at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1695543.
- [15] C. Bernard and M. Ghossoub. Static portfolio choice under cumulative prospect theory. *Mathematics and Financial Economics*, 2:277–306, 2010.
- [16] S. Biagini. Expected utility maximization: Duality methods. In: R. Cont (ed.), *Encyclopedia of Quantitative Finance*, 638–645. Wiley, Chichester, 2010.
- [17] M. Bichuch and S. Sturm. Portfolio optimization under convex incentive schemes. *Preprint, Princeton University*, 2011. Available at <http://arxiv.org/abs/1109.2945>.
- [18] P. Billingsley. *Convergence of Probability Measures*. Wiley, New York, 1968.
- [19] P. Billingsley. *Probability and Measure*. Wiley, New York, 1986.
- [20] A. Bizid, E. Jouini, and P.F. Koehl. Pricing of non-redundant derivatives in a complete market. *Review of Derivatives Research*, 2:287–314, 1998.
- [21] R.R. Bliss and N. Panigirtzoglou. Option-implied risk aversion estimates. *Journal of Finance*, 59:407–446, 2004.
- [22] B. Bouchard, N. Touzi, and A. Zeghal. Dual formulation of the utility maximization problem: The case of nonsmooth utility. *Annals of Applied Probability*, 14:678–717, 2004.
- [23] D.T. Breeden and R.H. Litzenberger. Prices of state-contingent claims implicit in option prices. *Journal of Business*, 51:621–651, 1978.
- [24] D. Brigo and F. Mercurio. *Interest Rate Models: Theory and Practice*. Springer, Berlin, 2nd edition, 2006.
- [25] D.P. Brown and M.R. Gibbons. A simple econometric approach for utility-based asset pricing models. *Journal of Finance*, 40:359–381, 1985.

- [26] D.P. Brown and J.C. Jackwerth. The pricing kernel puzzle: Reconciling index option data and economic theory. In: J. A. Batten and N. Wagner (eds.), *Derivative Securities Pricing and Modelling (Contemporary Studies in Economic and Financial Analysis)*, 155–183. Emerald Group Publishing Limited, 2012.
- [27] S. Browne. Reaching goals by a deadline: Digital options and continuous-time active portfolio management. *Advances in Applied Probability*, 31:551–577, 1999.
- [28] S. Browne. Risk-constrained dynamic active portfolio management. *Management Science*, 46:1188–1199, 2000.
- [29] L. Carassus and H. Pham. Portfolio optimization for piecewise concave criteria. In: S. Ogawa (ed.), *The 8th Workshop on Stochastic Numerics*, volume 1620 of *RIMS Kôkyûroku series*, 81–108. Research Institute for Mathematical Sciences, Kyoto, 2009.
- [30] L. Carassus and M. Rásonyi. On optimal investment for a behavioural investor in multiperiod incomplete market models. *Preprint, University of Edinburgh*. To appear in *Mathematical Finance*, 2011. Available at <http://arxiv.org/abs/1107.1617>.
- [31] G. Carlier and R.A. Dana. Optimal demand for contingent claims when agents have law invariant utilities. *Mathematical Finance*, 21:169–201, 2011.
- [32] J.N. Carpenter. Does option compensation increase managerial risk appetite? *Journal of Finance*, 55:2311–2331, 2000.
- [33] F. Chabi-Yo, R. Garcia, and E. Renault. State dependence can explain the risk aversion puzzle. *Review of Financial Studies*, 21:973–1011, 2008.
- [34] E. De Giorgi and T. Hens. Making prospect theory fit for finance. *Financial Markets and Portfolio Management*, 20:339–360, 2006.
- [35] E. De Giorgi and T. Post. Second-order stochastic dominance, reward-risk portfolio selection, and the CAPM. *Journal of Financial and Quantitative Analysis*, 43:525–546, 2008.
- [36] E. De Giorgi, T. Hens, and M.O. Rieger. Financial market equilibria with cumulative prospect theory. *Journal of Mathematical Economics*, 46:633–651, 2010.
- [37] G. Deelstra, H. Pham, and N. Touzi. Dual formulation of the utility maximization problem under transaction costs. *Annals of Applied Probability*, 11:1353–1383, 2001.

