Portfolio Theory

Sven Karbach
(sven@karbach.org)

Korteweg-de Vries Institute for Mathematics
University of Amsterdam

Version: March 2, 2025

Overview

These lecture notes accompany the course ”Portfolio Theory” offered at the University of Amsterdam as part of the *Master Stochastics and Financial Mathematics* program. This course provides a rigorous introduction to stochastic finance in finite discrete-time, including the pricing and hedging of European contingent claims, expected utility and theory of risk measures, as well as portfolio optimization. The only prerequisit is a course similar to ”Measure Theoretic Probability” or its equivalent. Although occasionally, we might refer to versions of the Hahn-Banach theorem, typically covered in functional analysis courses, a background in functional analysis or stochastic analysis is not strictly necessary and key results from (discrete-time) stochastic processes and analysis are provided in the appendix. The course comprises three main parts:

I. In the first part, we focus on the modelling of discrete-time financial markets, exploring the important concepts of arbitrage-free and complete markets. We also delve into their relationships with equivalent martingale measures (existence and uniqueness). The main results in this section are the first and second fundamental theorems of asset pricing. Additionally, we demonstrate that realistic financial models in finite-discrete time cannot be complete. This realization prompts the question of how to price and hedge European contingent claims in incomplete markets.

II. To address the challenges posed in the first part, we introduce expected utility theory and risk measures in the second part. These tools allow us to identify criteria beyond arbitrage-freeness to determine *optimal prices* and *optimal hedging strategies* for European contingent claims, as well as providing us with meaningful target functions for general portfolio optimization problems.

III. The third part then revolves around solving static and dynamic portfolio optimization problems. For this, we introduce a martingale method, relate the portfolio optimization to a dual measure-valued optimization problem and introduce the dynamic programming principle to solve optimal control problems numerically.

References

Numerous comprehensive books and lecture notes are available on portfolio theory and discrete-time mathematical finance. The most relevant to this course include:

- Föllmer and Schied’s ”Stochastic Finance”;
- The Lecture Notes “Portfolio Theory” by Peter Spreij;
- “Portfolio Optimization and Performance Analysis” by Prigent.

Disclaimers

- These lecture notes are intended solely for educational purposes.
- The content is subject to change and might contain errors.
- Feedback and corrections are highly appreciated.
- Parts marked by a ♣ are meant for self-studying.

Financial Markets

1 Financial Markets in Finite Discrete-Time

1.1 Basic Stochastic Concepts of Mathematical Finance

1.1.1 The Probability Space

Throughout this course, we work within a probability space $(\Omega, \mathcal{F}, \mathbb{P})$, where Ω represents the set of all **outcomes** $\omega \in \Omega$; $\mathcal{F}$ is a σ-algebra on Ω containing all **events** $A \in \mathcal{F}$; and $\mathbb{P}$ denotes a probability measure on the measurable space $(\Omega, \mathcal{F})$. We sometimes refer to $\mathbb{P}$ as the **physical probability measure** since it describes the real-world probabilities associated with the occurrence of events $A \in \mathcal{F}$. A **finite probability space** is any probability space $(\Omega, \mathcal{F}, \mathbb{P})$ for which there exists an integer $N \in \mathbb{N}$ and outcomes $\omega_1, \omega_2, \ldots, \omega_N$ such that $\Omega = \{\omega_1, \omega_2, \ldots, \omega_N\}$, and $\mathbb{P}(\{\omega_i\}) > 0$ for all $i \in \{1, 2, \ldots, N\}$.

1.1.2 Time and the Flow of Information

In these notes, we consider time to progress in discrete steps over a finite horizon. This means that there exists a $T \in \mathbb{N}$ and real numbers $0 \leq t_0 < t_1 < \ldots < t_T < \infty$ representing the **trading times**. Depending on the application, the trading times $t_0, t_1, t_2, \ldots, t_T$ could correspond to days, months, or years, as well as seconds, minutes, or hours. For notational simplicity, we use the index of the trading times instead of the trading days themselves, i.e., the set of trading times is given as $\mathbf{T} \coloneqq \{0, 1, \ldots, T\} \subseteq \mathbb{N}_0$ with finite **time horizon** $T \in \mathbb{N}$. When modeling time-dependent random phenomena, the natural chronological order implied by the flow of time must be respected. This means that an event observed yesterday reveals information that we still remember today, but today we do not necessarily know what will happen tomorrow. In probability theory, this chronological ordering of information is represented by a filtration $\mathbb{F} = (\mathcal{F}_t)_{t \in \mathbf{T}}$, with each $\mathcal{F}_t \subseteq \mathcal{F}$ for $t \in \mathbf{T}$ being a sub-σ-algebra such that $\mathcal{F}_s \subseteq \mathcal{F}_t$ for all $0 \leq s \leq t \leq T$. The filtration $\mathbb{F}$ is interpreted as a flow of information with $\mathcal{F}_t$ containing all information available up to and including the t-th trading time. Usually, we assume that: (i) $\mathcal{F}_0$ is $\mathbb{P}$-trivial, i.e., $\mathbb{P}(A) \in \{0, 1\}$ for all $A \in \mathcal{F}_0$, i.e., all $\mathcal{F}_0$-measurable random variables are constant $\mathbb{P}$-a.s.; (ii) $\mathcal{F} = \mathcal{F}_T$, i.e., all events occur in the finite time horizon. Given a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ and a filtration $\mathbb{F}$, we call the tuple $(\Omega, \mathcal{F}, \mathbb{P}, \mathbb{F})$ a **filtered probability space**.

1.1.3 Stochastic Processes and Financial Markets

An $\mathbb{R}^d$-valued stochastic process in finite discrete-time $\{0, 1, \ldots, T\}$ is any family of $\mathbb{R}^d$-valued random variables $X_0, X_1, \ldots, X_T$ defined on $(\Omega, \mathcal{F}, \mathbb{P})$. We denote a stochastic process by $(X_t)_{t \in \mathbf{T}}$ or simply X. To give meaning to time in this framework, we must relate the stochastic process with the flow of information, that is, the filtration $\mathbb{F} = (\mathcal{F}_t)_{t \in \mathbf{T}}$. A stochastic process X is said to be **adapted** to $\mathbb{F}$ if X_t is $\mathcal{F}_t$-measurable for all $t \in \mathbf{T}$ and it is called **predictable** with respect to $\mathbb{F}$ if X_t is $\mathcal{F}_{t-1}$-measurable for all $t \in \{1, \ldots, T\}$. In mathematical finance, stochastic processes are employed to describe the random evolution of financial quantities, including asset prices and dynamic trading strategies. An asset may encompass various entities, from a bank account to stocks, bonds, commodities, options, and futures, as long as the asset's price process $(S_t)_{t \in \mathbf{T}}$ is 'observable' at every trading time. In mathemtical terms, this translates to $(S_t)_{t \in \mathbf{T}}$ being adapted. Suppose there are $(d + 1) \in \mathbb{N}$ tradable assets. Given a filtered probability space $(\Omega, \mathcal{F}, \mathbb{P}, \mathbb{F})$ and an $\mathbb{R}^{d+1}$-valued asset price process $X = (S^{(0)}, S^{(1)}, \ldots, S^{(d)})$ adapted to $\mathbb{F}$, then we call the tuple $(\Omega, \mathcal{F}, \mathbb{P}, \mathbb{F}, X)$ a **financial market**.

Remark 1.1 (Interpretation of adapted and predictable processes in finance)**.** In this course, all processes under consideration will be at least adapted. An adapted process in mathematical finance may, for instance, be the price process $(S_t)_{t\in\mathbf{T}}$ of some stock. Here, adaptedness can be interpreted as follows: at each trading day $t \in \mathbf{T}$, the asset price S_t is known given the information in $\mathcal{F}_t$, which means that it can be "measured" by checking, e.g., your brokerage account. This literally translates to S_t being $\mathcal{F}_t$-measurable for all $t \in \mathbf{T}$, which is precisely the definition of adaptedness. For most risky assets, e.g., stocks or commodities, their future price is uncertain, indicating that S_t is not measurable with respect to $\mathcal{F}_{t-1}$, hence the price process of a risky asset $(S_t)_{t\in\mathbf{T}}$ is usually *not* predictable.
In contrast, if $(B_t)_{t\in\mathbf{T}}$ represents a bank account, its value can be measured at every trading day t, and is actually already foreseeable at time $t-1$ as the interest rate to be received the following trading period, say over a month, is already known today, hence the bank account process $(B_t)_{t\in\mathbf{T}}$ is usually assumed to be predictable and we say that $(B_t)_{t\in\mathbf{T}}$ is **locally risk-free**. If we would know the evolution of the interest rate for the entire period, then B_t is $\mathcal{F}_0$ measurable for all $t \in \mathbb{T}$, i.e., it is entirely riskless and deterministic.

Note here, that so far we have not imposed any specific model for the evolution of the asset price processes $\left\{(S_t^{(i)})_{t\in\mathbf{T}} \colon i = 0, 1, \ldots, d\right\}$ and for large parts of this course we will also not do so. However, in the finite discrete-time setting, the following model is popular:

Example 1.2 (The Multinomial and Cox-Ross-Rubinstein model)**.** Simple yet rather flexible discrete-time models can be constructed as *multinomial models*, in which asset prices are modeled as multiplicative cumulative processes as follows: Let $r_1, r_2, \ldots, r_T$ denote random variables representing interest rates for periods $(0,1], \ldots, (T-1,T]$, respectively, and let $R_1^{(i)}, R_2^{(i)}, \ldots, R_T^{(i)}$ for every $i = 1, 2, \ldots, d$ be sets of random variables describing the returns of the i-th risky asset with price process $(S_t^{(i)})_{t\in\mathbf{T}}$ in the same periods. We then model the bank account process B as

$$B_t \coloneqq B_0 \prod_{j=1}^{t} (1 + r_j), \quad \text{for all } t = 0, 1, \ldots, T, \tag{1}$$

with initial capital $B_0 \geq 0$ and the asset price processes $(S_t)_{t\in\mathbf{T}}$ as

$$S_t^{(i)} \coloneqq S_0^{(i)} \prod_{j=1}^{t} (1 + R_j^{(i)}), \quad \text{for all } t = 0, 1, \ldots, T, \text{ and } i = 1, 2, \ldots, d, \tag{2}$$

with $S_0^{(i)} > 0$ being the initial price of the i-th (risky) asset observed today. Moreover, we shall assume that $r_t > -1$ and $R_t^{(i)} > -1$ holds $\mathbb{P}$-almost surely for all $t = 1, 2, \ldots, T$ and $i = 1, 2, \ldots, d$. If $r_t = r$ for all $t = 1, 2, \ldots, T$ for some $r > -1$ and all $R_1^{(i)}, R_2^{(i)}, \ldots, R_T^{(i)}$ are independent and only assume finitely-many values, we call $(S_t)_{t\in\mathbf{T}}$ the **multinomial model**. If $d = 1$ and the returns are identically distributed according to $R_t^{(1)} = U$ $(U \in \mathbb{R})$ with probability $0 < p < 1$ and $R_t^{(1)} = D$ $(U > D > -1)$ with probability $1 - p$ for all $t = 1, \ldots, T$, then we call the model the **binomial** or **Cox-Ross-Rubinstein (CRR) model**, see also Section 2.5 below.

1.2 Portfolios and Trading Strategies

Throughout this section we fix a filtered probability space $(\Omega, \mathcal{F}, \mathbb{P}, \mathbb{F})$ in finite discrete-time $\mathbf{T}$ and let $d \in \mathbb{N}$ be the total number of tradable risky assets in the market (or at least the ones that we want to trade). This is exclusive the bank account, which is tradable as well and which we henceforth denote by $S^{(0)} = (S_t^{(0)})_{t\in\mathbf{T}}$. As before, it is reasonable to assume that the bank account process $S^{(0)}$ is predictable with respect to $\mathbb{F}$. Next, we denote by $S_t^{(i)}$ the price of the i-th asset at trading time $t \in \mathbf{T}$ and assume that for every $i = 1\dots, d$ and $t = 0, 1, \dots, T$ the random variable $S_t^{(i)}$ is $\mathcal{F}_t$-measurable.
We define the vector-valued asset price process as

$$X_t \coloneqq (S_t^{(0)}, S_t^{(1)}, S_t^{(2)}, \dots, S_t^{(d)})^\intercal, \quad \text{for } t = 0, 1, \dots, T.$$

The process $X = (X_t)_{t\in\mathbf{T}}$ is an $\mathbb{R}^{d+1}$-valued adapted stochastic process giving rise to the financial market $(\Omega, \mathcal{F}, \mathbb{P}, \mathbb{F}, X)$. In this market we want to trade and construct portfolios dynamically in time. This motivates the following definition:

Definition 1.3 (Trading strategies and wealth processes)**.** We call any $\mathbb{R}^{d+1}$-valued process $(\varphi_t)_{t\in\mathbf{T}}$ a **trading strategy** in the market $(\Omega, \mathcal{F}, \mathbb{P}, \mathbb{F}, X)$, if it is predictable with respect to $\mathbb{F}$. The $(d+1)$-components of a trading strategy φ will be denoted by $\varphi^{(i)}$ for $i = 0, 1, \dots, d$, i.e.

$$\varphi_t = (\varphi_t^{(0)}, \varphi_t^{(1)}, \dots, \varphi_t^{(d)})^\intercal, \quad \text{for } t = 0, 1, \dots, T. \tag{3}$$

For a trading strategy φ we define the associated **wealth process** $(W_t(\varphi))_{t\in\mathbf{T}}$ as

$$W_t(\varphi) \coloneqq \varphi_t^\intercal S_t = \sum_{i=0}^{d} \varphi_t^{(i)} S_t^{(i)}, \quad \text{for } t = 0, 1, \dots, T. \tag{4}$$

A trading strategy φ describes a dynamically evolving portfolio consisting of the $d+1$ assets available for trade, i.e., d risky assets with price processes $S^{(1)}, S^{(2)}, \dots, S^{(d)}$ and the bank account $S^{(0)}$. For every fixed trading day $t \in \mathbf{T}$, the random variable $\varphi_t^{(0)}$ describes the quantity (measured in some unit or currency) held in the bank account at time t, whereas $\varphi_t^{(i)}$ represents the quantity we hold of the i-th asset at time t. Intuitively, a trading strategy must be predictable, as the portfolio allocation for the trading day t is set up (hence known) at trading day $t-1$ and $W_t(\varphi)$, i.e., the wealth at time t using strategy φ, is the value of our portfolio *before* we adjust the portfolio at trading day t.

It is noteworthy that our portfolio weights at time t following strategy φ are $\mathbb{R}^{d+1}$-valued. Indeed, our discourse will navigate within the bounds of an idealized market framework, more formally known as a *frictionless market*. In a frictionless market, we assume that: (i) the buy price is equal to the sell price for any asset (no bid-ask spread, no exchange commission and no taxation), which justifies the use of a single price process for every asset; (ii) we can lend and borrow capital for the same interest rate, which justifies the use of a single interest rate r; (iii) we are able to buy and sell all tradable assets in arbitrary large or small fractions, and (iv) we can purchase negative quantities of an asset (called *shortselling*), hence φ is $\mathbb{R}^{d+1}$-valued; (v) our trading does not affect the asset price, i.e., the price process X in a frictionless market is *exogenous*. Due to the importance of these conditions we give it an own definition.

Definition 1.4 (Frictionless market)**.** We say that the financial market $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P}, X)$ is **frictionless**, if the following conditions are satisfied:

i) No transaction costs;

ii) Money can be lent and borrowed with the same interest rate;

iii) Assets are available in arbitrary quantities;

iv) Short-selling is possible;

v) Trading strategies do not impact prices.

Any deviation to incorporate market frictions would necessitate the imposition of additional constraints on the trading strategies under consideration. We shall, however, leave this intricacy for more advanced explorations beyond the scope of this introductory course and make the following assumption:

Assumption 1.5. As a standing assumption throughout these lecture notes, we confine our discourse to frictionless financial markets.

1.2.1 Gains and Costs of a Trading Strategy

For simplicity we assume that $\varphi_0^{(0)} B_0^{(0)} = W_0$ for some **initial capital** $W_0 \in \mathbb{R}$ and $\varphi_0^{(i)} = 0$ for all $i = 1, \dots, d$. This means at time $t = 0$, we hold $W_0 (B_0^{(0)})^{-1}$ units of our bank account and own none of the risky assets. Then at time $t = 1$, we set up the first portfolio using the initial capital W_0 and distribute this initial investment over the d risky assets and the bank account according to the vector φ_1, which is actually constant if $\mathcal{F}_0$ is trivial. The wealth process $W(\varphi)$ in (4) is a real-valued and adapted stochastic process. We need two more processes to develop the theory further:

Definition 1.6 (Gains and cost process)**.** Let φ be a trading strategy in the financial market $(\Omega, \mathcal{F}, \mathbb{P}, \mathbb{F}, X)$. We define the **gains process** $(G_t(\varphi))_{t \in \mathbf{T}}$ associated with the strategy φ as

$$G_t(\varphi) := \sum_{i=1}^{t} \varphi_i^\intercal \Delta X_i, \quad \text{for } t = 0, \dots, T, \tag{5}$$

where $\Delta X_i := X_i - X_{i-1}$ for $i = 1, 2, \dots, T$ and where we set $G_0(\varphi) = 0$ by convention. The **cost process** $(C_t(\varphi))_{t \in \mathbf{T}}$ of the strategy φ is defined as

$$C_t(\varphi) := W_t(\varphi) - G_t(\varphi), \quad \text{for } t = 0, 1, \dots, T. \tag{6}$$

The interpretation of the gains and cost processes of a trading strategy φ is relatively straightforward: The gains process describes the profit or loss (depending on its sign) that we make trading according to strategy φ, while the cost process represents the costs associated with strategy φ. Note that by definition of the wealth process (4) and the cost process (6), the initial cost C_0 is equal to the initial capital $W_0(\varphi) = W_0$.

We make one important observation here: The gains process $(G_t(\varphi))_{t\in\mathbf{T}}$ associated with a trading strategy φ can be reformulated as a **discrete stochastic integral**, which is henceforth denoted by $(\varphi \bullet X)_{t\in\mathbf{T}}$ and by definition satisfies:

$$(\varphi \bullet X)_t = \sum_{i=1}^{t} \varphi_i^\intercal \Delta X_i = \sum_{i=1}^{t}\sum_{j=1}^{d} \varphi_i^{(j)}(X_i^{(j)} - X_{i-1}^{(j)}), \quad \text{for all } t = 0, 1, \ldots, T.$$

In this course, we will collect several properties of martingale transforms in the Appendix A, but assume that the reader is familiar with the fundamental concepts of martingales, martingale transforms and discrete-time stochastic integration.

If our strategy involves further inflow of money, e.g., due to an active savings plan, then the cost process takes this into account. Note however, if the strategy may necessitate the injection of additional capital or the withdrawal of excess capital, then by our notation this must be transferred to or from an external capital pool, not our bank account $S^{(0)}$. Strategies that do not produce additional costs beyond the initial investment are important and have their own name:

Definition 1.7 (Self-financing strategy)**.** A trading strategy φ is called **self-financing** if its cost process $(C_t(\varphi))_{t\in\mathbf{T}}$ is constant over time, that is:

$$C_t(\varphi) = W_0(\varphi), \quad \forall t \in \mathbf{T}. \tag{7}$$

Remark 1.8. Note that according to (6) a trading strategy φ is self-financing, if and only if

$$W_t(\varphi) = W_0(\varphi) + G_t(\varphi), \quad \mathbb{P}\text{-a.s. for all } t = 0, 1, \ldots, T. \tag{8}$$

This means that the wealth process of a self-financing strategy, at every trading time $t \in \mathbf{T}$, is equal to the sum of the initial wealth $W_0(\varphi)$ and the gains up to time t.

1.2.2 Discounting

Let $(\Omega, \mathcal{F}, \mathbb{P}, \mathbb{F}, X)$ be a (frictionless) financial market with bank account $S^{(0)}$ such that $S_t^{(0)} > 0$ $\mathbb{P}$-almost surely for all $t \in \mathbf{T}$, and d risky assets with price processes $S^{(1)}, \ldots, S^{(d)}$. We define the **discounted price process** $\widetilde{X}$ as

$$\widetilde{X} := \frac{X}{S^{(0)}} = (1, \frac{S^{(1)}}{S^{(0)}}, \frac{S^{(2)}}{S^{(0)}}, \ldots, \frac{S^{(d)}}{S^{(0)}}). \tag{9}$$

We also write $\widetilde{S}^{(i)} = S^{(i)}/S^{(0)}$ and define the **discounted wealth process** $(\widetilde{W}_t(\varphi))_{t\in\mathbf{T}}$ as

$$\widetilde{W}_t(\varphi) := \frac{W_t(\varphi)}{S_t^{(0)}} = \varphi_t^\intercal \widetilde{X}_t, \quad \forall t \in \mathbf{T}. \tag{10}$$

In a finite discrete-time market, discounting a price process does not affect the mathematics beyond the numbers, but it is often more convenient to work with discounted quantities as bookmaking becomes simpler. Specifically, note that discounting implies expressing prices in quantities of the reference asset, i.e., relative to the bank account $S^{(0)}$ instead of absolute terms with respect to a reference currency.

Example 1.9. If, in equation (1), we let $B_0 = 1$ and assume that the interest rates $r_1, r_2, \dots, r_T$ are all deterministic and equal to some fixed $r > -1$, then the bank account $S_t^{(0)}$ satisfies $S_t^{(0)} = (1+r)^t$ for all $t = 0, 1, \dots, T$. In this case, we have $\widetilde{S}_t^{(i)} = (1+r)^{-t} S_t^{(i)}$, which represents the so called **present value** of $S_t^{(i)}$.

The gains process $\widetilde{G}(\varphi) = \varphi \bullet \widetilde{X}$ of the discounted asset price process does not depend on the bank account part $\varphi^{(0)}$ as $\Delta\widetilde{X} = (0, \Delta\widetilde{S}^{(1)}, \Delta\widetilde{S}^{(2)}, \dots, \Delta\widetilde{S}^{(d)})$ and the following lemma holds true:

Lemma 1.10. A strategy φ is self-financing if and only if $\widetilde{W}_t(\varphi) = \widetilde{W}_0(\varphi) + (\varphi \bullet \widetilde{X})_t$ holds $\mathbb{P}$-almost surely for all $t = 0, 1, \dots, T$.

Proof. This follows from Exercise 1.2. □

The following lemma examplifies the special role that the bank account process plays:

Lemma 1.11. For any $\mathbb{R}^d$-valued predictable process $(\varphi^{(1)}, \varphi^{(2)}, \dots, \varphi^{(d)})$ and $W_0 \in \mathbb{R}$, there exists a unique predictable process $\varphi^{(0)}$ such that the $\mathbb{R}^{d+1}$-valued process φ, defined as $\varphi_t := (\varphi_t^{(0)}, \varphi_t^{(1)}, \dots, \varphi_t^{(d)})$ for $t = 0, 1, \dots, T$, is a self-financing strategy with $W_0(\varphi) = W_0$.

Proof. Let $(\varphi^{(1)}, \varphi^{(2)}, \dots, \varphi^{(d)})$ be an $\mathbb{R}^d$-valued predictable process and let $W_0 \in \mathbb{R}$. We are looking for a real-valued predictable process $\varphi^{(0)}$ such that the process $\varphi := (\varphi^{(0)}, \varphi^{(1)}, \dots, \varphi^{(d)})$ is a self-financing straetgy that satisfies $W_0(\varphi) = W_0$. Note that by Lemma 1.10, any self-financing strategy φ must satisfy the identity

$$\widetilde{W}_t(\varphi) = \widetilde{W}_0(\varphi) + (\varphi \bullet \widetilde{X})_t \quad \mathbb{P}\text{-a.s.},$$

for all $t = 0, 1, \dots, T$, which is equivalent to

$$\varphi_t^{(0)} + \sum_{i=1}^{d} \varphi_t^{(i)} \frac{S_t^{(i)}}{S_t^{(0)}} = \widetilde{W}_0(\varphi) + \sum_{j=1}^{t} \varphi_j^{\intercal} \Big(\Delta \frac{S_j^{(0)}}{S_j^{(0)}}, \Delta \frac{S_j^{(1)}}{S_j^{(0)}}, \dots, \Delta \frac{S_j^{(d)}}{S_j^{(0)}}\Big), \quad \mathbb{P}\text{-a.s. } \forall t = 0, 1, \dots, T.$$

Now, for any $t \in \mathbf{T}$ and by rearranging some terms, we obtain

$$\begin{aligned}
\varphi_t^{(0)} &= \widetilde{W}_0(\varphi) + \sum_{j=1}^{t} (\varphi_j^{(1)}, \varphi_j^{(2)}, \dots, \varphi_j^{(d)}) \Big(\Delta \frac{S_j^{(1)}}{S_j^{(0)}}, \dots, \Delta \frac{S_j^{(d)}}{S_j^{(0)}}\Big)^{\intercal} - \sum_{i=1}^{d} \varphi_t^{(i)} \frac{S_t^{(i)}}{S_t^{(0)}} \\
&= \widetilde{W}_0(\varphi) + \sum_{j=1}^{t-1} (\varphi_j^{(1)}, \varphi_j^{(2)}, \dots, \varphi_j^{(d)}) \Big(\Delta \frac{S_j^{(1)}}{S_j^{(0)}}, \dots, \Delta \frac{S_j^{(d)}}{S_j^{(0)}}\Big)^{\intercal} + \sum_{i=1}^{d} \varphi_t^{(i)} \Big(\frac{S_t^{(i)}}{S_t^{(0)}} - \frac{S_{t-1}^{(i)}}{S_{t-1}^{(0)}}\Big) \\
&\quad - \sum_{i=1}^{d} \varphi_t^{(i)} \frac{S_t^{(i)}}{S_t^{(0)}} \\
&= \widetilde{W}_0(\varphi) + \Big((\varphi^{(1)}, \varphi^{(2)}, \dots, \varphi^{(d)})^{\intercal} \bullet (\tilde{S}^{(1)}, \tilde{S}^{(2)}, \dots, \tilde{S}^{(d)})^{\intercal}\Big)_{t-1} - \sum_{i=1}^{d} \varphi_t^{(i)} \frac{S_{t-1}^{(i)}}{S_{t-1}^{(0)}}. \qquad (11)
\end{aligned}$$

Thus for every $t \in \mathbf{T}$ and given the predictable process $(\varphi^{(1)}, \varphi^{(2)}, \ldots, \varphi^{(d)})$ and W_0, we see that the random variables on the right-hand side of (11) are all $\mathcal{F}_{t-1}$-measurbale, which implies that the process $(\varphi_t^{(0)})_{t\in\mathbf{T}}$ given by the left-hand side of (11) is $\mathcal{F}_{t-1}$-measurable and by setting $\widetilde{W}_0(\varphi) = W_0$, we therefore found the unique predictable process $(\varphi_t^{(0)})_{t\in\mathbf{T}}$ satisfying the asserted properties. □

Remark 1.12. Note that by Lemma 1.11, we may identify any strategy $(\varphi^{(1)}, \ldots, \varphi^{(d)})$ with the self-financing strategy $\varphi = (\varphi^{(0)}, \varphi^{(1)}, \ldots, \varphi^{(d)})$ satisfying $W_0(\varphi) = W_0$ for some initial capital $W_0 \in \mathbb{R}$, i.e., the initial capital W_0 and the trading $(\varphi^{(1)}, \ldots, \varphi^{(d)})$ in the d risky assets uniquely determines the trading in the bank account $\varphi^{(0)}$. We therefore often denote a self-financing straetgy φ as $\varphi = (W_0, \phi)$, where $W_0 \in \mathbb{R}$ and ϕ is an $\mathbb{R}^d$-valued predictable process from Lemma 1.11, and write $\widetilde{W}_t(\varphi) = \widetilde{W}_0 + \widetilde{G}_t(\phi)$.

1.3 Arbitrage and Equivalent Martingale Measures

In this section, we study the important financial concept of *arbitrage*, respectively *arbitrage opportunities* and the consequences of their absence in financial markets. This theoretical concept is essential for the asset pricing theory that we will develop in this course.

1.3.1 Arbitrage

We begin with the definition of an **arbitrage opportunity** and the absence of such in a financial market:

Definition 1.13. A self-financing strategy φ is called an arbitrage opportunity if the following conditions hold true:

i) the initial wealth is zero, i.e., $W_0(\varphi) = 0$ $\mathbb{P}$-a.s.,

ii) the terminal wealth is non-negative, i.e., $W_T(\varphi) \geq 0$ $\mathbb{P}$-a.s.,

iii) there is a positive probability that the terminal wealth is strictly positive, i.e., $\mathbb{P}(W_T(\varphi) > 0) > 0$.

We call a financial market $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P}, X)$ **arbitrage free**, if there exists no such arbitrage opportunity.

A first implication of the absence of arbitrage in a financial market is given by the following lemma:

Lemma 1.14 (Law of one price)**.** Let φ and ψ be two self-financing strategies such that $W_T(\varphi) \leq W_T(\psi)$ $\mathbb{P}$-almost surely. If the financial market $(\Omega, \mathcal{F}, \mathbb{P}, \mathbb{F}, X)$ is arbitrage free, then also $W_t(\varphi) \leq W_t(\psi)$ $\mathbb{P}$-almost surely for all $t = 0, 1, \ldots, T-1$. In particular in an arbitrage-free market, it holds that if $W_T(\varphi) = W_T(\psi)$ $\mathbb{P}$-almost surely, then also $W_t(\varphi) = W_t(\psi)$ $\mathbb{P}$-almost surely for all $t = 0, 1, \ldots, T-1$.

Proof. Suppose that there exists a time $t \in \{0, 1, \ldots, T-1\}$ such that $W_t(\varphi) > W_t(\psi)$ with positive probability, i.e., $\mathbb{P}(W_t(\varphi) > W_t(\psi)) > 0$.
Consider a new trading strategy $\theta = \psi - \varphi$. Notice that θ is also a self-financing strategy since both φ and ψ are self-financing. The wealth process of θ is given by $W_t(\theta) = W_t(\psi) - W_t(\varphi)$.
By our assumption, there is a positive probability that $W_t(\theta) < 0$. But since $W_T(\varphi) \leq W_T(\psi)$ $\mathbb{P}$-almost surely, we have $W_T(\theta) = W_T(\varphi) - W_T(\psi) \geq 0$ $\mathbb{P}$-almost surely.
Now, let's consider the strategy θ' that starts investing according to θ at time t, if $W_t(\varphi) > W_t(\psi)$, and does not trade afterwards, i.e., just holds the position. Note that at time t, we can afford the position $\theta = \psi - \varphi$ with zero initial capital, since we only buy it if $W_t(\varphi) > W_t(\psi)$ and in this case we have $W_t(\varphi) - W_t(\psi) > 0$ units that we invest in any asset at time t, e.g., the bank account.
Since θ is self-financing, θ' is also self-financing. Moreover, the initial wealth of θ' is zero, and its terminal wealth is non-negative, i.e., $W_T(\theta') \geq 0$. Additionally, there is a positive probability that the terminal wealth of θ' is strictly positive (it is, if the situation $W_t(\varphi) > W_t(\psi)$ happens, the $(W_t(\varphi) - W_t(\psi))S_T^{(i)}$ the price of the i-th asset that we invested in), i.e., $\mathbb{P}(W_T(\theta') > 0) > 0$. Thus θ' is an arbitrage opportunity, which contradicts the assumption that the market is arbitrage-free and our initial assumption that there exists a $t \in \mathbf{T}$ with $W_t(\varphi) > W_t(\psi)$ with positive probability must be false. □

1.3.2 Equivalent Martingale Measures

In this section, we come to a purely mathematical description of arbitrage free markets. This is closely related to the conecpt of martingales and martingale measures. First, we recall that two probability measures $\mathbb{P}$ and $\mathbb{Q}$, defined on the same measurable space $(\Omega, \mathcal{F})$, are called equivalent if for all $A \in \mathcal{F}$ we have $\mathbb{P}(A) = 0$ if and only if $\mathbb{Q}(A) = 0$. In this case we write $\mathbb{P} \sim \mathbb{Q}$ and we may define the density process $Z_t := \frac{d\mathbb{P}}{d\mathbb{Q}}|_{\mathcal{F}_t}$ for $t = 0, 1, \ldots, T$, which is a martingale with respect to $\mathbb{Q}$.

Definition 1.15. An equivalent martingale measure for the market $(\Omega, \mathcal{F}, \mathbb{P}, \mathbb{F}, X)$ is any probability measure $\mathbb{Q} \sim \mathbb{P}$ on $(\Omega, \mathcal{F}_T)$ such that the discounted price process $\widetilde{X}$ is a $\mathbb{Q}$-martingale with respect to the filtration $\mathbb{F}$.

Theorem 1.16 (Doob's system theorem ♣)**.** For a probability measure $\mathbb{Q}$, the following conditions are equivalent:

i) $\mathbb{Q}$ is a martingale measure;

ii) If $\varphi = (W_0, \phi)$ is a self-financing strategy and ϕ is bounded, then the discounted wealth process $\widetilde{W}(\varphi)$ of φ is a $\mathbb{Q}$-martingale;

iii) If $\varphi = (W_0, \phi)$ is a self-financing strategy and its discounted wealth process $\widetilde{W}(\varphi)$ satisfies $\mathbb{E}_{\mathbb{Q}}\left[\widetilde{W}_t(\varphi)^-\right] < \infty$ for all $t \in \mathbf{T}$, then $\widetilde{W}(\varphi)$ is a $\mathbb{Q}$-martingale;

iv) If $\varphi = (W_0, \phi)$ is a self-financing strategy and its discounted wealth process $\widetilde{W}(\varphi)$ satisfies $\widetilde{W}_T(\varphi) \geq 0$ $\mathbb{P}$-a.s., then $\mathbb{E}_{\mathbb{Q}}\left[\widetilde{W}_T(\varphi)\right] = W_0$.

Proof. **i) ⇒ ii):** Let $(\widetilde{W}_t(\varphi))_{t\in\mathbf{T}}$ be the discounted value process of a self-financing trading strategy $\varphi = (W_0, \phi)$ such that there exists a constant c with $|\phi_t^{(i)}| \le c$ for all $i = 0, 1, \dots, d$ and $t = 0, 1, \dots, T$. Then,

$$|\widetilde{W}_t(\varphi)| \le |\widetilde{W}_0| + \sum_{k=1}^{t} c(|\widetilde{S}_k| + |\widetilde{S}_{k-1}|).$$

Since $\mathbb{Q}$ is a martingale measure, each $|\widetilde{S}_k|$ belongs to $L^1(\mathbb{Q})$ and we have $\mathbb{E}_{\mathbb{Q}}\left[|\widetilde{W}_t(\varphi)|\right] < \infty$. Moreover, for $0 \le t \le T-1$,

$$\begin{aligned}\mathbb{E}_{\mathbb{Q}}\left[\widetilde{W}_{t+1}(\varphi)|\mathcal{F}_t\right] &= \mathbb{E}_{\mathbb{Q}}\left[\widetilde{W}_t(\varphi) + \phi_{t+1}^{\intercal}(\widetilde{S}_{t+1} - \widetilde{S}_t)|\mathcal{F}_t\right] \\ &= \widetilde{W}_t(\varphi) + \phi_{t+1}^{\intercal}\mathbb{E}_{\mathbb{Q}}\left[\widetilde{S}_{t+1} - \widetilde{S}_t|\mathcal{F}_t\right] \\ &= \widetilde{W}_t(\varphi),\end{aligned}$$

where ϕ_{t+1} is $\mathcal{F}_t$-measurable and bounded by assumption.
ii) ⇒ iii): We will show the following implication:

$$\text{If } \mathbb{E}_{\mathbb{Q}}\left[\widetilde{W}_t(\varphi)^-\right] < \infty, \text{ then } \mathbb{E}_{\mathbb{Q}}\left[\widetilde{W}_t(\varphi)|\mathcal{F}_{t-1}\right] = \widetilde{W}_{t-1}(\varphi). \tag{12}$$

Since $\mathbb{E}_{\mathbb{Q}}\left[\widetilde{W}_T(\varphi)^-\right] < \infty$ by assumption, we will then get

$$\mathbb{E}_{\mathbb{Q}}\left[\widetilde{W}_{T-1}(\varphi)^-\right] = \mathbb{E}_{\mathbb{Q}}\left[\mathbb{E}_{\mathbb{Q}}\left[\widetilde{W}_T(\varphi)|\mathcal{F}_{T-1}\right]^-\right] \le \mathbb{E}_{\mathbb{Q}}\left[\widetilde{W}_T(\varphi)^-\right] < \infty,$$

due to Jensen's inequality for conditional expectations. Repeating this argument will yield $\mathbb{E}_{\mathbb{Q}}\left[\widetilde{W}_t(\varphi)^-\right] < \infty$ and $\mathbb{E}_{\mathbb{Q}}\left[\widetilde{W}_t(\varphi)|\mathcal{F}_{t-1}\right] = \widetilde{W}_{t-1}(\varphi)$ for all $t \in \mathbf{T}$. Since $\widetilde{W}_0(\varphi) = \widetilde{W}_0$ is a finite constant, we also get $\mathbb{E}_{\mathbb{Q}}\left[\widetilde{W}_t(\varphi)\right] = \widetilde{W}_0(\varphi)$, which together with the fact that $\mathbb{E}_{\mathbb{Q}}\left[\widetilde{W}_t(\varphi)^-\right] < \infty$ implies $\widetilde{W}_t(\varphi) \in L^1(\mathbb{Q})$ for all $t \in \mathbf{T}$. Thus, the martingale property of $\widetilde{W}(\varphi)$ will follow.
To prove (12), note first that $\mathbb{E}_{\mathbb{Q}}\left[\widetilde{W}_t(\varphi)|\mathcal{F}_{t-1}\right]$ is well-defined due to our assumption $\mathbb{E}_{\mathbb{Q}}\left[\widetilde{W}_t(\varphi)^-\right] < \infty$. Next, let $\phi_t^{(a)} := \phi_t \mathbf{1}_{\{|\phi_t| \le a\}}$ for a constant $a > 0$. Then $(\phi_t^{(a)})^{\intercal}(\widetilde{S}_t - \widetilde{S}_{t-1})$ is a martingale increment by condition ii). In particular, $(\phi_t^{(a)})^{\intercal}(\widetilde{S}_t - \widetilde{S}_{t-1}) \in L^1(\mathbb{Q})$ and $\mathbb{E}_{\mathbb{Q}}\left[(\phi_t^{(a)})^{\intercal}(\widetilde{S}_t - \widetilde{S}_{t-1})|\mathcal{F}_{t-1}\right] = 0$. Hence,

$$\begin{aligned}\mathbb{E}_{\mathbb{Q}}\left[\widetilde{W}_t(\varphi)|\mathcal{F}_{t-1}\right]\mathbf{1}_{\{|\phi_t| \le a\}} &= \mathbb{E}_{\mathbb{Q}}\left[\widetilde{W}_t(\varphi)\mathbf{1}_{\{|\phi_t| \le a\}}|\mathcal{F}_{t-1}\right] - \mathbb{E}_{\mathbb{Q}}\left[(\phi_t^{(a)})^{\intercal}(\widetilde{S}_t - \widetilde{S}_{t-1})|\mathcal{F}_{t-1}\right] \\ &= \widetilde{W}_{t-1}(\varphi)\mathbf{1}_{\{|\phi_t| \le a\}}.\end{aligned}$$