- [38] F. Delbaen and W. Schachermayer. *The Mathematics of Arbitrage*. Springer, Berlin, 2006.
- [39] D. Denneberg. *Non-Additive Measure and Integral*. Kluwer Academic Publishers, Dordrecht, 1994.
- [40] K. Detlefsen, W. Härdle, and R. Moro. Empirical pricing kernels and investor preferences. *Mathematical Methods in Economics and Finance*, 3:19–48, 2010.
- [41] M. Dierkes. Do cumulative prospect theory type preferences explain asset prices? *Unpublished preprint, University of Muenster*, 2012.
- [42] P.H. Dybvig. Distributional analysis of portfolio choice. *Journal of Business*, 61:369–393, 1988.
- [43] H. Föllmer and A. Schied. *Stochastic Finance: An Introduction in Discrete Time*. Walter de Gruyter, Berlin, 3rd edition, 2011.
- [44] G. Franke, R.C. Stapleton, and M.G. Subrahmanyam. Who buys and who sells options: The role of options in an economy with background risk. *Journal of Economic Theory*, 82:89–109, 1998.
- [45] M. Friedman and L.J. Savage. The utility analysis of choices involving risk. *Journal of Political Economy*, 56:279–304, 1948.
- [46] R. Garcia, R. Luger, and E. Renault. Empirical assessment of an intertemporal option pricing model with latent variables. *Journal of Econometrics*, 116:49–83, 2003.
- [47] D. German. Corrections to the prices of derivatives due to market incompleteness. *Applied Mathematical Finance*, 18:155–187, 2011.
- [48] C. Gollier. Portfolio choices and asset prices: The comparative statics of ambiguity aversion. *Review of Economic Studies*, 78:1329–1344, 2011.
- [49] Y. Golubev, W. Härdle, and R. Timonfeev. Testing monotonicity of pricing kernels. *Preprint, Humboldt University Berlin*, 2008. Available at <http://sfb649.wiwi.hu-berlin.de/papers/pdf/SFB649DP2008-001.pdf>.
- [50] F.J. Gomes. Portfolio choice and trading volume with loss-averse investors. *Journal of Business*, 78:675–706, 2005.
- [51] M. Grith, W. Härdle, and V. Krätschmer. A microeconomic explanation of the EPK paradox. *Unpublished preprint, Humboldt University Berlin*, 2011.

- [52] L.P. Hansen and K.J. Singleton. Stochastic consumption, risk aversion, and the temporal behavior of asset returns. *Journal of Political Economy*, 91:249–265, 1983.
- [53] R. Hartley and L. Farrell. Can expected utility theory explain gambling? *American Economic Review*, 92:613–624, 2002.
- [54] H. He. Convergence from discrete- to continuous-time contingent claims prices. *Review of Financial Studies*, 3:523–546, 1990.
- [55] H. He. Optimal consumption-portfolio policies: a convergence from discrete to continuous time models. *Journal of Economic Theory*, 55:340–363, 1991.
- [56] X.D. He and X.Y. Zhou. Portfolio choice via quantiles. *Mathematical Finance*, 21:203–231, 2011.
- [57] X.D. He and X.Y. Zhou. Portfolio choice under cumulative prospect theory: An analytical treatment. *Management Science*, 57:315–331, 2011.
- [58] T. Hens and C. Reichlin. Three solutions to the pricing kernel puzzle. *Preprint, University of Zurich*. To appear in *Review of Finance*, 2012. Available at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1582888.
- [59] J.C. Jackwerth. Recovering risk aversion from option prices and realized returns. *Review of Financial Studies*, 13:433–451, 2000.
- [60] J.C. Jackwerth and M. Rubinstein. Recovering probability distributions from option prices. *Journal of Finance*, 51:1611–1631, 1996.
- [61] M. Jeanblanc, M. Yor, and M. Chesney. *Mathematical Methods for Financial Markets*. Springer, Berlin, 2009.
- [62] H. Jin, Z.Q. Xu, and X.Y. Zhou. A convex stochastic optimization problem arising from portfolio selection. *Mathematical Finance*, 18:171–183, 2008.
- [63] H. Jin, S. Zhang, and X.Y. Zhou. Behavioral portfolio selection with loss control. *Acta Mathematica Sinica*, 27:255–274, 2011.
- [64] H. Jin and X.Y. Zhou. Behavioral portfolio selection in continuous time. *Mathematical Finance*, 18:385–426, 2008.
- [65] E. Jouini. Convergence of the equilibrium prices in a family of financial models. *Finance and Stochastics*, 7:491–507, 2003.