By letting $a \uparrow \infty$, we obtain (12).
iii) ⇒ iv):
Since $\mathcal{F}_0 = \{\emptyset, \Omega\}$, we have

$$\widetilde{W}_0 = \mathbb{E}_{\mathbb{Q}}\left[\widetilde{W}_T(\varphi)|\mathcal{F}_0\right] = \mathbb{E}_{\mathbb{Q}}\left[\widetilde{W}_T(\varphi)\right].$$

This proves this direction.
iv) ⇒ i): To prove that $\widetilde{S}_t^{(i)} \in L^1(\mathbb{Q})$ for given $i \in \{1, \dots, d\}$ and $t \in \mathbf{T}$, consider the deterministic process φ defined by

$$\phi_s^{(i)} := \mathbf{1}_{\{s \le t\}} \quad \text{and} \quad \phi_s^{(j)} := 0 \quad \text{for} \quad j \ne i.$$

It follows from Lemma 1.11 that ϕ can be complemented with a predictable process $\varphi^{(0)}$ such that $\varphi := (\varphi^{(0)}, \phi)$ is a self-financing strategy with initial investment $W_0 = \widetilde{S}_i^{(0)}$. The corresponding value process satisfies

$$\widetilde{W}_T(\varphi) = \widetilde{W}_0 + \sum_{s=1}^{T} \varphi_s^{\intercal}(\widetilde{S}_s - \widetilde{S}_{s-1}) = X_t^{(i)} \geq 0.$$

From iv) we get

$$\mathbb{E}_{\mathbb{Q}}\left[\widetilde{S}_t^{(i)}\right] = \mathbb{E}_{\mathbb{Q}}\left[\widetilde{W}_T(\varphi)\right] = W_0 = \widetilde{S}_0^{(i)}, \tag{13}$$

which yields $\widetilde{S}_t^{(i)} \in L^1(\mathbb{Q})$.
Condition i) will follow if we can show that

$$\mathbb{E}_{\mathbb{Q}}\left[\widetilde{S}_t^{(i)}\mathbf{1}_A\right] = \mathbb{E}_{\mathbb{Q}}\left[\widetilde{S}_{t-1}^{(i)}\mathbf{1}_A\right]$$

for $i = 1, \dots, d$, $t = 1, \dots, T$ and $A \in \mathcal{F}_{t-1}$. To this end, we define a d-dimensional predictable process φ by

$$\phi_s^{(i)} := \mathbf{1}_{\{s<t\}} + \mathbf{1}_{A^c}\mathbf{1}_{\{s=t\}} \quad \text{and} \quad \phi_s^{(j)} := 0 \quad \text{for} \quad j \neq i.$$

As above, we take a predictable process $\varphi^{(0)}$ such that $\varphi := (\varphi^{(0)}, \phi)$ is a self-financing strategy with initial investment $\widetilde{W}_0' = \widetilde{S}_0^{(i)}$. Its terminal value is given by

$$\widetilde{W}_T(\varphi) = \widetilde{W}_0' + \sum_{s=1}^{T} \varphi_s^{\intercal}(\widetilde{S}_s - \widetilde{S}_{s-1}) = \widetilde{S}_t^{(i)}\mathbf{1}_{A^c} + \widetilde{S}_{t-1}^{(i)}\mathbf{1}_A \geq 0.$$

Using iv) yields

$$\widetilde{S}_0^{(i)} = \widetilde{W}_0' = \mathbb{E}_{\mathbb{Q}}\left[\widetilde{W}_T(\varphi)\right] = \mathbb{E}_{\widetilde{S}_t^{(i)}\mathbf{1}_{A^c}}\left[+\right]\mathbb{E}_{\mathbb{Q}}\left[\widetilde{S}_{t-1}^{(i)}\mathbf{1}_A\right].$$

By comparing this identity with (13), we conclude that

$$\mathbb{E}_{\mathbb{Q}}\left[\widetilde{S}_t^{(i)}\mathbf{1}_A\right] = \mathbb{E}_{\mathbb{Q}}\left[\widetilde{S}_{t-1}^{(i)}\mathbf{1}_A\right],$$

which also concludes the proof. □

The following lemma is remarkable as it establishes a connection between the purely mathematical notion of an equivalent martingale measure and the normative financial condition of arbitrage-free markets.

Lemma 1.17. Let $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P}, X)$ be a financial market. If there exists an equivalent martingale measure $\mathbb{Q}$, then $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P}, X)$ is arbitrage free.

Proof. Let φ denote a self-financing strategy with $W_0(\varphi) = 0$ and $W_T(\varphi) \geq 0$ $\mathbb{P}$-a.s., i.e. φ is a potential arbitrage opportunity if also the fourth condition $\mathbb{P}(W_T(\varphi) > 0) > 0$ would hold. However, we prove that if there exists an equivalent martingale measure $\mathbb{Q}$ of $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P}, X)$ this fourth condition can not hold. Indeed, let $\mathbb{Q}$ be a equivalent martingale measure and denote the expectation with respect to $\mathbb{Q}$ by $\mathbb{E}_{\mathbb{Q}}[\cdot]$. Then the expected terminal wealth under $\mathbb{Q}$ satisfies:

$$\mathbb{E}_{\mathbb{Q}}\left[\widetilde{W}_T(\varphi)\right] = \mathbb{E}_{\mathbb{Q}}\left[\widetilde{W}_0(\varphi) + (\varphi \bullet \widetilde{X})_T\right] = \mathbb{E}_{\mathbb{Q}}\left[(\varphi \bullet \widetilde{X})_0\right] = 0,$$

where we used the fact that $\widetilde{W}_0(\varphi) = 0$, the equivalence of $\mathbb{P}$ and $\mathbb{Q}$, the definition of the discounted wealth process as martingale transform $\widetilde{W}(\varphi) = (\varphi \bullet \widetilde{X})$, and that $\widetilde{X}$ and thus also $(\varphi \bullet \widetilde{X})$ is a $\mathbb{Q}$-martingale. Since $\mathbb{P} \sim \mathbb{Q}$, we have $\mathbb{E}_{\mathbb{P}}\left[\widetilde{W}_T(\varphi)\right] = \mathbb{E}_{\mathbb{P}}[W_T(\varphi)] = 0$. But we also know that $W_T(\varphi) \geq 0$ $\mathbb{P}$-a.s. Combining these facts, we have that $W_T(\varphi) = 0$ $\mathbb{P}$-a.s.. Therefore, $\mathbb{P}(W_T(\varphi) > 0) = 0$, which means that the third condition for an arbitrage opportunity does not hold. Hence, if there exists an equivalent martingale measure $\mathbb{Q}$, the market $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P}, X)$ is arbitrage-free. □

1.3.3 The First Fundamental Theorem of Asset Pricing

We can now state the following dynamic version of the so called *first fundamental theorem of asset pricing*, which relates the absence of arbitrage opportunities in a market to the existence of equivalent martingale measures. In the following, for a sub-σ-algebra $\mathcal{G} \subseteq \mathcal{F}$, we denote by $L^0(\Omega, \mathcal{G}, \mathbb{P})$ the set of $\mathcal{G}$-measurable random variables, identifying all $\mathbb{P}$-almost surely equal random variables.

Theorem 1.18 (First Fundamental Theorem of Asset Pricing)**.** The financial market $(\Omega, \mathcal{F}, \mathbb{P}, \mathbb{F}, X)$ is arbitrage-free if and only if there exists an equivalent martingale measure $\mathbb{Q}$ of $(\Omega, \mathcal{F}, \mathbb{P}, \mathbb{F}, X)$. In this case, there exists an equivalent martingale measure $\mathbb{Q}$ which has a bounded density $\frac{\mathrm{d}\mathbb{Q}}{\mathrm{d}\mathbb{P}}$.

Proof. If there exists an equivalent martingale measure $\mathbb{Q}$ for $(\Omega, \mathcal{F}, \mathbb{P}, \mathbb{F}, X)$, then it follows immediately from Lemma 1.17 that the market is arbitrage-free.

The reverse direction is the tricky one and we will only prove it for the case of a finite probability space $(\Omega, \mathcal{F}, \mathbb{P})$. Therefore, assume that $\Omega = \{\omega_1, \omega_2, \ldots, \omega_N\}$ for some $N \in \mathbb{N}$ and $\mathbb{P}(\{\omega_i\}) > 0$ for all $i = 1, \ldots, N$.
Suppose that the market $(\Omega, \mathcal{F}, \mathbb{P}, \mathbb{F}, X)$ is free of arbitrage. Note that this means that the set $\mathcal{G} := \left\{\widetilde{G}_T(\phi) \colon \phi \text{ is a predictable } \mathbb{R}^d - \text{valued process}\right\}$ of discounted gains satisfies

$$\mathcal{G} \cap L^{0,+}(\Omega, \mathcal{F}_T) = \{0\}, \tag{14}$$

with $L^{0,+}(\Omega, \mathcal{F}_T)$ denoting the set of all non-negative $\mathcal{F}_T$-measurable random variables. Further note that

$$\mathcal{G} = \left\{\widetilde{W}_T(\varphi) \colon \varphi = (W_0, \phi) \text{ is a self-financing strategy with } W_0 = 0\right\},$$

as every self-financing strategy φ with $\widetilde{W}_0(\varphi) = 0$ satisfies $\widetilde{W}_T(\varphi) = \widetilde{G}_T(\varphi)$ and (14) then tells us that all self-financing strategies φ with initial wealth zero and non-negative wealth at time T lead to zero terminal wealth, i.e., $W_T(\varphi) = G_T(\varphi) = 0$.
Next, note that since $\Omega = \{\omega_1, \omega_2, \ldots, \omega_N\}$, every random variable Y on $(\Omega, \mathcal{F}_T)$ is of the form $Y(\omega) = \sum_{i=1}^N y_i \mathbf{1}_{\{\omega_i\}}(\omega)$ where $y_i \in \mathbb{R}$ for $i = 1, 2, \ldots, N$. Hence, we identify every $Y \in L^0(\Omega, \mathcal{F}_T)$ with the vector $y = (y_1, y_2, \ldots, y_N)^\intercal$, where, in particular, $Y = 0$ is identified with $0 \in \mathbb{R}^N$ and $L^{0,+}(\Omega, \mathcal{F}_T)$ with the non-negative orthant $\mathbb{R}^N_+$.

Note further, that for every trading strategy φ the discounted terminal gain $\widetilde{G}_T(\varphi)$ is $\mathcal{F}_T$ measurable as well, i.e., it is also of the form

$$\widetilde{G}_T(\varphi)(\omega) = \sum_{i=1}^{N} g_i(\varphi)\mathbf{1}_{\{\omega_i\}}(\omega).$$

Using the same identification here, means that $\mathcal{G}$ is identified with the set

$$\mathcal{G}' = \{(g_1(\varphi), g_2(\varphi), \ldots, g_N(\varphi))^\intercal : \varphi \text{ is a trading strategy}\}.$$

Even if this set seems to be a little obscure, it is a linear subspace of $\mathbb{R}^N$, since the gains process is linear in φ (the discrete stochastic integral is linear in its integrand), and moreover the zero is also mapped to the zero vector in $\mathbb{R}^N$. Since, we are using the same identification we see that the property (14) is preserved and we have

$$\mathcal{G}' \cap \mathbb{R}_+^N = \{0\}.$$

Now, the following idea to 'construct' an equivalent martingale measure out of the geometric condition (14) is a rather nice idea and a similar (but more involved) approach works for the infinite probability space case:
First, define the convex, closed and bounded set

$$\mathcal{C} := \left\{ y = (y_1, y_2, \ldots, y_N) \in \mathbb{R}_+^N : \sum_{i=1}^{N} y_i = 1 \right\}. \tag{15}$$

We observe that $\mathcal{C} \subseteq \mathbb{R}_+^N$, but $0 \in \mathcal{C}^c$ and therefore also $\mathcal{C} \cap \mathcal{G}' = \emptyset$! This actually allows us to separate the compact convex set $\mathcal{C}$ from the subspace $\mathcal{G}'$ by a *separating hyperplane*, that is, there exists a $v \in \mathbb{R}^N \setminus \{0\}$ such that

$$v^\intercal g = 0 \quad \forall g \in \mathcal{G}', \tag{16}$$

$$\text{and} \quad v^\intercal f > 0 \quad \forall f \in \mathcal{C}. \tag{17}$$

Note further that the for all $i = 1, 2, \ldots, N$ the unit basis vectors $f_i = (0, \ldots, 0, 1, 0, \ldots, 0)^\intercal$ are in $\mathcal{C}$ and therefore by (17) we have $v^\intercal f_i = v_i > 0$ for all $i = 1, 2, \ldots, N$. We thus define

$$p_i := \frac{v_i}{\sum_{i=1}^{N} v_i} \quad \forall i = 1, 2, \ldots, N. \tag{18}$$

Note that for all $i = 1, 2, \ldots, N$ we have $p_i \in (0, 1)$ and $\sum_{i=1}^{N} p_i = 1$, i.e., we can interpret the p_i's as probabilities, which puts us close to the desired equivalent martingale measure. Indeed, we set $\mathbb{Q}(\{\omega_i\}) = p_i$ and note that $\mathbb{Q}$ is equivalent to $\mathbb{P}$, as both measure agrees that all ω_i have a positive probability to occure.

Next, we show that $\mathbb{Q}$ is also an martingale measure, i.e., the discounted asset price vector $\widetilde{X}$ is a martingale under $\mathbb{Q}$. This can be seen as follows: Note first, that given the definition of $\mathbb{Q}$ and dividing both sides of the identity (16) by $\sum_{i=1}^{N} v_i$, we see that it is equivalent to $\sum_{i=1}^{N} p_i g_i = 0$ or in other words $\sum_{i=1}^{N} \mathbb{Q}(\{\omega_i\}) g_i = 0$, which means $\mathbb{E}_{\mathbb{Q}}[g] = 0$ for the random variable $g = \sum_{i=1}^{N} g_i \mathbf{1}_{\{\omega_i\}}$. In particular, this holds for the $\widetilde{G}_T(\varphi) = \sum_{i=1}^{N} g_i(\varphi)\mathbf{1}_{\{\omega_i\}}$ from above and hence for any trading strategy φ we have

$$\mathbb{E}_{\mathbb{Q}}\left[\widetilde{G}_T(\varphi)\right] = 0. \tag{19}$$

Now, let $A \in \mathcal{F}_{k-1}$ be arbitrary and define the strategy $\widetilde{\varphi}_t(\omega) := \mathbf{1}_k(t)\mathbf{1}_A(\omega)f_i$, i.e., the strategy is that at time k we buy one unit of the i-th asset if also the event $A \in \mathcal{F}_{k-1}$ occurs.
This process is indeed a trading strategy, i.e., predictable, and inserting this into (19) yields:

$$\begin{aligned}
\mathbb{E}_{\mathbb{Q}}\left[\widetilde{G}_T(\widetilde{\varphi})\right] &= \mathbb{E}_{\mathbb{Q}}\left[\sum_{t=1}^{T}\widetilde{\varphi}_t^{\intercal}(\widetilde{X}_t - \widetilde{X}_{t-1})\right] \\
&= \mathbb{E}_{\mathbb{Q}}\left[\widetilde{\varphi}_k^{(i)}(\widetilde{S}_k^{(i)} - \widetilde{S}_{k-1}^{(i)})\right] \\
&= \mathbb{E}_{\mathbb{Q}}\left[\mathbf{1}_A(\widetilde{S}_k^{(i)} - \widetilde{S}_{k-1}^{(i)})\right] = 0.
\end{aligned}$$

Hence, $\mathbb{E}_{\mathbb{Q}}\left[\mathbf{1}_A\widetilde{S}_k^{(i)}\right] = \mathbb{E}_{\mathbb{Q}}\left[\mathbf{1}_A\widetilde{S}_{k-1}^{(i)}\right]$ and since $A \in \mathcal{F}_{k-1}$ was arbitrary, this proves the martingale property for the asset price process $\widetilde{S}^{(i)}$ (and $i \in \{1, 2, \ldots, d\}$ was also arbitrary) so also the discounted process $\widetilde{X}$ is a $\mathbb{Q}$-martingale. This concludes the proof as we found an equivalent martingale measure $\mathbb{Q}$ as asserted. Since the probability space is finite, the Radon Nikodym derivative is immediately bounded. □

Remark 1.19. The proof of the first fundamental theorem of asset pricing is based on the geometric description of the no-arbitrage condition (14) and the idea of separating a closed convex cone inside $\mathcal{G}$ from the positive orthant. This idea can be extended towards the general case of an infinite probability space and a version of Hahn-Banach's Separation Theorem C.3. However, for this, the ambient space of $\mathcal{G}$ should be a Banach space, i.e., one could think of considering $\mathcal{L}^p$-spaces for some $p \geq 1$.

2 Pricing and Hedging of European Contingent Claims

In this section, we consider a finite discrete-time financial market $(\Omega, \mathcal{F}, \mathbb{P}, \mathbb{F}, X)$, where $\mathbf{T} = \{0, 1, \dots, T\}$ represents the trading times with time horizon $T \in \mathbb{N}$, and $d \in \mathbb{N}$ indicates the total number of risky assets. In this setting, the multivariate asset price process $X = (S^{(0)}, S^{(1)}, \dots, S^{(d)})$ is often referred to as the price process of the *underlying securities* $S^{(i)}$ for $i = 0, 1, \dots, d$, as we also explore securities whose prices are contingent on the evolution of these underlying prices. Such securities are then called **derivatives** written on the underlyings $S^{(0)}, S^{(1)}, \dots, S^{(d)}$. Derivatives are typically contracts that specify particular transactions (dependent on the underlying securities) to be executed at predetermined prices in the future.

Suppose a derivative with price process denoted by $S^{(d+1)}$ is introduced into the arbitrage-free market $(\Omega, \mathcal{F}, \mathbb{P}, \mathbb{F}, X)$ as an additional tradable asset. Then a critical question emerges: if the market was initially arbitrage-free and the derivative is liquidly tradable, what conditions must the price process of the derivative fulfill to maintain the absence of arbitrage opportunities? To address this question, we first refine the concept of derivatives in the next section and introduce the fundamental **pricing** and **hedging problem** that we face in mathematical finance.

2.1 European Contingent Claims

In this section we introduce a particular class of contingent contracts, which are characterized by the fact that they have a fixed *expiry date*, i.e., a fixed date where the contract can be executed.

Definition 2.1 (European Contingent Claim)**.** We call any $\mathcal{F}_T$-measurable non-negative random variable H a **European contingent claim**. If a European contingent claim H is $\sigma(X_0, X_1, \dots, X_T)$-measurable, then it is considered a (European) **derivative** of X.

A European contingent claim describes the pay-off that the holder of the instrument receives at time T, which may depend on the entire information up to time T. If it solely depends on the history of the underlying price process X, it is classified as a derivative written on X. In the context of contingent claims, the *European* refers to the payoff being fixed at the terminal date T as opposed to *American contingent claims*, which can be executed at any trading time.
In the following, we present a few examples of popular derivatives in real-world financial markets. The derivatives are specified by their contingent claim H:

Example 2.2 (European call option)**.** A *European call option* written on asset i with maturity date T and strike price K grants its owner the right, but not the obligation, to purchase at time T one unit of asset i for the price K. The contingent claim of a European call option is thus given by

$$H(\omega) \triangleq \max(0, S_T^{(i)}(\omega) - K) = \left(S_T^{(i)}(\omega) - K\right)^+. \tag{20}$$

Example 2.3 (Digital Barrier option)**.** A *Digital Barrier option* written on asset i with maturity date T, barrier level B, and strike price K provides a fixed payoff if the asset price reaches or exceeds the barrier level B during the option's life. The contingent claim of a Digital Barrier option is given by

$$H(\omega) \triangleq \mathbf{1}_{\{S_t^{(i)}(\omega) \geq B \text{ for some } t \in \{0,\dots,T\}\}}(\omega). \tag{21}$$

Example 2.4 (Asian call option)**.** An *Asian call option* written on asset i with maturity date T and strike price K gives its owner the right, but not the obligation, to purchase one unit of asset i for the price K at time T, based on the average price of the asset over a specific period. The contingent claim of an Asian call option is

$$H(\omega) \triangleq \max\left(0, \frac{1}{T}\sum_{t=0}^{T} S_t^{(i)}(\omega) - K\right) = \left(\frac{1}{T}\sum_{t=0}^{T} S_t^{(i)}(\omega) - K\right)^{+}. \tag{22}$$

2.1.1 The Pricing and Hedging Problem

For the buyer, derivatives serve as powerful instruments within financial markets, offering both retail and institutional buyers the opportunity to manage future cash flows, hedge commodity exposures, speculate, and more. Simultaneously, a counterparty (often a financial institution such as a bank) must undertake the inverse role in the derivative trade, i.e., selling the derivative to the client. This trade positions the seller of the claim with the liability of $-H$ at time T, which could potentially be huge and should be insured (or in financial terms, *hedged*). Indeed, the intricacies of this interaction lead to the two following problems that we aim to solve in this course:

Pricing Problem The *asset pricing problem* inquires: given a contingent claim H with maturity date T, what determines the *fair value* of the associated derivative at any time t prior to the maturity date T? In this context, *fair value* denotes a price that precludes arbitrage opportunities in the market upon the introduction of this newly priced asset.

Hedging Problem In every derivative contract, a counterparty assumes the opposite position. The hedging problem for the seller can be formulated as follows: upon selling the contingent claim H, what measures can be taken to mitigate the risk associated with the obligation to pay the random and uncertain amount H at time T? What price should the seller of the claim charge prior to maturity?

These problems are inherently intertwined. The fundamental principle for addressing both questions involves utilizing only the underlying assets $S^{(0)}, S^{(1)}, \dots, S^{(d)}$ to construct a synthetic product that *replicates* H as accurately as possible. As this product's value is derived from the given assets, it should serve as a reliable approximation for the value of H in the absence of arbitrage. We formalize this idea in the next definition.

Definition 2.5. A European contingent claim H is called **attainable** or *replicable* if there exists a self-financing strategy φ such that the terminal wealth of this strategy satisfies $W_T(\varphi) = H$ $\mathbb{P}$-almost surely. In this case, we say that the strategy φ replicates the contingent claim H, and we refer to φ as a **replicating** or *perfect hedging strategy* for H.

2.1.2 Risk-Neutral Pricing

In this section, we show that for any attainable European contingent claim H, the asset pricing problem can be solved in arbitrage-free markets. Furthermore, a unique "fair price" can be derived for each trading day $t \in \mathbf{T}$ prior to maturity, along with a pricing formula, where fair price means that the extended market is still free-of-arbitrage after introducing the contigent claim. This is demonstrated in the following theorem, where we denote the discounted contingent claim by

$$\widetilde{H} \coloneqq \frac{H}{S_T^{(0)}}.$$

Theorem 2.6 (Risk-Neutral Pricing)**.** Consider an arbitrage-free financial market $(\Omega, \mathcal{F}, \mathbb{P}, \mathbb{F}, X)$ and let $\mathbb{Q}$ be an equivalent martingale measure for this market. Then, for any attainable European contingent claim H, its discounted unique fair value $\widetilde{W}_t^H$ at time t is given by

$$\widetilde{W}_t^H = \mathbb{E}_{\mathbb{Q}}\left[\widetilde{H} \mid \mathcal{F}_t\right] = \widetilde{W}_t(\varphi), \quad \mathbb{P}\text{-a.s.} \quad \forall t = 0, 1, \ldots, T, \tag{23}$$

where φ is any replicating strategy for H.

Proof. Let φ be a self-financing strategy that replicates the contingent claim H, so that $W_T(\varphi) = H$ $\mathbb{P}$-almost surely, and hence $\widetilde{W}_T(\varphi) = \widetilde{H}$. Since the market is arbitrage-free, the First Fundamental Theorem of Asset Pricing ensures the existence of an equivalent martingale measure $\mathbb{Q}$ under which the discounted wealth process $\widetilde{W}_t(\varphi)$ is a $\mathbb{Q}$-martingale. Therefore,

$$\widetilde{W}_t(\varphi) = \mathbb{E}_{\mathbb{Q}}\left[\widetilde{W}_T(\varphi) \mid \mathcal{F}_t\right] = \mathbb{E}_{\mathbb{Q}}\left[\widetilde{H} \mid \mathcal{F}_t\right], \quad \mathbb{P}\text{-a.s.} \quad \forall t = 0, 1, \ldots, T.$$

By the *Law of One Price* (Lemma 1.14), the fair value $\widetilde{W}_t^H$ of the claim H at time t must equal the wealth of any replicating strategy at that time, i.e.,

$$\widetilde{W}_t^H = \widetilde{W}_t(\varphi) = \mathbb{E}_{\mathbb{Q}}\left[\widetilde{H} \mid \mathcal{F}_t\right], \quad \mathbb{P}\text{-a.s.} \quad \forall t = 0, 1, \ldots, T.$$

This completes the proof. □

It is crucial to observe that W_t^H represents the fair value or price of the contingent claim H at time t, which implies that the extended market $(\Omega, \mathcal{F}, \mathbb{P}, \mathbb{F}, (X, S^{(d+1)}))$ with $S_t^{(d+1)} \coloneqq W_t^H$ for $t = 0, 1, \ldots, T$, is arbitrage-free. A significant assertion of Theorem 2.6 is that the fair price is uniquely determined by the formula (23). In other words, assuming the absence of arbitrage opportunities in a market yields a unique price for any attainable European contingent claim. This provides an answer to the pricing problem for all attainable claims. Nonetheless, a challenge remains: determining whether a contingent claim is attainable and effectively constructing corresponding replicating strategies φ.

2.1.3 Arbitrage-Free Prices

We have solved the pricing problem for any attainable European contingent claim H in Theorem 2.6, where we observed that there exists a unique fair price $\widetilde{W}_0^H$ for the claim at $t = 0$. In this section, we show that if a claim is non-attainable in an arbitrage-free market, then there exists an entire open interval of prices such that the extended market remains arbitrage-free.

Definition 2.7. A real number $\pi^H \geq 0$ is called an **arbitrage-free price** of a contingent claim H if there exists a stochastic process $S^{(d+1)}$ satisfying:

i) $S_0^{(d+1)} = \pi^H$,

ii) $S_t^{(d+1)} \geq 0$ $\mathbb{P}$-almost surely for all $t = 0, 1, \ldots, T$,

iii) $S_T^{(d+1)} = H$ $\mathbb{P}$-almost surely,

and such that the extended market $(\Omega, \mathcal{F}, \mathbb{P}, \mathbb{F}, (S^{(0)}, \ldots, S^{(d)}, S^{(d+1)}))$ is arbitrage-free. Furthermore, we denote the set of all such arbitrage-free prices of H as $\Pi(H)$.

We denote by $\mathcal{P}$ the set of all equivalent martingale measures (EMMs) of the financial market $(\Omega, \mathcal{F}, \mathbb{P}, \mathbb{F}, X)$. The following theorem is fundamental for solving the pricing problem for non-attainable contingent claims.

Theorem 2.8. Let $\Pi(H)$ denote the set of arbitrage-free prices of a European contingent claim H. Then $\Pi(H)$ is non-empty and can be expressed as

$$\Pi(H) = \left\{ \mathbb{E}_{\mathbb{Q}} \left[\widetilde{H} \right] : \mathbb{Q} \in \mathcal{P},\ \mathbb{E}_{\mathbb{Q}} \left[\widetilde{H} \right] < \infty \right\}. \tag{24}$$

Furthermore, we have $\inf \Pi(H) = \inf_{\mathbb{Q} \in \mathcal{P}} \mathbb{E}_{\mathbb{Q}} \left[\widetilde{H} \right]$ and $\sup \Pi(H) = \sup_{\mathbb{Q} \in \mathcal{P}} \mathbb{E}_{\mathbb{Q}} \left[\widetilde{H} \right]$.

Proof. Suppose $\pi^H \in \Pi(H)$. By the definition of an arbitrage-free price, the extended market $(\Omega, \mathcal{F}, \mathbb{P}, \mathbb{F}, (S^{(0)}, S^{(1)}, \ldots, S^{(d)}, S^{(d+1)}))$ is arbitrage-free. Consequently, by the First Fundamental Theorem of Asset Pricing (Theorem 1.18), there exists an equivalent martingale measure $\mathbb{Q}$ such that, for all $i = 1, 2, \ldots, d+1$, we have

$$\widetilde{S}_t^{(i)} = \mathbb{E}_{\mathbb{Q}} \left[\widetilde{S}_T^{(i)} \mid \mathcal{F}_t \right], \quad \forall t = 0, 1, \ldots, T.$$

In particular, for $i = 1, 2, \ldots, d$, the processes $(\widetilde{S}_t^{(i)})_{t=0}^T$ are $\mathbb{Q}$-martingales, indicating that $\mathbb{Q} \in \mathcal{P}$. Thus, $\mathbb{Q}$ is also an equivalent martingale measure for the original market, and we have

$$\pi^H = \widetilde{S}_0^{(d+1)} = \mathbb{E}_{\mathbb{Q}} \left[\widetilde{S}_T^{(d+1)} \mid \mathcal{F}_0 \right] = \mathbb{E}_{\mathbb{Q}} \left[\widetilde{H} \right] < \infty,$$

which implies that $\pi^H \in \left\{ \mathbb{E}_{\mathbb{Q}} \left[\widetilde{H} \right] : \mathbb{Q} \in \mathcal{P},\ \mathbb{E}_{\mathbb{Q}} \left[\widetilde{H} \right] < \infty \right\}$.
Conversely, suppose $\pi^H = \mathbb{E}_{\mathbb{Q}} \left[\widetilde{H} \right]$ for some $\mathbb{Q} \in \mathcal{P}$. Define the process

$$S_t^{(d+1)} := S_t^{(0)} \mathbb{E}_{\mathbb{Q}} \left[\widetilde{H} \mid \mathcal{F}_t \right], \quad \forall t = 0, 1, \ldots, T.$$

Then $S_0^{(d+1)} = S_0^{(0)}\mathbb{E}_{\mathbb{Q}}\left[\widetilde{H}\right] = \pi^H$, and $S_T^{(d+1)} = S_T^{(0)}\widetilde{H} = H$. The process $(\widetilde{S}_t^{(d+1)})_{t=0}^T$, where $\widetilde{S}_t^{(d+1)} = \mathbb{E}_{\mathbb{Q}}\left[\widetilde{H} \mid \mathcal{F}_t\right]$, is a $\mathbb{Q}$-martingale. Since $\mathbb{Q}$ is an EMM for the original market, and $(\widetilde{S}_t^{(d+1)})_{t=0}^T$ is a $\mathbb{Q}$-martingale, $\mathbb{Q}$ is an EMM for the extended market as well. By the First Fundamental Theorem of Asset Pricing, the extended market is arbitrage-free, and thus $\pi^H \in \Pi(H)$.
To show that $\Pi(H)$ is non-empty, note that since the market is arbitrage-free, there exists at least one EMM $\mathbb{Q} \in \mathcal{P}$. If $\mathbb{E}_{\mathbb{Q}}\left[\widetilde{H}\right] < \infty$, then $\pi^H = \mathbb{E}_{\mathbb{Q}}\left[\widetilde{H}\right]$ is an arbitrage-free price. If $\mathbb{E}_{\mathbb{Q}}\left[\widetilde{H}\right] = \infty$, we can consider truncated payoffs $\widetilde{H} \wedge n$ for $n \in \mathbb{N}$, and proceed similarly. Furthermore, the expressions for $\inf \Pi(H)$ and $\sup \Pi(H)$ follow directly from (25). □

Note that if H is attainable, then it follows from Theorem 2.6 that $\Pi(H) = \left\{W_0^H\right\}$, i.e., the set of arbitrage-free prices is a singleton, where W_0^H is given by (23). In sharp contrast, the following theorem asserts that the set of arbitrage-free prices for a non-attainable contingent claim is an open interval.

2.1.4 Arbitrage-Free Prices

We have solved the pricing problem for any attainable European contingent claim H in Theorem 2.6, where we observed that there exists a unique fair price $\widetilde{W}_0^H$ for the claim at $t = 0$. In this section, we show that if a claim is non-attainable in an arbitrage-free market, then there exists an entire open interval of prices such that the extended market remains arbitrage-free.

Definition 2.9. A real number $\pi^H \geq 0$ is called an **arbitrage-free price** of a contingent claim H if there exists a stochastic process $S^{(d+1)}$ satisfying

i) $S_0^{(d+1)} = \pi^H$,

ii) $S_t^{(d+1)} \geq 0$ $\mathbb{P}$-almost surely for all $t = 0, 1, \ldots, T$,

iii) $S_T^{(d+1)} = H$ $\mathbb{P}$-almost surely,

and such that the extended market $(\Omega, \mathcal{F}, \mathbb{P}, \mathbb{F}, (S^{(0)}, \ldots, S^{(d)}, S^{(d+1)}))$ is arbitrage-free. Furthermore, we denote the set of all such arbitrage-free prices of H as $\Pi(H)$.

Remark 2.10. Note that the definition of an arbitrage-free price assumes that the extended market $(\Omega, \mathcal{F}, \mathbb{P}, \{\mathcal{F}_t\}_{t\geq 0}, (S^{(0)}, \ldots, S^{(d)}, S^{(d+1)}))$ is still a financial market. By our convention of frictionless markets, we implicitly assume that we can trade the derivative $S^{(d+1)}$ at any trading time in arbitrary quantities without transaction costs or market impact. This assumption is valid for plain vanilla options (such as European put and call options) on large, liquid underlyings like the S&P 500 index, where options markets are deep and liquid. However, for less liquid options, this arbitrage-free pricing criterion is not directly applicable because the derivative cannot be traded freely in arbitrary quantities without affecting the market price. In such cases, the assumption of frictionless trading breaks down. Despite this, the arbitrage-free pricing methodology is still used in practice as an approximation.

We denote by $\mathcal{P}$ the set of all equivalent martingale measures (EMMs) of the financial market $(\Omega, \mathcal{F}, \mathbb{P}, \mathbb{F}, X)$. The following theorem is fundamental for solving the pricing problem for non-attainable contingent claims.

Theorem 2.11. Let $\Pi(H)$ denote the set of arbitrage-free prices of a European contingent claim H. Then $\Pi(H)$ is non-empty and can be expressed as

$$\Pi(H) = \left\{ \mathbb{E}_{\mathbb{Q}} \left[\widetilde{H} \right] : \mathbb{Q} \in \mathcal{P},\ \mathbb{E}_{\mathbb{Q}} \left[\widetilde{H} \right] < \infty \right\}. \tag{25}$$

Furthermore, we have $\inf \Pi(H) = \inf_{\mathbb{Q} \in \mathcal{P}} \mathbb{E}_{\mathbb{Q}} \left[\widetilde{H} \right]$ and $\sup \Pi(H) = \sup_{\mathbb{Q} \in \mathcal{P}} \mathbb{E}_{\mathbb{Q}} \left[\widetilde{H} \right]$.

Proof. Suppose $\pi^H \in \Pi(H)$. By the definition of an arbitrage-free price π^H, the market $(\Omega, \mathcal{F}, \mathbb{P}, \mathcal{F}, (S^{(0)}, S^{(1)}, \dots, S^{(d)}, S^{(d+1)}))$ is arbitrage-free. Consequently, by the First Fundamental Theorem of Asset Pricing (Theorem 1.18), there exists an equivalent martingale measure $\mathbb{Q}$ such that for all $i = 1, 2, \dots, d+1$, we have

$$\widetilde{S}_t^{(i)} = \mathbb{E}_{\mathbb{Q}} \left[\widetilde{S}_T^{(i)} | \mathcal{F}_t \right], \quad \forall t = 0, 1, \dots, T.$$

In particular, when $i = 1, 2, \dots, d$, we observe that $(\widetilde{S}_t^{(i)})_{t \in \mathbf{T}}$ are $\mathbb{Q}$-martingales, indicating that $\mathbb{Q} \in \mathcal{P}$. Thus, $\mathbb{Q}$ is also an equivalent martingale measure for the original market and we have

$$\pi^H = \mathbb{E}_{\mathbb{Q}} \left[\widetilde{S}_T^{(d+1)} | \mathcal{F}_0 \right] = \mathbb{E}_{\mathbb{Q}} \left[\widetilde{H} | \mathcal{F}_0 \right] = \mathbb{E}_{\mathbb{Q}} \left[\widetilde{H} \right] < \infty,$$

which implies that $\pi^H \in \left\{ \mathbb{E}_{\mathbb{Q}} \left[\widetilde{H} \right] : \mathbb{Q} \in \mathcal{P} \wedge \mathbb{E}_{\mathbb{Q}} \left[\widetilde{H} \right] < \infty \right\}$.