- [66] E. Jouini and H. Kallal. Efficient trading strategies in the presence of market frictions. *Review of Financial Studies*, 14:343–369, 2001.
- [67] E. Jouini and C. Napp. Continuous time equilibrium pricing of nonredundant assets. *NYU Working Paper No. FIN-99-008*, 1999. Available at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1296393.
- [68] E. Jouini and V. Porte. Efficient trading strategies. *Preprint, Ceremade Paris*, 2005. Available at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1000208.
- [69] D. Kahneman and A. Tversky. Prospect theory: An analysis of decision under risk. *Econometrica*, 47:263–292, 1979.
- [70] I. Karatzas, J.P. Lehoczky, S.E. Shreve, and G.L. Xu. Martingale and duality methods for utility maximization in an incomplete market. *SIAM Journal on Control and Optimization*, 29:702–730, 1991.
- [71] I. Karatzas and S.E. Shreve. *Methods of Mathematical Finance*. Springer, Berlin, 1998.
- [72] C. Kardaras and G. Žitković. Stability of the utility maximization problem with random endowment in incomplete markets. *Mathematical Finance*, 21:313–333, 2011.
- [73] D. Kramkov and W. Schachermayer. The asymptotic elasticity of utility functions and optimal investment in incomplete markets. *Annals of Applied Probability*, 9:904–950, 1999.
- [74] D. Kramkov and W. Schachermayer. Necessary and sufficient conditions in the problem of optimal investment in incomplete markets. *Annals of Applied Probability*, 13:1504–1516, 2003.
- [75] D. Kramkov and M. Sirbu. Sensitivity analysis of utility-based prices and risk-tolerance wealth processes. *Annals of Applied Probability*, 16:2140–2194, 2006.
- [76] M. Kulldorff. Optimal control of favorable games with a time limit. *SIAM Journal on Control and Optimization*, 31:52–69, 1993.
- [77] K. Larsen. Optimal portfolio delegation when parties have different coefficients of risk aversion. *Quantitative Finance*, 5:503–512, 2005.
- [78] K. Larsen and H. Yu. Horizon dependence of utility optimizers in incomplete models. *Finance and Stochastics*, 16:779–801, 2012.

- [79] K. Larsen and G. Žitković. Stability of utility-maximization in incomplete markets. *Stochastic Processes and their Applications*, 117:1642–1662, 2007.
- [80] H. Levy and M. Levy. Prospect theory and mean-variance analysis. *Review of Financial Studies*, 17:1015–1041, 2004.
- [81] M. Levy and H. Levy. Prospect theory: Much ado about nothing? *Management Science*, 48:1334–1349, 2002.
- [82] M. Magill and M. Quinzii. *Theory of Incomplete Markets, Volume 1*. The MIT Press, Cambridge, Massachusetts, 1996.
- [83] A.J. McNeil, R. Frey, and P. Embrechts. *Quantitative Risk Management: Concepts, Techniques and Tools*. Princeton University Press, Princeton, N.J., 2005.
- [84] T.A. Pirvu and K. Schulze. Multi-stock portfolio optimization under prospect theory. *Mathematics and Financial Economics*, 6:337–362, 2012.
- [85] S.R. Pliska. A stochastic calculus model of continuous trading: Optimal portfolios. *Mathematics of Operations Research*, 11:371–382, 1986.
- [86] V. Polkovnichenko and F. Zhao. Probability weighting functions implied by options prices. *Preprint, University of Texas at Dallas*. To appear in *Journal of Financial Economics*, 2010. Available at <http://business.utsa.edu/finance/files/ZhaoPaper.pdf>.
- [87] D. Prelec. The probability weighting function. *Econometrica*, 66:497–527, 1998.
- [88] J.L. Prigent. *Weak Convergence of Financial Markets*. Springer, Berlin, 2003.
- [89] J. Quiggin. *Generalized Expected Utility Theory: The Rank-dependent Model*. Kluwer Academic Publishers, Amsterdam, 1993.
- [90] M. Rásonyi and A. Rodrigues. Optimal portfolio choice for a behavioural investor in continuous-time markets. *Preprint, University of Edinburgh*. To appear in *Annals of Finance*, 2011. Available at <http://arxiv.org/abs/1202.0628>.
- [91] C. Reichlin. Non-concave utility maximization with a given pricing measure. *NCCR FINRISK working paper No. 517, ETH Zurich*. To appear in *Mathematics and Financial Economics*, 2011. Available at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1940277.