Conversely, if $\pi^H = \mathbb{E}_{\mathbb{Q}} \left[\widetilde{H} \right]$ for some $\mathbb{Q} \in \mathcal{P}$, we define

$$X_t^{(d+1)} := \mathbb{E}_{\mathbb{Q}} \left[\widetilde{H} | \mathcal{F}_t \right], \quad \text{for } t = 0, 1, \dots, T.$$

Note that the so defined process $(X_t^{(d+1)})_{t \in \mathbf{T}}$ satisfies all the requirements of Defintion 2.9. In particular, it is a $\mathbb{Q}$-martingale, which means that $\mathbb{Q}$ is also an equivalent martingale measure for the extended market $(\Omega, \mathcal{F}, \mathbb{P}, \mathbb{F}, (S^{(0)}, S^{(1)}, \dots, S^{(d)}, S^{(d+1)}))$. Hence, by the FFTAP, the market is arbitrage-free and $\pi^H \in \Pi(H)$.
To demonstrate that $\Pi(H)$ is non-empty, let $\mathbb{P} \sim \widetilde{\mathbb{P}}$ such that $\mathbb{E}_{\widetilde{P}} \left[\widetilde{H} \right] < \infty$, e.g, we could choose $\mathrm{d}\widetilde{P} = c(1 + \widetilde{H})^{-1} \, \mathrm{d}\mathbb{P}$ for some normalizing constant c. If the original market is arbitrage-free under $\mathbb{P}$, then it is also arbitrage-free under $\widetilde{\mathbb{P}}$. Therefore, by the Fundamental Theorem of Asset Pricing, there exists an equivalent martingale measure $\mathbb{Q} \sim \widetilde{\mathbb{P}}$ such that the Radon-Nikodym derivative $\frac{\mathrm{d}\mathbb{Q}}{\mathrm{d}\widetilde{\mathbb{P}}}$ is bounded, and we have

$$\mathbb{E}_{\mathbb{Q}} \left[\widetilde{H} \right] = \mathbb{E}_{\widetilde{\mathbb{P}}} \left[\widetilde{H} \frac{\mathrm{d}\mathbb{Q}}{\mathrm{d}\widetilde{\mathbb{P}}} \right] < \infty.$$

This confirms the non-emptiness of $\Pi(H)$.
That $\inf \Pi(H) = \inf_{\mathbb{Q} \in \mathcal{P}} \mathbb{E}_{\mathbb{Q}} \left[\widetilde{H} \right]$ follows immediately from (25) and $\Pi(H) \neq \emptyset$.
To establish the upper bound, we need to show that if there exists $\mathbb{Q} \in \mathcal{P}$ such that $\mathbb{E}_{\mathbb{Q}} \left[\widetilde{H} \right] = \infty$, then for all $R > 0$ there exists an arbitrage-free price $\pi^H \in \Pi(H)$ with

$\pi^H > R$. To this end, let $n \in \mathbb{N}$ be such that $\widetilde{\pi}_n := \mathbb{E}_{\mathbb{Q}}\left[\widetilde{H} \wedge n\right] > R$ (which exists by monotone convergence) and define

$$X_t^{(d+1)} := \mathbb{E}_{\mathbb{Q}}\left[\widetilde{H} \wedge n | \mathcal{F}_t\right], \quad \text{for } t = 0, 1, \dots, T.$$

Then $X_0^{(d+1)} = \widetilde{\pi}_n$ and $\mathbb{Q}$ is an equivalent martingale measure for the extended market $(\Omega, \mathcal{F}, \mathbb{P}, \mathbb{F}, (S^{(0)}, S^{(1)}, \dots, S^{(d)}, S^{(d+1)}))$. By the FFTAP, it is arbitrage-free. In particular, by the first part of this theorem (applied to the extended market), there exists an $\mathbb{Q}' \sim \mathbb{Q}$ such that $\widetilde{\pi}_n = \mathbb{E}_{\mathbb{Q}'}\left[\widetilde{H}\right] < \infty$. However, $\mathbb{Q}'$ is also an equivalent martingale measure for the original market $(\Omega, \mathcal{F}, \mathbb{P}, \mathbb{F}, (S^{(0)}, S^{(1)}, \dots, S^{(d)}))$. Thus, $\widetilde{\pi} := \mathbb{E}_{\mathbb{Q}'}\left[\widetilde{H}\right]$ is also an aribtrage-free price for the original market, and hence

$$\pi^H = \mathbb{E}_{\mathbb{Q}'}\left[\widetilde{H}\right] \geq \mathbb{E}_{\mathbb{Q}'}\left[\widetilde{H} \wedge n\right] = \mathbb{E}_{\mathbb{Q}'}\left[X_T^{(d+1)}\right] = X_0^{(d+1)} = \widetilde{\pi}_n > R,$$

which yields a price π^H with the desired properties.

□

Note that if H is attainable, then it follows from Theorem 2.6 that $\Pi(H) = \left\{W_0^H\right\}$, i.e., the set of arbitrage-free prices is a singleton, where W_0^H is given by (23). In sharp contrast, the following theorem asserts that the set of arbitrage-free prices for a non-attainable contingent claim is an open interval.

Theorem 2.12. For any European contingent claim H, if H is not attainable, then $\Pi(H)$ is an open interval, i.e.,

$$\Pi(H) = (\inf \Pi(H), \sup \Pi(H)),$$

and moreover, $\inf \Pi(H) < \sup \Pi(H)$.

Proof. First, recall that

$$\Pi(H) = \left\{\mathbb{E}_{\mathbb{Q}}\left[\widetilde{H}\right] \mid \mathbb{Q} \in \mathcal{P},\ \mathbb{E}_{\mathbb{Q}}\left[\widetilde{H}\right] < \infty\right\},$$

where $\widetilde{H} = H/S_T^{(0)}$ is the discounted contingent claim.
Since $\mathcal{P}$ is convex and the mapping $\mathbb{Q} \mapsto \mathbb{E}_{\mathbb{Q}}\left[\widetilde{H}\right]$ is linear, it follows that $\Pi(H)$ is an interval. We will show that $\Pi(H)$ is open by constructing, for any $\pi \in \Pi(H)$, arbitrage-free prices $\check{\pi}$ and $\hat{\pi}$ such that $\check{\pi} < \pi < \hat{\pi}$.
Let $\mathbb{Q} \in \mathcal{P}$ be such that $\pi = \mathbb{E}_{\mathbb{Q}}\left[\widetilde{H}\right]$. Define the process

$$U_t := \mathbb{E}_{\mathbb{Q}}\left[\widetilde{H} \mid \mathcal{F}_t\right], \quad t = 0, 1, \dots, T.$$

Then U_t is a $\mathbb{Q}$-martingale with $U_0 = \pi$ and $U_T = \widetilde{H}$. We can write

$$\widetilde{H} = U_0 + \sum_{t=1}^{T} (U_t - U_{t-1}).$$

Since H is not attainable, there must exist some $t \in \{1, 2, \dots, T\}$ such that $U_t - U_{t-1} \notin \mathcal{K}_t$, where

$$\mathcal{K}_t := \left\{\eta^\intercal\left(\widetilde{S}_t - \widetilde{S}_{t-1}\right) \mid \eta \in L^0\left(\Omega, \mathcal{F}_{t-1}, \mathbb{Q}; \mathbb{R}^d\right)\right\}.$$

Here, $L^0\left(\Omega, \mathcal{F}_{t-1}, \mathbb{Q}; \mathbb{R}^d\right)$ denotes the set of $\mathbb{R}^d$-valued, $\mathcal{F}_{t-1}$-measurable random variables. Note that $\mathcal{K}_t$ is a closed linear subspace of $L^1(\Omega, \mathcal{F}_t, \mathbb{Q})$. By the Hahn-Banach theorem, there exists a $Z \in L^\infty(\Omega, \mathcal{F}_t, \mathbb{Q})$ such that

$$\mathbb{E}_{\mathbb{Q}}[WZ] = 0 \quad \text{for all } W \in \mathcal{K}_t, \quad \text{and} \quad \mathbb{E}_{\mathbb{Q}}[(U_t - U_{t-1}) Z] > 0.$$

Without loss of generality, we can assume that $\|Z\|_\infty \leq \frac{1}{3}$. Define

$$\hat{Z} \coloneqq 1 + Z - \mathbb{E}_{\mathbb{Q}}[Z \mid \mathcal{F}_{t-1}].$$

Then $\hat{Z}$ is positive and satisfies $\mathbb{E}_{\mathbb{Q}}\left[\hat{Z} \mid \mathcal{F}_{t-1}\right] = 1$. Therefore, $\hat{Z}$ can be used to define a new probability measure $\hat{\mathbb{Q}}$ equivalent to $\mathbb{Q}$ via

$$\left.\frac{\mathrm{d}\hat{\mathbb{Q}}}{\mathrm{d}\mathbb{Q}}\right|_{\mathcal{F}_T} = \hat{Z}.$$

We now compute

$$\begin{aligned}
\mathbb{E}_{\hat{\mathbb{Q}}}\left[\widetilde{H}\right] &= \mathbb{E}_{\mathbb{Q}}\left[\widetilde{H}\hat{Z}\right] \\
&= \mathbb{E}_{\mathbb{Q}}\left[\left(U_0 + \sum_{s=1}^{T}(U_s - U_{s-1})\right)\hat{Z}\right] \\
&= U_0 \mathbb{E}_{\mathbb{Q}}\left[\hat{Z}\right] + \sum_{s=1}^{T}\mathbb{E}_{\mathbb{Q}}\left[(U_s - U_{s-1})\hat{Z}\right] \\
&= U_0 + \sum_{s=1}^{T}\mathbb{E}_{\mathbb{Q}}\left[(U_s - U_{s-1})\hat{Z}\right],
\end{aligned}$$

since $\mathbb{E}_{\mathbb{Q}}\left[\hat{Z}\right] = 1$.
For $s \neq t$, note that $\hat{Z}$ is $\mathcal{F}_t$-measurable, and $U_s - U_{s-1}$ is $\mathcal{F}_s$-measurable, with $s \neq t$. Therefore, for $s \neq t$, $\mathbb{E}_{\mathbb{Q}}\left[(U_s - U_{s-1})\hat{Z}\right] = \mathbb{E}_{\mathbb{Q}}[U_s - U_{s-1}]\,\mathbb{E}_{\mathbb{Q}}\left[\hat{Z}\right] = 0$.
Thus,

$$\begin{aligned}
\mathbb{E}^{\hat{\mathbb{Q}}}\left[\widetilde{H}\right] &= U_0 + \mathbb{E}_{\mathbb{Q}}\left[(U_t - U_{t-1})\hat{Z}\right] \\
&= \pi + \mathbb{E}_{\mathbb{Q}}[(U_t - U_{t-1})(1 + Z - \mathbb{E}_{\mathbb{Q}}[Z \mid \mathcal{F}_{t-1}])] \\
&= \pi + \mathbb{E}_{\mathbb{Q}}[(U_t - U_{t-1}) Z] - \mathbb{E}_{\mathbb{Q}}[(U_t - U_{t-1})\,\mathbb{E}_{\mathbb{Q}}[Z \mid \mathcal{F}_{t-1}]].
\end{aligned}$$

Since U_{t-1} is $\mathcal{F}_{t-1}$-measurable, we have

$$\mathbb{E}_{\mathbb{Q}}[(U_t - U_{t-1})\,\mathbb{E}_{\mathbb{Q}}[Z \mid \mathcal{F}_{t-1}]] = \mathbb{E}_{\mathbb{Q}}[\mathbb{E}_{\mathbb{Q}}[U_t - U_{t-1} \mid \mathcal{F}_{t-1}]\,\mathbb{E}_{\mathbb{Q}}[Z \mid \mathcal{F}_{t-1}]] = 0,$$

because U_t is a $\mathbb{Q}$-martingale.
Therefore,

$$\mathbb{E}_{\hat{\mathbb{Q}}}\left[\widetilde{H}\right] = \pi + \mathbb{E}_{\mathbb{Q}}[(U_t - U_{t-1}) Z] > \pi,$$

since $\mathbb{E}_{\mathbb{Q}}[(U_t - U_{t-1}) Z] > 0$ by construction.
Moreover, since $\|Z\|_\infty \leq \frac{1}{3}$, we have

$$1 - \frac{2}{3} \leq \hat{Z} \leq 1 + \frac{1}{3} - \left(-\frac{1}{3}\right) = \frac{5}{3},$$

so the Radon-Nikodym derivative $\mathrm{d}\hat{\mathbb{Q}}/\mathrm{d}\mathbb{Q}$ is bounded between $\frac{2}{3}$ and $\frac{5}{3}$.
We now need to verify that $\hat{\mathbb{Q}} \in \mathcal{P}$, i.e., that under $\hat{\mathbb{Q}}$, the discounted asset price process $\widetilde{S}$ is a martingale.
For $k \neq t$, since $\hat{Z}$ is $\mathcal{F}_t$-measurable, and $\hat{Z}$ and $\widetilde{S}_k - \widetilde{S}_{k-1}$ are independent given $\mathcal{F}_{k-1}$, we have

$$\mathbb{E}_{\hat{\mathbb{Q}}}\left[\widetilde{S}_k - \widetilde{S}_{k-1} \mid \mathcal{F}_{k-1}\right] = \mathbb{E}_{\mathbb{Q}}\left[\widetilde{S}_k - \widetilde{S}_{k-1} \mid \mathcal{F}_{k-1}\right] = 0.$$

For $k = t$, we have

$$\begin{aligned}
\mathbb{E}_{\hat{\mathbb{Q}}}\left[\widetilde{S}_t - \widetilde{S}_{t-1} \mid \mathcal{F}_{t-1}\right] &= \mathbb{E}_{\mathbb{Q}}\left[\left(\widetilde{S}_t - \widetilde{S}_{t-1}\right)\hat{Z} \mid \mathcal{F}_{t-1}\right] \\
&= \mathbb{E}_{\mathbb{Q}}\left[\left(\widetilde{S}_t - \widetilde{S}_{t-1}\right)(1 + Z - \mathbb{E}_{\mathbb{Q}}[Z \mid \mathcal{F}_{t-1}]) \mid \mathcal{F}_{t-1}\right] \\
&= \mathbb{E}_{\mathbb{Q}}\left[\widetilde{S}_t - \widetilde{S}_{t-1} \mid \mathcal{F}_{t-1}\right] + \mathbb{E}_{\mathbb{Q}}\left[\left(\widetilde{S}_t - \widetilde{S}_{t-1}\right) Z \mid \mathcal{F}_{t-1}\right] \\
&\quad - \mathbb{E}_{\mathbb{Q}}\left[\widetilde{S}_t - \widetilde{S}_{t-1} \mid \mathcal{F}_{t-1}\right]\mathbb{E}_{\mathbb{Q}}[Z \mid \mathcal{F}_{t-1}] \\
&= 0 + 0 - 0 = 0,
\end{aligned}$$

where we have used the fact that $\mathbb{E}_{\mathbb{Q}}\left[\widetilde{S}_t - \widetilde{S}_{t-1} \mid \mathcal{F}_{t-1}\right] = 0$ (since $\widetilde{S}$ is a $\mathbb{Q}$-martingale) and $\mathbb{E}_{\mathbb{Q}}\left[\left(\widetilde{S}_t - \widetilde{S}_{t-1}\right) Z \mid \mathcal{F}_{t-1}\right] = 0$ (since Z is orthogonal to $\mathcal{K}_t$). Therefore, under $\hat{\mathbb{Q}}$, the discounted asset price process $\widetilde{S}$ is a martingale, so $\hat{\mathbb{Q}} \in \mathcal{P}$. Similarly, we can construct an equivalent martingale measure $\check{\mathbb{Q}}$ such that $\check{\pi} := \mathbb{E}^{\check{\mathbb{Q}}}\left[\widetilde{H}\right] < \pi$.
Define

$$\frac{\mathrm{d}\check{\mathbb{Q}}}{\mathrm{d}\mathbb{Q}} := 2 - \frac{\mathrm{d}\hat{\mathbb{Q}}}{\mathrm{d}\mathbb{Q}}.$$

Since $\frac{\mathrm{d}\hat{\mathbb{Q}}}{\mathrm{d}\mathbb{Q}}$ takes values between $\frac{2}{3}$ and $\frac{5}{3}$, the density $\frac{\mathrm{d}\check{\mathbb{Q}}}{\mathrm{d}\mathbb{Q}}$ takes values between $\frac{1}{3}$ and $\frac{4}{3}$, so $\check{\mathbb{Q}}$ is a probability measure equivalent to $\mathbb{Q}$. Furthermore, since $\mathbb{Q} \in \mathcal{P}$ and $\hat{\mathbb{Q}} \in \mathcal{P}$, it follows that $\check{\mathbb{Q}} \in \mathcal{P}$ (the set $\mathcal{P}$ is convex).
Finally, we have

$$\check{\pi} = \mathbb{E}_{\check{\mathbb{Q}}}\left[\widetilde{H}\right] = \mathbb{E}_{\mathbb{Q}}\left[\widetilde{H}\left(2 - \frac{\mathrm{d}\hat{\mathbb{Q}}}{\mathrm{d}\mathbb{Q}}\right)\right] = 2\mathbb{E}_{\mathbb{Q}}\left[\widetilde{H}\right] - \mathbb{E}_{\hat{\mathbb{Q}}}\left[\widetilde{H}\right] < \pi,$$

since $\mathbb{E}_{\hat{\mathbb{Q}}}\left[\widetilde{H}\right] > \mathbb{E}_{\mathbb{Q}}\left[\widetilde{H}\right]$.
Thus, we have found $\check{\pi} < \pi < \hat{\pi}$ with $\check{\pi}, \hat{\pi} \in \Pi(H)$, showing that $\Pi(H)$ is an open interval. Moreover, since $\Pi(H)$ is non-empty and contains more than one point, we have $\inf \Pi(H) < \sup \Pi(H)$. □

2.2 EMM Consistent with Market Prices

Theorem 2.12 tells us that for a non-attainable claim, there is an entire interval of arbitrage-free prices. Hence, speaking of *the price* of a contingent claim is not possible in such cases, and the pricing problem is not uniquely solved. In practice, one can use an equivalent martingale measure that is consistent with observed market prices of liquidly traded options.
Indeed, suppose we observe market prices of a set of European call options with different strikes K_i and expiry dates T_i, for $i = 1, \ldots, M$. Denote the observed market prices by

$\{\hat{C}(T_i, K_i) : i = 1, \ldots, M\}$. We can attempt to choose an equivalent martingale measure $\mathbb{Q}$ such that, for each i,

$$\hat{C}(T_i, K_i) = e^{-rT_i}\, \mathbb{E}_{\mathbb{Q}}\left[H_{T_i,K_i}\right], \tag{26}$$

where $H_{T_i,K_i} = (S_{T_i} - K_i)^+$ is the payoff of the European call option at time T_i, and r is the risk-free interest rate. By matching the discounted expected payoffs under $\mathbb{Q}$ to the observed option prices, we ensure that the equivalent martingale measure $\mathbb{Q}$ is consistent with the market prices of these options.

However, it might be the case that no such $\mathbb{Q}$ exists, especially if the market prices do not align perfectly with any model. In such situations, we can choose $\mathbb{Q}$ that minimizes the discrepancy between the model prices and the observed market prices. Specifically, we can solve the optimization problem:

$$\mathbb{Q}^* = \arg\min_{\mathbb{Q}\in\mathcal{P}} \sum_{i=1}^{M} \left(e^{-rT_i}\, \mathbb{E}_{\mathbb{Q}}\left[H_{T_i,K_i}\right] - \hat{C}(T_i, K_i)\right)^2, \tag{27}$$

where $\mathcal{P}$ is a set of equivalent martingale measures, often parameterized by a set of parameters in a chosen model for the asset price process $(S_t)_{t\in\mathbf{T}}$. This procedure is called **calibration to market option prices**.

Typically, the set of equivalent martingale measures $\mathcal{P}$ is parameterized by the parameters of the asset price model we are using. For example, in a stochastic volatility model like the Heston model, $\mathcal{P}$ would depend on parameters such as the initial volatility, mean-reversion rate, long-term variance level, volatility of volatility, and the correlation between the asset price and volatility processes. By adjusting these parameters, we can generate different equivalent martingale measures within the model.

Calibration involves finding the parameter values that make the model's theoretical option prices match the observed market prices as closely as possible. This is typically done using numerical optimization methods, as the relationship between the model parameters and the option prices can be complex and non-linear.

Example: Calibrating the Heston Model

Consider calibrating the Heston stochastic volatility model to market option prices. The Heston model assumes that the asset price S_t and the variance process v_t evolve according to:

$$\begin{aligned} dS_t &= rS_t\, dt + \sqrt{v_t}\, S_t\, dW_t^S, \\ dv_t &= \kappa(\theta - v_t)\, dt + \sigma_v \sqrt{v_t}\, dW_t^v, \end{aligned}$$

where W_t^S and W_t^v are Brownian motions with correlation ρ, r is the risk-free interest rate, κ is the mean-reversion rate of the variance, θ is the long-term variance, σ_v is the volatility of volatility, and v_t is the instantaneous variance.

Under the risk-neutral measure $\mathbb{Q}$, the drift of S_t is adjusted to r, ensuring that discounted asset prices are martingales. The model parameters $(\kappa, \theta, \sigma_v, \rho, v_0)$ define the dynamics under $\mathbb{Q}$.

To calibrate the model, we adjust the parameters $(\kappa, \theta, \sigma_v, \rho, v_0)$ to minimize the difference between the model prices and the observed market prices:

$$\min_{\kappa,\theta,\sigma_v,\rho,v_0} \sum_{i=1}^{M} \left(C_{\text{model}}(T_i, K_i; \kappa, \theta, \sigma_v, \rho, v_0) - \hat{C}(T_i, K_i)\right)^2, \tag{28}$$

where $C_{\text{model}}(T_i, K_i; \kappa, \theta, \sigma_v, \rho, v_0)$ denotes the theoretical price of the option computed under the Heston model with the given parameters.

2.3 Complete Markets

An arbitrage-free market in which every European contingent claim is attainable would be ideal from an efficiency perspective, as every claim would then have a unique price process. It makes sense to define such an efficient market:

Definition 2.13. An arbitrage-free market $(\Omega, \mathcal{F}, \mathbb{P}, \{\mathcal{F}_t\}_{t=0}^T, X)$ is called **complete** if every European contingent claim $H \in \mathcal{L}^{0,+}(\mathcal{F}_T)$ is attainable.

Following Definitions 2.5 and 2.13, an arbitrage-free market $(\Omega, \mathcal{F}, \mathbb{P}, \{\mathcal{F}_t\}_{t=0}^T, X)$ is complete if and only if for all $H \in \mathcal{L}^{0,+}(\mathcal{F}_T)$ there exists a self-financing trading strategy $\varphi = (W_0, \phi)$ such that

$$H = W_0 + (\phi \bullet X)_T, \quad \mathbb{P}\text{-almost surely.}$$

If the market is complete, then by risk-neutral pricing we see that

$$\widetilde{W}_t^H = \mathbb{E}_{\mathbb{Q}}\left[\widetilde{H} \mid \mathcal{F}_t\right] = \widetilde{W}_0 + (\phi \bullet \widetilde{S})_t,$$

for any $H \in \mathcal{L}^{0,+}(\mathcal{F}_T)$ and $\mathbb{Q} \in \mathcal{P}$, i.e., the $\mathbb{Q}$-martingale $(\widetilde{W}_t^H)_{t=0}^T$ can be represented as a sum of a stochastic integral $(\phi \bullet \widetilde{S})$ and a constant $\widetilde{W}_0$. It is therefore valuable to make the connection between the definition of complete markets and the concept of integral representations discussed in, e.g., Section 8 of Peter Spreij's lecture notes on "Stochastic Integration".

2.3.1 The Second Fundamental Theorem of Asset Pricing

Theorem 2.14 (Second Fundamental Theorem of Asset Pricing)**.** An arbitrage-free financial market $(\Omega, \mathcal{F}, \mathbb{P}, \mathbb{F}, X)$ is complete if and only if there exists a unique equivalent martingale measure.

Proof. Assume first that the market is complete. Then every contingent claim $H \in \mathcal{L}^{0,+}(\mathcal{F}_T)$ is attainable, and therefore by Theorem 2.6 there exists a unique fair price process $(\widetilde{W}_t^H)_{t=0}^T$ such that for any $\mathbb{Q} \in \mathcal{P}$ we have

$$\widetilde{W}_t^H = \mathbb{E}_{\mathbb{Q}}\left[\widetilde{H} \mid \mathcal{F}_t\right], \quad \text{for every } t = 0, 1, \ldots, T.$$

In particular, this holds for the claims $\widetilde{H} = \mathbf{1}_F$ for any $F \in \mathcal{F}_T$, which for $t = 0$ yields

$$\widetilde{W}_0^H = \mathbb{E}_{\mathbb{Q}}[\mathbf{1}_F \mid \mathcal{F}_0] = \mathbb{E}_{\mathbb{Q}}[\mathbf{1}_F] = \mathbb{Q}(F),$$

for any $\mathbb{Q} \in \mathcal{P}$. Since the initial price $\widetilde{W}_0^H$ is uniquely determined (due to the Law of One Price), it follows that $\mathbb{Q}(F)$ is uniquely determined for all $F \in \mathcal{F}_T$. Therefore, all equivalent martingale measures coincide, and hence there exists a unique equivalent martingale measure.

Conversely, suppose that there exists only one equivalent martingale measure $\mathbb{Q}'$. By Theorem 2.11, the set of arbitrage-free prices of any European contingent claim H is given by

$$\Pi(H) = \left\{\mathbb{E}_{\mathbb{Q}}\left[\widetilde{H}\right] : \mathbb{Q} \in \mathcal{P}\right\}.$$

Since $\mathcal{P} = \{\mathbb{Q}'\}$, we have $\Pi(H) = \left\{\mathbb{E}^{\mathbb{Q}'}\left[\widetilde{H}\right]\right\}$, which is a singleton. By Theorem 2.12, if H were not attainable, then $\Pi(H)$ would be an open interval. Therefore, every contingent claim H must be attainable, and thus the market is complete. □

Theorem 2.15 (Characterization of Complete Markets)**.** For an equivalent martingale measure $\mathbb{Q} \in \mathcal{P}$, the following conditions are equivalent:

i) $\mathcal{P} = \{\mathbb{Q}\}$.

ii) $\mathbb{Q}$ is an extreme point of $\mathcal{P}$, the set of all equivalent martingale measures.

iii) $\mathbb{Q}$ is an extreme point of $\mathcal{Q}$, the set of all martingale measures.

iv) Every $\mathbb{Q}$-martingale $(M_t)_{t=0}^T$ can be represented as a discrete stochastic integral with respect to $\widetilde{S}$, i.e., there exists a predictable process $(\phi_t)_{t=1}^T$ such that

$$M_t = M_0 + \sum_{k=1}^{t} \phi_k^\intercal(\widetilde{S}_k - \widetilde{S}_{k-1}), \quad \text{for } t = 0, 1, \dots, T.$$

Proof. **i) ⇒ iii):** Suppose that $\mathbb{Q}$ is not an extreme point of $\mathcal{Q}$, i.e., there exist $\mathbb{Q}_1, \mathbb{Q}_2 \in \mathcal{Q}$, $\mathbb{Q}_1 \neq \mathbb{Q}_2$, and $\alpha \in (0,1)$ such that

$$\mathbb{Q} = \alpha\mathbb{Q}_1 + (1-\alpha)\mathbb{Q}_2.$$

Since $\mathbb{Q}_1$ and $\mathbb{Q}_2$ are martingale measures, but not necessarily equivalent to $\mathbb{P}$, they may not be in $\mathcal{P}$. However, since $\mathbb{Q}$ is equivalent to $\mathbb{P}$, both $\mathbb{Q}_1$ and $\mathbb{Q}_2$ are absolutely continuous with respect to $\mathbb{Q}$, and hence with respect to $\mathbb{P}$. Therefore, they are also in $\mathcal{P}$, contradicting the assumption that $\mathcal{P} = \{\mathbb{Q}\}$. Thus, $\mathbb{Q}$ must be an extreme point of $\mathcal{Q}$.
iii) ⇒ ii): This follows immediately since $\mathcal{P} \subseteq \mathcal{Q}$, and any extreme point of $\mathcal{Q}$ that belongs to $\mathcal{P}$ is also an extreme point of $\mathcal{P}$.
ii) ⇒ i): Suppose that $\mathbb{Q}$ is an extreme point of $\mathcal{P}$, but that $\mathcal{P}$ contains some $\hat{\mathbb{Q}} \neq \mathbb{Q}$. We will derive a contradiction. The idea is to first show that if a distinct measure $\hat{\mathbb{Q}}$ exists, we can choose it so that the density $d\hat{\mathbb{Q}}/d\mathbb{Q}$ is bounded by some constant c.
Once we have such a measure $\hat{\mathbb{Q}}$, define another measure $\mathbb{Q}'$ by

$$\frac{d\mathbb{Q}'}{d\mathbb{Q}} := 1 + \epsilon\left(\frac{d\hat{\mathbb{Q}}}{d\mathbb{Q}} - 1\right)$$

for some $\epsilon > 0$ small enough (specifically, $\epsilon < 1/c$). Then $\mathbb{Q}'$ is also an EMM and $\mathbb{Q}$ can be expressed as a strict convex combination of $\hat{\mathbb{Q}}$ and $\mathbb{Q}'$, contradicting the extremity of $\mathbb{Q}$. Thus, no other measure $\hat{\mathbb{Q}}$ exists, proving that $\mathcal{P} = \{\mathbb{Q}\}$.
i) ⇒ iv): Let $(M_t)_{t=0}^T$ be any $\mathbb{Q}$-martingale. Then the terminal value M_T is $\mathcal{F}_T$-measurable and integrable under $\mathbb{Q}$. Since the market is complete, any contingent claim is attainable, so there exists a self-financing strategy ϕ such that

$$M_T = M_0 + \sum_{k=1}^{T} \phi_k^\intercal(\widetilde{S}_k - \widetilde{S}_{k-1}), \quad \mathbb{Q}\text{-almost surely}.$$

Therefore, the process

$$M_t = M_0 + \sum_{k=1}^{t} \phi_k^\intercal (\widetilde{S}_k - \widetilde{S}_{k-1}), \quad t = 0, 1, \ldots, T,$$

is a $\mathbb{Q}$-martingale, and since M_T equals both representations, they must coincide. Hence, every $\mathbb{Q}$-martingale can be represented as a stochastic integral with respect to $\widetilde{S}$.
iv) $\Rightarrow$ i): Suppose that every $\mathbb{Q}$-martingale admits a representation as a stochastic integral with respect to $\widetilde{S}$. Let H be any non-negative $\mathcal{F}_T$-measurable random variable with $\mathbb{E}_{\mathbb{Q}}[H] < \infty$. Then the process

$$M_t = \mathbb{E}_{\mathbb{Q}}\left[H \mid \mathcal{F}_t\right], \quad t = 0, 1, \ldots, T,$$

is a $\mathbb{Q}$-martingale, and by assumption, can be represented as

$$M_t = M_0 + \sum_{k=1}^{t} \phi_k^\intercal (\widetilde{S}_k - \widetilde{S}_{k-1}).$$

Therefore, $H = M_T$ can be replicated by the self-financing strategy ϕ, and hence is attainable. Since this holds for any H, the market is complete, and by the Second Fundamental Theorem of Asset Pricing, there is a unique equivalent martingale measure, i.e., $\mathcal{P} = \{\mathbb{Q}\}$. □

2.3.2 Complete Markets in Finite Discrete Time Have Finite Probability Spaces

We call $A \in \mathcal{F}$ an **atom** of $(\Omega, \mathcal{F}, \mathbb{P})$ whenever $\mathbb{P}(A) > 0$ and for each $B \in \mathcal{F}$ with $B \subseteq A$, we have either $\mathbb{P}(B) = 0$ or $\mathbb{P}(B) = \mathbb{P}(A)$. We have the following important result on the structure of complete discrete-time financial markets with finite time horizon:

Proposition 2.16. Let $(\Omega, \mathcal{F}, \mathbb{P}, \mathbb{F}, X)$ be a discrete-time financial market with finite time horizon $T \in \mathbb{N}_0$ and $d \in \mathbb{N}$ assets. If $(\Omega, \mathcal{F}, \mathbb{P}, \mathbb{F}, X)$ is complete, then $(\Omega, \mathcal{F}, \mathbb{P})$ can be decomposed into at most $(d+1)^T$ atoms, and $\dim \mathcal{L}^0(\Omega, \mathcal{F}, \mathbb{P}) \leq (d+1)^T$ holds true. d

Proof. We proceed by induction over $T \in \mathbb{N}_0$.
Base case: For $T = 1$ and assuming the market is complete, then for every $H \in \mathcal{L}^{0,+}(\Omega, \mathcal{F}_1)$, we have

$$H = W_0(\varphi) + \varphi_1^\intercal \Delta S_1, \quad \mathbb{P}\text{-a.s.}$$

As $\mathcal{F}_0$ is trivial and $W_0(\varphi)$, φ_1, and S_0 are $\mathcal{F}_0$-measurable, they are constants, yielding

$$H = \text{constant} + \varphi_1^\intercal \Delta S_1.$$

The random variable ΔS_1 is $\mathcal{F}_1$-measurable, and therefore $\mathcal{L}^{0,+}(\Omega, \mathcal{F}_1)$ is spanned by the $d+1$ random variables $\Delta S_1^{(0)}, \ldots, \Delta S_1^{(d)}$. Consequently, we have $\dim \mathcal{L}^0(\Omega, \mathcal{F}) \leq d+1$, and by by Proposition C.1 the dimension of $\mathcal{L}^0(\Omega, \mathcal{F})$ coincides with the number of atoms of $(\Omega, \mathcal{F}, \mathbb{P})$, the assertion follows for $T = 1$.
Inductive step: Assume the statement holds for $T-1$. We will show it holds for T.

For $H \in \mathcal{L}^0(\Omega, \mathcal{F}_T)$, since the market is complete, H can be expressed as

$$H = W_0(\varphi) + \sum_{k=1}^{T} \varphi_k^\intercal \Delta S_k = W_{T-1}(\varphi) + \varphi_T^\intercal \Delta S_T,$$

where $W_{T-1}(\varphi)$, φ_T, and S_{T-1} are all $\mathcal{F}_{T-1}$-measurable and hence constant on each of the at most $(d+1)^{T-1}$ atoms $(A_i)_{i=1}^{(d+1)^{T-1}}$ of $(\Omega, \mathcal{F}_{T-1}, \mathbb{P})$.
On each atom A_i, the values of $W_{T-1}(\varphi)$ and φ_T are constants. Therefore, on each A_i, H can be expressed as

$$H|_{A_i} = \text{constant} + \varphi_T^\intercal \Delta S_T|_{A_i}.$$

The function $H|_{A_i}$ is again represented by the linear combination of the $d+1$ variables $\Delta S_T^{(0)}|_{A_i}, \ldots, \Delta S_T^{(d)}|_{A_i}$ on A_i. Therefore, by the same argument as before, $\mathcal{F}_T|_{A_i}$ has at most $d+1$ atoms.
Since there are at most $(d+1)^{T-1}$ atoms A_i at time $T-1$, and each can be partitioned into at most $d+1$ atoms at time T, the total number of atoms at time T is at most $(d+1)^{T-1} \times (d+1) = (d+1)^T$. Therefore, $\dim \mathcal{L}^0(\Omega, \mathcal{F}_T, \mathbb{P}) \leq (d+1)^T$. This completes the induction. □

2.4 Superhedging

One might consider that instead of searching for a replicating strategy, the seller of a contingent claim could construct a strategy that ensures the terminal wealth exceeds the claim payoff. This leads to the concept of superhedging.

Definition 2.17. A self-financing strategy $\varphi = (W_0, \phi)$ is called a **superhedge** of a European contingent claim H if $W_T(\varphi) \geq H$ $\mathbb{P}$-almost surely. The initial capital W_0 is called a **superhedge price**.

Although using a superhedge is beneficial for the seller (since it covers the liability of paying H at time T), the superhedge price that the seller charges is generally not an arbitrage-free price, as the following proposition shows.

Proposition 2.18. For any European contingent claim H we define the **cheapest superhedge price** $\pi_{\text{super}}(H)$ as

$$\pi_{\text{super}}(H) := \inf \{W_0 \in \mathbb{R} \mid \exists \text{ self-financing } \varphi = (W_0, \phi) \text{ s.t.} W_T(\varphi) \geq H \ \mathbb{P}\text{-a.s.}\}.$$

Then,

$$\pi_{\text{super}}(H) \geq \sup \Pi(H).$$

If there exists a self-financing strategy $\varphi = (\pi_{\text{super}}(H), \phi)$ such that $W_T(\varphi) \geq H$ $\mathbb{P}$-almost surely, we call φ a **cheapest superhedge** of the claim H.

Proof. Let $\varphi = (V_0, \phi)$ be a self-financing strategy such that $W_T(\varphi) \geq H$ $\mathbb{P}$-almost surely. Let $\mathbb{Q}$ be any equivalent martingale measure (i.e., $\mathbb{Q} \in \mathcal{P}$). Then, since the discounted wealth process $\widetilde{W}_t(\varphi) = W_t(\varphi)/S_t^{(0)}$ is a $\mathbb{Q}$-supermartingale (because it is self-financing and $\widetilde{S}$ is a $\mathbb{Q}$-martingale), we have

$$W_0 = \widetilde{W}_0(\varphi) \geq \mathbb{E}_{\mathbb{Q}}\left[\widetilde{W}_T(\varphi)\right] \geq \mathbb{E}_{\mathbb{Q}}\left[\widetilde{H}\right].$$

Taking the supremum over all $\mathbb{Q} \in \mathcal{P}$, we obtain

$$W_0 \geq \sup_{\mathbb{Q}\in\mathcal{P}} \mathbb{E}_{\mathbb{Q}}\left[\widetilde{H}\right] = \sup \Pi(H).$$

Since $\pi_{\text{super}}(H)$ is the infimum of such W_0, it follows that $\pi_{\text{super}}(H) \geq \sup \Pi(H)$. □

2.5 Pricing and Hedging in the Cox-Ross-Rubinstein Model

We now construct the Cox-Ross-Rubinstein (CRR) model on a so-called canonical probability space. In this model, we have $d = 1$ and set the bank account $S_t^{(0)} = (1+r)^t$ for $t = 0, 1, \dots, T$, for some fixed interest rate $r > -1$. We have only one risky asset with price process $(S_t^{(1)})_{t=0}^T \coloneqq (S_t)_{t=0}^T$. In the CRR model, we assume that the return series $R_t = \frac{\Delta S_t}{S_{t-1}}$ for $t = 1, 2, \dots, T$ can only assume two possible values called U (Up) and D (Down) satisfying $-1 < D < U$. We define this model by

$$\Omega \coloneqq \{-1, +1\}^T = \{\omega = (y_1, y_2, \dots, y_T) \colon y_t \in \{-1, +1\}\},$$

let $Y_t(\omega) = y_t$ for every $\omega = (y_1, y_2, \dots, y_T)$, and set

$$R_t(\omega) = D\frac{1 - Y_t(\omega)}{2} + U\frac{1 + Y_t(\omega)}{2},$$

and for some fixed $S_0 > 0$, set

$$S_t(\omega) = S_0 \prod_{k=1}^{t} (1 + R_k(\omega)), \quad \widetilde{S}_t(\omega) = S_0 \prod_{k=1}^{t} \frac{1 + R_k(\omega)}{1 + r}, \quad t = 0, 1, \dots, T. \tag{29}$$

As a filtration $\mathbb{F}$, we choose $\mathcal{F}_t = \sigma(S_0, S_1, \dots, S_t) = \sigma(Y_1, Y_2, \dots, Y_t)$ for $t = 0, 1, \dots, T$, and note that $\mathcal{F}_0 = \{\emptyset, \Omega\}$ and set $\mathcal{F} \coloneqq \mathcal{F}_T$. Lastly, let $\mathbb{P}$ be a probability measure on $(\Omega, \mathcal{F})$ such that $\mathbb{P}(\{\omega\}) > 0$ for all $\omega \in \Omega$. We call the financial market $(\Omega, \mathcal{F}, \mathbb{P}, \mathbb{F}, (S^{(0)}, S^{(1)}))$ the *canonical CRR model*.

2.5.1 The Arbitrage-Free and Complete CRR Model

Note that according to Proposition 2.16, there is hope that the CRR model is complete since the number of atoms in $(\Omega, \mathcal{F}, \mathbb{P})$ is exactly $(d+1)^T = 2^T$, so right at the boundary. In particular, if there were a time when the risky asset could assume three different values with non-trivial probability, then such a multinomial model could not be complete. The following theorem yields a sufficient and necessary condition for the CRR model being complete.