- [92] C. Reichlin. Behavioural portfolio selection: Asymptotics and stability along a sequence of models. *NCCR FINRISK working paper No. 794*, ETH Zurich, 2012. Available at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2120376.
- [93] M.O. Rieger. Co-monotonicity of optimal investments and the design of structured financial products. *Finance and Stochastics*, 15:27–55, 2011.
- [94] M.O. Rieger. Optimal financial investments for non-concave utility functions. *Economic Letters*, 114:239–240, 2012.
- [95] R.T. Rockafellar. *Convex Analysis*. Princeton University Press, Princeton, N.J., 1972.
- [96] R.T. Rockafellar and R.J-B. Wets. *Variational Analysis*. Springer, Berlin, 1998.
- [97] J.V. Rosenberg and R.F. Engle. Empirical pricing kernels. *Journal of Financial Economics*, 64:341–372, 2002.
- [98] L. Rüschendorf. On the minimum discrimination information theorem. *Statistics & Decisions, Supplement Issue*, 1:263–283, 1984.
- [99] W. Schachermayer, M. Sîrbu, and E. Taflin. In which financial markets do mutual fund theorems hold true? *Finance and Stochastics*, 13:49–77, 2009.
- [100] D. Schmeidler. Integral representation without additivity. *Proceedings of the American Mathematical Society*, 97:255–261, 1986.
- [101] H. Shefrin. *A Behavioral Approach to Asset Pricing*. Elsevier/North-Holland, Amsterdam, 2nd edition, 2008.
- [102] H. Shefrin and M. Statman. Behavioral portfolio theory. *Journal of Financial and Quantitative Analysis*, 35:127–151, 2000.
- [103] P. Siorpaes. Do arbitrage-free prices come from utility maximization? *Preprint, University of Vienna*, 2012. Available at <http://arxiv.org/pdf/1207.4749.pdf>.
- [104] T. Strömberg. A note on the differentiability of conjugate functions. *Archiv der Mathematik*, 93:481–485, 2009.
- [105] A. Tversky and D. Kahneman. Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, 5: 297–323, 1992.
- [106] P.P. Wakker. *Prospect Theory: For Risk and Ambiguity*. Cambridge University Press, Cambridge, 2010.

- [107] N. Westray and H. Zheng. Constrained nonsmooth utility maximization without quadratic inf convolution. *Stochastic Processes and their Applications*, 119:1561–1579, 2009.
- [108] N. Westray and H. Zheng. Minimal sufficient conditions for a primal optimizer in nonsmooth utility maximization. *Finance and Stochastics*, 15:501–512, 2011.
- [109] Y. Xi. Existence of Arrow–Debreu equilibrium with S-shaped utility function. *Acta Mathematicae Applicatae Sinica*, 23:149–154, 2007.
- [110] A. Yatchew and W. Härdle. Nonparametric state price density estimation using constrained least squares and the bootstrap. *Journal of Econometrics*, 133:579–599, 2006.
- [111] A. Ziegler. Why does implied risk aversion smile? *Review of Financial Studies*, 20:859–904, 2007.

Curriculum Vitae

Christian Rochus August Reichlin
born April 25, 1984
Swiss citizen

Education

| | |
|---|-----------------|
| Ph.D. Studies in Mathematics, ETH Zurich | 10/2008–11/2012 |
| Master in Mathematics, ETH Zurich (Diploma with distinction) | 04/2007–07/2008 |
| Bachelor in Mathematics, ETH Zurich | 10/2003–03/2007 |
| Matura at Kantonsschule Alpenquai, Luzern | 08/1997–06/2003 |

Academic Employment

| | |
|---|------------------------------------|
| Teaching assistant in Economics, University of Zurich | 09/2009–11/2012 |
| Teaching assistant in Mathematics, ETH Zurich | 10/2008–11/2012 |
| Tutor in Mathematics, ETH Zurich | 10/2005–06/2006 09/2007–12/2007 |