Theorem 2.19. Let $(\Omega, \mathcal{F}, \mathbb{P}, \mathbb{F}, (S^{(0)}, S^{(1)}))$ denote the Cox-Ross-Rubinstein (CRR) model as above. Then the model is arbitrage-free if and only if $D < r < U$. Under this condition, the model is complete, and its unique equivalent martingale measure $\mathbb{Q}$ is characterized by the fact that the random variables $R_1, R_2, \dots, R_T$ are independent under $\mathbb{Q}$ and satisfy

$$\mathbb{Q}(R_t = U) = \frac{r - D}{U - D}, \quad \forall t = 1, 2, \ldots, T.$$

Proof. Recall that a measure $\mathbb{Q}$ on $(\Omega, \mathcal{F})$ is a martingale measure for the financial market $(\Omega, \mathcal{F}, \mathbb{P}, \mathbb{F}, (S^{(0)}, S^{(1)}))$ if and only if the discounted price process $\widetilde{S} := (\widetilde{S}_t)_{t=0}^T$ is a martingale under $\mathbb{Q}$. In the context of the CRR model, this condition implies that for all $t = 0, 1, \ldots, T-1$, we have

$$\widetilde{S}_t = \mathbb{E}_{\mathbb{Q}}\left[\widetilde{S}_{t+1} \mid \mathcal{F}_t\right] \quad \mathbb{Q}\text{-almost surely.}$$

By inserting (29), this condition reads as

$$\widetilde{S}_t = \widetilde{S}_t \mathbb{E}_{\mathbb{Q}}\left[\frac{1 + R_{t+1}}{1 + r} \mid \mathcal{F}_t\right],$$

which simplifies to

$$\mathbb{E}_{\mathbb{Q}}\left[R_{t+1} \mid \mathcal{F}_t\right] = r \quad \mathbb{Q}\text{-almost surely.}$$

Since R_{t+1} takes only the values U and D, we have

$$\mathbb{E}_{\mathbb{Q}}\left[R_{t+1} \mid \mathcal{F}_t\right] = U\mathbb{Q}(R_{t+1} = U \mid \mathcal{F}_t) + D\left(1 - \mathbb{Q}(R_{t+1} = U \mid \mathcal{F}_t)\right).$$

Solving for $\mathbb{Q}(R_{t+1} = U \mid \mathcal{F}_t)$, we obtain

$$\mathbb{Q}(R_{t+1} = U \mid \mathcal{F}_t) = \frac{r - D}{U - D}.$$

Since the right-hand side does not depend on $\mathcal{F}_t$, it follows that $\mathbb{Q}(R_{t+1} = U \mid \mathcal{F}_t)$ is constant, and thus R_{t+1} is independent of $\mathcal{F}_t$ under $\mathbb{Q}$. Therefore, the R_t are independent under $\mathbb{Q}$, and

$$\mathbb{Q}(R_t = U) = \frac{r - D}{U - D} \quad \text{for all } t = 1, 2, \ldots, T.$$

For $\mathbb{Q}$ to be a probability measure (i.e., with $0 < \mathbb{Q}(R_t = U) < 1$), we require $D < r < U$. Therefore, if the model is arbitrage-free (i.e., there exists an equivalent martingale measure $\mathbb{Q}$), then necessarily $D < r < U$. Conversely, if $D < r < U$, then we can define $\mathbb{Q}$ as above, and since $\mathbb{Q}$ is equivalent to $\mathbb{P}$ (because $\mathbb{Q}(\omega) > 0$ for all ω), and the discounted price process $\widetilde{S}$ is a $\mathbb{Q}$-martingale, the model is arbitrage-free. Moreover, since there is a unique $\mathbb{Q}$ satisfying the martingale condition, the model is complete. □

2.5.2 Pricing in the CRR Model

Let us consider an arbitrage-free CRR model and denote by $\mathbb{Q}$ its unique equivalent martingale measure according to Theorem 2.19. As before, given any European contingent claim $H \in \mathcal{L}^{0,+}(\Omega, \mathcal{F}_T)$, we define its discounted version as $\widetilde{H} = (S_T^{(0)})^{-1} H$. Observe that we can express any contingent claim $\widetilde{H}$ as a function of the risky asset paths since $\mathcal{F}_T = \sigma(S_1, \ldots, S_T)$, i.e., $\widetilde{H} = h(S_0, \ldots, S_T)$ for some function h. The following proposition holds true.

Proposition 2.20. The unique fair price process $(\widetilde{W}_t^H)_{t=0}^T$ of the claim H can be written as

$$\widetilde{W}_t^H(\omega) = w_t(S_0, S_1(\omega), \dots, S_t(\omega)),$$

where

$$w_t(x_0, \dots, x_t) = \mathbb{E}_{\mathbb{Q}}\left[h\left(x_0, \dots, x_t, x_t\frac{S_1}{S_0}, \dots, x_t\frac{S_{T-t}}{S_0}\right)\right]. \tag{30}$$

Proof. Using risk-neutral pricing, we have

$$\begin{aligned}
\widetilde{W}_t^H &= \mathbb{E}_{\mathbb{Q}}\left[\widetilde{H} \mid \mathcal{F}_t\right] = \mathbb{E}_{\mathbb{Q}}\left[h(S_0, \dots, S_T) \mid \mathcal{F}_t\right] \\
&= \mathbb{E}_{\mathbb{Q}}\left[h\left(S_0, \dots, S_t, S_{t+1}, \dots, S_T\right) \mid \mathcal{F}_t\right] \\
&= \mathbb{E}_{\mathbb{Q}}\left[h\left(S_0, \dots, S_t, S_t\frac{S_{t+1}}{S_t}, \dots, S_t\frac{S_T}{S_t}\right) \mid \mathcal{F}_t\right].
\end{aligned}$$

Notice that for every $s \in \mathbb{N}$, the random variable S_{t+s}/S_t is independent of $\mathcal{F}_t$ under $\mathbb{Q}$ and has the same distribution as $S_s/S_0 = \prod_{k=1}^s(1 + R_k)$. Therefore, we conclude, with the help of Fubini's theorem, that:

$$\begin{aligned}
&\mathbb{E}_{\mathbb{Q}}\left[h\left(S_0, \dots, S_t, S_t\frac{S_{t+1}}{S_t}, \dots, S_t\frac{S_T}{S_t}\right) \mid \mathcal{F}_t\right](\omega) \\
&= \mathbb{E}_{\mathbb{Q}}\left[h\left(S_0(\omega), \dots, S_t(\omega), S_t(\omega)\frac{S_1}{S_0}, \dots, S_t(\omega)\frac{S_{T-t}}{S_0}\right)\right],
\end{aligned}$$

which yields the assertion on the form of w_t in (30). □

Following Proposition 2.20, we observe that the wealth process $\widetilde{W}^H$ is characterized by the following recursion:

$$\begin{cases} \widetilde{W}_T^H = \widetilde{H}, \\ \widetilde{W}_t^H = \mathbb{E}_{\mathbb{Q}}\left[\widetilde{W}_{t+1}^H \mid \mathcal{F}_t\right], \quad t = 0, 1, \dots, T-1. \end{cases} \tag{31}$$

Consequently, for $p = \dfrac{r - D}{U - D}$, the functions w_t for $t = 0, 1, \dots, T$ satisfy the recursion:

$$w_T(x_0, x_1, \dots, x_T) = h(x_0, \dots, x_T), \tag{32}$$

$$\begin{aligned}
w_t(x_0, x_1, \dots, x_t) &= p\, w_{t+1}(x_0, \dots, x_t, x_t(1+U)) \\
&\quad + (1-p)\, w_{t+1}(x_0, \dots, x_t, x_t(1+D)), \quad t = T-1, T-2, \dots, 0.
\end{aligned} \tag{33}$$

2.5.3 Hedging in the CRR Model

Definition 2.21. Let H be a European contingent claim and let $(\Omega, \mathcal{F}, \mathbb{P}, \mathbb{F}, (S^{(0)}, S^{(1)}))$ denote a complete CRR model. We define the **delta hedge** φ of H in this model as

$\varphi_t(\omega) = \Delta_t(S_0, S_1(\omega), \ldots, S_{t-1}(\omega))$, where for $t = 1, 2, \ldots, T$ we set

$$\Delta_t(x_0, \ldots, x_{t-1}) := \frac{w_t(x_0, \ldots, x_{t-1}, x_{t-1}(1+U)) - w_t(x_0, \ldots, x_{t-1}, x_{t-1}(1+D))}{x_{t-1}(U-D)}. \tag{34}$$

Next, we show that the delta hedge is a replicating strategy for any claim H:

Proposition 2.22. The delta hedge φ in Definition 2.21 is a replicating strategy for the European contingent claim H.

Proof. For all $\omega = (y_1, y_2, \ldots, y_T)$, the strategy φ must satisfy

$$\varphi_t(\omega)\left(\widetilde{S}_t(\omega) - \widetilde{S}_{t-1}(\omega)\right) = \widetilde{W}_t^H(\omega) - \widetilde{W}_{t-1}^H(\omega), \tag{35}$$

where φ_t, $\widetilde{S}_{t-1}$, and $\widetilde{W}_{t-1}^H$ depend only on the first $t-1$ components of ω. Fix $t \in \{1, 2, \ldots, T\}$ and consider two scenarios:

$$\begin{aligned}
\omega^+ &:= (y_1, y_2, \ldots, y_{t-1}, +1, y_{t+1}, \ldots, y_T), \\
\omega^- &:= (y_1, y_2, \ldots, y_{t-1}, -1, y_{t+1}, \ldots, y_T).
\end{aligned}$$

In these scenarios, the asset price moves up or down at time t. The corresponding changes in wealth are:

$$\begin{aligned}
\widetilde{W}_t^H(\omega^+) - \widetilde{W}_{t-1}^H(\omega) &= \varphi_t(\omega)\left(\widetilde{S}_t(\omega^+) - \widetilde{S}_{t-1}(\omega)\right), \\
\widetilde{W}_t^H(\omega^-) - \widetilde{W}_{t-1}^H(\omega) &= \varphi_t(\omega)\left(\widetilde{S}_t(\omega^-) - \widetilde{S}_{t-1}(\omega)\right).
\end{aligned}$$

Subtracting these two equations, we obtain:

$$\widetilde{W}_t^H(\omega^+) - \widetilde{W}_t^H(\omega^-) = \varphi_t(\omega)\left(\widetilde{S}_t(\omega^+) - \widetilde{S}_t(\omega^-)\right).$$

Solving for $\varphi_t(\omega)$, we get:

$$\varphi_t(\omega) = \frac{\widetilde{W}_t^H(\omega^+) - \widetilde{W}_t^H(\omega^-)}{\widetilde{S}_t(\omega^+) - \widetilde{S}_t(\omega^-)}.$$

Note that:

$$\begin{aligned}
\widetilde{S}_t(\omega^+) &= \widetilde{S}_{t-1}(\omega) \times \frac{1+U}{1+r}, \\
\widetilde{S}_t(\omega^-) &= \widetilde{S}_{t-1}(\omega) \times \frac{1+D}{1+r}.
\end{aligned}$$

Similarly, since $\widetilde{W}_t^H(\omega) = w_t(S_0, S_1(\omega), \ldots, S_t(\omega))$, and $S_t(\omega^\pm) = S_{t-1}(\omega)(1+U \text{ or } 1+D)$, we have:

$$\begin{aligned}
\widetilde{W}_t^H(\omega^+) &= w_t(S_0, \ldots, S_{t-1}(\omega), S_{t-1}(\omega)(1+U)), \\
\widetilde{W}_t^H(\omega^-) &= w_t(S_0, \ldots, S_{t-1}(\omega), S_{t-1}(\omega)(1+D)).
\end{aligned}$$

Therefore,

$$\varphi_t(\omega) = \frac{w_t(x_0, \ldots, x_{t-1}, x_{t-1}(1+U)) - w_t(x_0, \ldots, x_{t-1}, x_{t-1}(1+D))}{\widetilde{S}_{t-1}(\omega)\left(\frac{1+U}{1+r} - \frac{1+D}{1+r}\right)}.$$

Simplifying the denominator:

$$\widetilde{S}_{t-1}(\omega)\left(\frac{1+U-(1+D)}{1+r}\right) = \widetilde{S}_{t-1}(\omega)\frac{U-D}{1+r}.$$

Therefore,

$$\varphi_t(\omega) = \frac{w_t(x_0, \ldots, x_{t-1}, x_{t-1}(1+U)) - w_t(x_0, \ldots, x_{t-1}, x_{t-1}(1+D))}{\widetilde{S}_{t-1}(\omega)\frac{U-D}{1+r}}$$

This confirms the formula in (34), and thus φ replicates H. □

Remark 2.23. Note that for a discounted claim $\widetilde{H} = h(S_T)$ with a payoff function that increases with S_T, the function w_t given as

$$w_t(x) = \mathbb{E}_{\mathbb{Q}}\left[h\left(x\frac{S_{T-t}}{S_0}\right)\right], \quad t = 0, 1, \ldots, T,$$

is also increasing in x. Therefore, the delta hedge (34) is non-negative and does not involve short selling.

Example 2.24 (Hedging and pricing a European call option in the CRR model)**.** Suppose we have a European call option with a payoff $H = h(S_T)$, where S_T represents the stock price at maturity T. The wealth process W^H depends only on S_t for all $t = 0, 1, \ldots, T$, and $W_t^H(\omega) = w_t(S_t(\omega))$ for $(w_t)_{t\in\mathbf{T}}$ as given by (30), i.e.,

$$w_t(x_t) = \sum_{k=0}^{T-t} h\Big(x_t(1+D)^{T-t-k}(1+U)^k\Big)\binom{T-t}{k}\left(\frac{r-D}{U-D}\right)^k\left(1-\frac{r-D}{U-D}\right)^{T-t-k},$$

and in particular, the unique arbitrage-free price π^H is given by:

$$\pi^H = w_0(S_0) = \sum_{k=0}^{T} h\Big(S_0(1+D)^{T-k}(1+U)^k\Big)\binom{T}{k}\left(\frac{r-D}{U-D}\right)^k\left(1-\frac{r-D}{U-D}\right)^{T-k}.$$

In particular, if H is the payoff of a discounted European call option, then

$$h(x) = \frac{(x-K)^+}{(1+r)^T},$$

where K is the strike price.

The unique arbitrage-free price of a European call option with strike K and maturity T in the CRR model is given by:

$$\pi_{\mathrm{EU}}^{H} = \frac{1}{(1+r)^T} \sum_{k=0}^{T} \Big(S_0(1+D)^{T-k}(1+U)^k - K\Big)^+ \binom{T}{k} \Big(\frac{r-D}{U-D}\Big)^k \Big(1 - \frac{r-D}{U-D}\Big)^{T-k}.$$

The hedging strategy, φ, is given by $\varphi_t(\omega) = \Delta_t(S_0, S_1(\omega), \dots, S_{t-1}(\omega))$, where Δ_t is the hedge ratio (or the option's delta) at time t. The hedge ratio for all $t = 1, 2, \dots, T$ is determined as follows:

$$\Delta_t(x_0, \dots, x_{t-1}) = (1+r)^t \frac{\Big(w_t\Big(x_0, \dots, x_{t-1}(1+U)\Big) - w_t\Big(x_0, \dots, x_{t-1}(1+D)\Big)\Big)}{x_{t-1}(1+U) - x_{t-1}(1+D)}. \tag{36}$$

The hedge ratio, Δ_t, represents the number of shares of the underlying stock that need to be held at each time step to replicate the option's payoff.

Incomplete Markets

Preface to Part II: In this second part of the course "Portfolio Theory", we assume an arbitrage-free financial market, denoted as $(\Omega, \mathcal{F}, \mathbb{P}, \mathbb{F}, X)$, encompassing $(d+1) \in \mathbb{N}$ tradable assets with a finite time horizon $T \in \mathbb{N}$. The main assumptions in this part—and indeed, the central challenge that we like to adress—is that the market is supposed to be incomplete. Indeed, we saw in Proposition 2.16 that this is the typical situation in the finite-discrete time setting, at least if one wants to model returns using continuous distributions. But also in the continuous-time setting, in the realm of stochastic volatility models market incompletness poses major challenge, see the example of the Heston stochastic volatility model.

The incompleteness of the market means that there exists a European contingent claim, with pay-off function denoted by $H \in \mathcal{L}^{0,+}(\mathcal{F}_T)$, that is *not attainable*. In fact, if a market is incomplete, then verifying whether an contingent claim is attainable is an infeasible problem as one would need to identify or construct a replicating strategy for H. However, there is in general no algorithm for that. Therefore, in markets characterized by incompleteness, we should anticipate an entire interval of arbitrage-free prices for any claim H, as described in Theorem 2.12. The incompleteness poses a new challenge for agents in a financial market:

i) Agent Bank: Assume that a bank considers the option of selling an contigent claim with pay-off H in an incomplete market. The bank is then confronted with the pricing and hedging problem, i.e., it must decide what price π_0 it charges for the claim H and, secondly, it must trade according to a reasonable hedging strategy. In the incomplete market setting these problems do not have a unique solution and in many situations, mere arbitrage-free pricing does not yield a satisfactory answer to the two problems, as, e.g., in the standard model the interval of arbitrage-free prices is very broad. The bank therefore needs an additional criteria to decide on the price of the claim at time $t = 0$ charge, as well as what hedging strategy to follow.

ii) Agent Investor: Related to the problem that a bank faces when selling a contingent claim, a investor who invests an initial capital W_0 at time $t = 0$ into the d-risky assets wants to maximize her terminal wealth $W_T(\varphi) = W_0 + G_T(\phi)$, or some function of that, following some optimal strategy ϕ. However, the terminal wealth $W_T(\varphi)$ is uncertain and it is not clear what strategy φ is *optimal*, and in what sense the investor measures optimality. Therefore, also the investor, without selling any contingent claim and knowing what the inital investment is, is looking for a criteria to decide what an opitmal strategy is.

Utility and Risk: To address the pricing, hedging and optimal investment problem in the incomplete market setting, we introduce three criteria to find appropriate prices, hedging and investment strategies. The first criteria that we discuss in this course is the minimization of the L^2-error or variance of hedging error; next we consider *Risk Measures* (see Section 4) and *Expected Utility* (see Section 5).

3 Variance-optimal Hedging

In an incomplete market setting, it is generally not possible to perfectly hedge every European contingent claim. As a result, the fair price of an option is not unique but spans an interval. One approach to determine a pricing and hedging criterion in this context is to find an initial price W_0 and a trading strategy ϕ that minimize the hedging error $W_0 + G_T(\phi) - H$ at the terminal time T. A common way to quantify this error is through the L^2-norm. This method is known as *quadratic* or *variance-optimal hedging*.
Throughout this section, we assume that H is a non-negative, square-integrable European contingent claim, i.e.,

$$H \in \mathcal{L}^2(\Omega, \mathcal{F}_T, \mathbb{P}) \cap \mathcal{L}^{0,+}(\Omega, \mathcal{F}_T).$$

Moreover, let $X_t^{(i)} \in \mathcal{L}^2(\mathbb{P})$ for all $t = 0, 1, \dots, T$ and $i = 1, 2, \dots, d$ denote the discounted asset price processes. We define the set of square-integrable strategies $\mathcal{S}^2$ by

$$\mathcal{S}^2 := \left\{\phi\colon \phi \text{ is } \mathbb{R}^d\text{-valued predictable process s.t. } G_t(\phi) \in \mathcal{L}^2(\Omega, \mathcal{F}_t, \mathbb{P})\ \forall t = 0, 1, \dots, T\right\}. \tag{37}$$

Here, $G_T(\phi)$ is the cumulative gain process up to time T, given by

$$G_T(\phi) = \sum_{t=1}^{T} \phi_t \Delta X_t,$$

where $\Delta X_t = X_t - X_{t-1}$.

Definition 3.1 (Variance-Optimal Strategy)**.** A self-financing strategy $\varphi^* = (W_0^*, \phi^*)$ with $W_0^* \in \mathbb{R}$ and $\phi^* \in \mathcal{S}^2$ is called a **variance-optimal strategy** for H if

$$\mathbb{E}_{\mathbb{P}}\left[(W_0^* + G_T(\phi^*) - H)^2\right] \leq \mathbb{E}_{\mathbb{P}}\left[(W_0 + G_T(\phi) - H)^2\right], \tag{38}$$

for all $W_0 \in \mathbb{R}$ and $\phi \in \mathcal{S}^2$. That is, it minimizes the expected quadratic hedging error among all self-financing strategies.

Recall that according to the self-financing condition, any pair $(W_0, \phi) \in \mathbb{R} \times \mathcal{S}^2$ defines a self-financing strategy φ with wealth process

$$W_t(\varphi) = W_0 + G_t(\phi), \quad t = 0, 1, \dots, T.$$

Consequently, $W_0 + G_T(\phi)$ in equation (38) equals $W_T(\varphi)$.
A variance-optimal strategy is not necessarily unique in terms of the trading strategy ϕ. However, the following lemma shows that the (discounted) wealth processes of any two variance-optimal strategies coincide.

Lemma 3.2. Any two variance-optimal hedging strategies φ and φ' will have the same (discounted) terminal wealth, i.e., $W_T(\varphi) = W_T(\varphi')$ $\mathbb{P}$-a.s.

Proof. Assume, for contradiction, that there exist two variance-optimal strategies φ and φ' such that $W_T(\varphi) \neq W_T(\varphi')$ on a set of positive probability. Consider the convex combination of these strategies:

$$\psi = \frac{1}{2}\varphi + \frac{1}{2}\varphi'.$$

Then ψ is also a self-financing strategy in $\mathcal{S}^2$ with initial wealth W_0^*. The terminal wealth of ψ is

$$W_T(\psi) = \frac{1}{2}W_T(\varphi) + \frac{1}{2}W_T(\varphi').$$

Since the function $x \mapsto (x - H)^2$ is strictly convex, we have

$$\begin{aligned}\mathbb{E}_{\mathbb{P}}\left[\Big(W_T(\psi) - H\Big)^2\right] &< \frac{1}{2}\mathbb{E}_{\mathbb{P}}\left[\Big(W_T(\varphi) - H\Big)^2\right] + \frac{1}{2}\mathbb{E}_{\mathbb{P}}\left[\Big(W_T(\varphi') - H\Big)^2\right] \\ &= \mathbb{E}_{\mathbb{P}}\left[\Big(W_T(\varphi) - H\Big)^2\right],\end{aligned}$$

since φ and φ' are both variance-optimal strategies with the same minimal expected squared error. This inequality contradicts the assumption that φ is variance-optimal. Therefore, it must be that $W_T(\varphi) = W_T(\varphi')$ $\mathbb{P}$-a.s. □

Remark 3.3. While the terminal wealth $W_T(\varphi)$ is uniquely determined for variance-optimal strategies, the trading strategies ϕ themselves may not be unique.

To simplify the discussion, we will henceforth assume that $d = 1$. In other words, we are considering a market consisting of only one risky asset (alongside the bank account). A variance-optimal hedge is any solution to the minimization problem:

$$\min\left\{\mathbb{E}_{\mathbb{P}}\left[\Big(W_0 + G_T(\phi) - H\Big)^2\right] : W_0 \in \mathbb{R},\, \phi \in \mathcal{S}^2\right\}.$$

Our strategy is to first optimize over all $\phi \in \mathcal{S}^2$ for a fixed $W_0 \in \mathbb{R}$ and then subsequently optimize over W_0, hoping that this yields a variance-optimal hedge. Define the subspace:

$$\mathcal{G}_T := \left\{G_T(\phi) \colon \phi \in \mathcal{S}^2\right\} \subseteq \mathcal{L}^2(\Omega, \mathcal{F}_T, \mathbb{P}). \tag{39}$$

If $\mathcal{G}_T$ is a closed subset of $\mathcal{L}^2(\Omega, \mathcal{F}_T, \mathbb{P})$, then a variance-optimal strategy in the risky asset would be any ϕ^* such that $G_T(\phi^*) = \mathsf{P}(H)$ $\mathbb{P}$-a.s., where P denotes the orthogonal projection onto $\mathcal{G}_T$. However, it is not obvious whether $\mathcal{G}_T$ is closed or not. It turns out that a simple relation between the conditional variance and the conditional expectation of the increments of X yields a sufficient condition for the closedness of $\mathcal{G}_T$. For this, let us denote the conditional variance of the increments of the asset price X_t as

$$\sigma_t^2 := \mathrm{Var}(\Delta X_t | \mathcal{F}_{t-1}), \quad t = 1, \ldots, T. \tag{40}$$

Similarly, for the conditional expectation of the increments, we write

$$\mu_t := \mathbb{E}_{\mathbb{P}}\left[\Delta X_t | \mathcal{F}_{t-1}\right], \quad t = 1, \ldots, T. \tag{41}$$

Proposition 3.4. If there exists a constant $K > 0$ such that

$$\mu_t^2 \le K\sigma_t^2 \quad \mathbb{P}\text{-a.s. for all } t = 1, \ldots, T, \tag{42}$$

then $\mathcal{G}_T$ is a closed subspace of $\mathcal{L}^2(\Omega, \mathcal{F}_T, \mathbb{P})$. The condition (42) is sometimes referred to as the **bounded mean-variance trade-off**.

Proof. Since X is square-integrable and adapted, there exists a martingale $(M_t)_{t\in\mathbf{T}}$ and a predictable process $(A_t)_{t\in\mathbf{T}}$ as per the Doob decomposition theorem such that $A_0 = 0$ and

$$X_t = A_t + M_t, \quad \text{for all } t = 0, 1, \dots, T.$$

Note that

$$\sigma_t^2 = \mathrm{Var}(\Delta X_t|\mathcal{F}_{t-1}) = \mathrm{Var}(\Delta M_t|\mathcal{F}_{t-1}) = \mathbb{E}_{\mathbb{P}}\left[\left(\Delta M_t\right)^2|\mathcal{F}_{t-1}\right],$$

since $\Delta A_t = A_t - A_{t-1}$ is $\mathcal{F}_{t-1}$-measurable.
For any $\phi \in \mathcal{S}^2$, the gain process satisfies

$$G_T(\phi) = G_{T-1}(\phi) + \phi_T \Delta X_T.$$

Thus,

$$\begin{aligned}\mathbb{E}_{\mathbb{P}}\left[G_T(\phi)^2\right] &= \mathbb{E}_{\mathbb{P}}\left[(G_{T-1}(\phi) + \phi_T\Delta X_T)^2\right]\\ &= \mathbb{E}_{\mathbb{P}}\left[G_{T-1}(\phi)^2\right] + 2\mathbb{E}_{\mathbb{P}}\left[G_{T-1}(\phi)\phi_T\Delta X_T\right] + \mathbb{E}_{\mathbb{P}}\left[\phi_T^2(\Delta X_T)^2\right]. \qquad (43)\end{aligned}$$

The cross term $\mathbb{E}_{\mathbb{P}}\left[G_{T-1}(\phi)\phi_T\Delta X_T\right]$ can be simplified using the tower property and the fact that $G_{T-1}(\phi)$ and ϕ_T are $\mathcal{F}_{T-1}$-measurable:

$$\mathbb{E}_{\mathbb{P}}\left[G_{T-1}(\phi)\phi_T\Delta X_T\right] = \mathbb{E}_{\mathbb{P}}\left[G_{T-1}(\phi)\phi_T\mu_T\right].$$

Similarly, the last term in (43) can be written as

$$\mathbb{E}_{\mathbb{P}}\left[\phi_T^2(\Delta X_T)^2\right] = \mathbb{E}_{\mathbb{P}}\left[\phi_T^2\mathbb{E}_{\mathbb{P}}\left[(\Delta X_T)^2|\mathcal{F}_{T-1}\right]\right] = \mathbb{E}_{\mathbb{P}}\left[\phi_T^2(\sigma_T^2 + \mu_T^2)\right].$$

Combining these, we get

$$\mathbb{E}_{\mathbb{P}}\left[G_T(\phi)^2\right] = \mathbb{E}_{\mathbb{P}}\left[G_{T-1}(\phi)^2\right] + 2\mathbb{E}_{\mathbb{P}}\left[G_{T-1}(\phi)\phi_T\mu_T\right] + \mathbb{E}_{\mathbb{P}}\left[\phi_T^2(\sigma_T^2 + \mu_T^2)\right]. \qquad (44)$$

Now, consider a sequence $(\phi^n)_{n\in\mathbb{N}} \subseteq \mathcal{S}^2$ such that $G_T(\phi^n)$ converges in $\mathcal{L}^2(\Omega, \mathcal{F}_T, \mathbb{P})$ to some $Z \in \mathcal{L}^2(\Omega, \mathcal{F}_T, \mathbb{P})$. Our goal is to show that $Z \in \mathcal{G}_T$.
First, observe that the sequence $(G_T(\phi^n))$ is Cauchy in $\mathcal{L}^2$. Using (44), we can write the difference between $G_T(\phi^n)$ and $G_T(\phi^m)$ as

$$G_T(\phi^n) - G_T(\phi^m) = \sum_{t=1}^{T}(\phi_t^n - \phi_t^m)\Delta X_t.$$

Then,

$$\begin{aligned}\mathbb{E}_{\mathbb{P}}\left[(G_T(\phi^n) - G_T(\phi^m))^2\right] &= \sum_{t=1}^{T}\mathbb{E}_{\mathbb{P}}\left[(\phi_t^n - \phi_t^m)^2\left(\sigma_t^2 + \mu_t^2\right)\right]\\ &\geq \sum_{t=1}^{T}\mathbb{E}_{\mathbb{P}}\left[(\phi_t^n - \phi_t^m)^2\,\sigma_t^2\right].\end{aligned}$$

This inequality implies that $(\phi_t^n\sigma_t)_{n\in\mathbb{N}}$ is a Cauchy sequence in $\mathcal{L}^2(\Omega, \mathcal{F}_T, \mathbb{P})$ for each t.

Define $\phi_t := \lim_{n\to\infty} \phi_t^n$ in $\mathcal{L}^2$ (up to a subsequence if necessary). Then, define

$$\Psi_t := \begin{cases} \dfrac{\phi_t \sigma_t}{\sigma_t^2}, & \text{on } \{\sigma_t > 0\}, \\ 0, & \text{on } \{\sigma_t = 0\}. \end{cases}$$

Note that Ψ_t is well-defined and $\mathcal{F}_{t-1}$-measurable.
Using the bounded mean-variance trade-off condition (42), we have

$$\sigma_t^2 + \mu_t^2 \leq (1+K)\sigma_t^2.$$

Therefore,

$$\begin{aligned} \mathbb{E}_{\mathbb{P}}\left[(\phi_t^n - \phi_t)^2 \left(\sigma_t^2 + \mu_t^2\right)\right] &\leq (1+K)\mathbb{E}_{\mathbb{P}}\left[(\phi_t^n - \phi_t)^2 \sigma_t^2\right] \\ &= (1+K)\mathbb{E}_{\mathbb{P}}\left[(\phi_t^n \sigma_t - \phi_t \sigma_t)^2\right] \\ &\to 0 \quad \text{as } n \to \infty. \end{aligned}$$

This implies that $G_T(\phi^n)$ converges to $G_T(\phi)$ in $\mathcal{L}^2$, so $Z = G_T(\phi) \in \mathcal{G}_T$. Therefore, $\mathcal{G}_T$ is closed in $\mathcal{L}^2(\Omega, \mathcal{F}_T, \mathbb{P})$. □

Remark 3.5. The bounded mean-variance trade-off condition (42) ensures that the variance of the asset's returns dominates its mean squared returns uniformly over time. This condition is crucial for the closedness of $\mathcal{G}_T$, as it prevents the trading strategies from becoming too "explosive" in the presence of high expected returns relative to volatility.

Proposition 3.6. Under the bounded mean-variance trade-off condition (42), there exists a variance-optimal strategy $\varphi^* = (W_0^*, \phi^*)$. Moreover, such a variance-optimal strategy is $\mathbb{P}$-almost surely unique up to modifications of ϕ_t^* on the set $\{\sigma_t = 0\}$ for all $t = 1, 2, \ldots, T$.

Proof. Since the bounded mean-variance trade-off condition (42) holds, Proposition 3.4 ensures that $\mathcal{G}_T$ is a closed subspace of $L^2(\Omega, \mathcal{F}_T, \mathbb{P})$.
Let $\mathsf{P} : L^2(\mathbb{P}) \to \mathcal{G}_T$ denote the orthogonal projection onto $\mathcal{G}_T$. For any $Z \in L^2(\mathbb{P})$, $\mathsf{P}(Z)$ is the unique element in $\mathcal{G}_T$ satisfying

$$\mathbb{E}_{\mathbb{P}}\left[(Z - \mathsf{P}(Z))^2\right] = \min_{Y \in \mathcal{G}_T} \mathbb{E}_{\mathbb{P}}\left[(Z - Y)^2\right].$$

For any $W_0 \in \mathbb{R}$, there exists $\phi^{W_0} \in \mathcal{S}^d$ such that

$$G_T(\phi^{W_0}) = \mathsf{P}(H - W_0).$$

This is because $\mathsf{P}(H - W_0) \in \mathcal{G}_T$ by definition, and there exists a strategy ϕ^{W_0} such that $G_T(\phi^{W_0}) = \mathsf{P}(H - W_0)$.
The expected squared hedging error is then

$$\mathbb{E}_{\mathbb{P}}\left[\left(W_0 + G_T(\phi^{W_0}) - H\right)^2\right] = \mathbb{E}_{\mathbb{P}}\left[(W_0 + \mathsf{P}(H - W_0) - H)^2\right].$$

Using the linearity of the projection operator, we have

$$\mathsf{P}(H - W_0) = \mathsf{P}(H) - W_0\mathsf{P}(1),$$

where 1 denotes the constant function equal to 1.
Therefore,

$$\begin{aligned} W_0 + \mathsf{P}(H - W_0) - H &= W_0 + (\mathsf{P}(H) - W_0\mathsf{P}(1)) - H \\ &= (\mathsf{P}(H) - H) + W_0\,(1 - \mathsf{P}(1))\,. \end{aligned}$$

Thus, the expected squared hedging error becomes

$$J(W_0) = \mathbb{E}_{\mathbb{P}}\left[((\mathsf{P}(H) - H) + W_0\,(1 - \mathsf{P}(1)))^2\right].$$

This is a quadratic function in W_0, since $\mathsf{P}(H) - H$ is independent of W_0.
Since $J(W_0)$ is quadratic with respect to W_0 and the coefficient of W_0^2 is non-negative (and positive unless $1 - \mathsf{P}(1) = 0$ almost surely), it attains its minimum at a unique value $W_0^* \in \mathbb{R}$. Therefore, the variance-optimal initial investment is W_0^*, and the corresponding trading strategy is $\phi^* = \phi^{W_0^*}$.
Suppose there are two variance-optimal strategies $\varphi^* = (W_0^*, \phi^*)$ and $\tilde{\varphi}^* = (W_0^*, \tilde{\phi}^*)$ that achieve the same minimal variance, but differ on a set of positive probability.
Since both are variance-optimal:

$$G_T(\phi^*) = \mathcal{P}(H - W_0^*) = G_T(\tilde{\phi}^*).$$

Thus,

$$G_T(\phi^* - \tilde{\phi}^*) = 0.$$

Recall that $G_T(\phi) = \sum_{t=1}^{T} \phi_t \Delta X_t$. So,

$$\sum_{t=1}^{T} (\phi_t^* - \tilde{\phi}_t^*)\Delta X_t = 0 \quad \mathbb{P}\text{-a.s.}$$

Conditioning on $\mathcal{F}_{t-1}$, the increment ΔX_t has variance σ_t^2. If $\sigma_t > 0$ almost surely, the only way for the above equality to hold is if $\phi_t^* - \tilde{\phi}_t^* = 0$ almost surely on that event. Indeed, if $\sigma_t > 0$, ΔX_t can vary, forcing $\phi_t^* = \tilde{\phi}_t^*$ almost surely.
If $\sigma_t = 0$ on some set, ϕ_t cannot be identified uniquely from $G_T(\phi)$, and hence ϕ_t^* may differ from $\tilde{\phi}_t^*$ on $\{\sigma_t = 0\}$ without affecting the variance.
Therefore, the variance-optimal strategy is unique up to modifications on the sets where $\sigma_t = 0$. □

Under additional assumptions, we can construct variance-optimal strategies explicitly. In particular, when the discounted asset price process X is a martingale under the probability measure $\mathbb{P}$, the variance-optimal strategy can be derived using martingale representation theorems.
Recall that in the variance-optimal hedging problem, the goal is to find a self-financing trading strategy $\varphi = (W_0, \phi)$ that minimizes the expected squared hedging error given by $\mathbb{E}_{\mathbb{P}}\left[(W_T(\varphi) - H)^2\right]$. Note that the wealth process of a variance-optimal strategy corresponds to the orthogonal projection of the discounted payoff H onto the space of attainable terminal wealths generated by trading in the asset X. Specifically, we can decompose the martingale $\widehat{W}_t = \mathbb{E}_{\mathbb{P}}[H \mid \mathcal{F}_t]$ into components driven by X and orthogonal to X.
We have the following important result:

Lemma 3.7 (Kunita-Watanabe Decomposition)**.** Let $(Y_t)_{t=0}^T$ be a square-integrable $\mathbb{P}$-martingale, and let $(X_t)_{t=0}^T$ be a square-integrable $\mathbb{P}$-martingale. Then there exists a predictable process ϕ and a square-integrable martingale M orthogonal to X (i.e., $\mathbb{E}_{\mathbb{P}}\left[\Delta M_t \Delta X_t \mid \mathcal{F}_{t-1}\right] = 0$ for all t), such that

$$Y_t = Y_0 + \sum_{s=1}^{t} \phi_s \Delta X_s + M_t, \quad t = 0, 1, \dots, T. \tag{45}$$

Moreover, the decomposition is unique.

Proof. Since both Y and X are square-integrable martingales, we can consider their predictable covariation processes. Define the predictable process ϕ by

$$\phi_t = \frac{\mathbb{E}_{\mathbb{P}}\left[\Delta Y_t \Delta X_t \mid \mathcal{F}_{t-1}\right]}{\mathbb{E}_{\mathbb{P}}\left[(\Delta X_t)^2 \mid \mathcal{F}_{t-1}\right]} \mathbb{1}_{\{\mathbb{E}_{\mathbb{P}}[(\Delta X_t)^2 \mid \mathcal{F}_{t-1}] > 0\}}, \quad t = 1, \dots, T.$$

Define the process M by

$$\Delta M_t = \Delta Y_t - \phi_t \Delta X_t, \quad t = 1, \dots, T.$$

Note that

$$\begin{aligned} \mathbb{E}_{\mathbb{P}}\left[\Delta M_t \Delta X_t \mid \mathcal{F}_{t-1}\right] &= \mathbb{E}_{\mathbb{P}}\left[\left(\Delta Y_t - \phi_t \Delta X_t\right) \Delta X_t \mid \mathcal{F}_{t-1}\right] \\ &= \mathbb{E}_{\mathbb{P}}\left[\Delta Y_t \Delta X_t \mid \mathcal{F}_{t-1}\right] - \phi_t \mathbb{E}_{\mathbb{P}}\left[(\Delta X_t)^2 \mid \mathcal{F}_{t-1}\right] = 0, \end{aligned}$$

by the definition of ϕ_t. Therefore, M is orthogonal to X. The Remainder is left as Exercise 3.4. □

Using this decomposition, we can construct the variance-optimal strategy explicitly when X is a martingale.

Theorem 3.8. Assume that under $\mathbb{P}$, the discounted asset price process X is a square-integrable martingale, i.e., $\mathbb{E}_{\mathbb{P}}\left[\Delta X_t \mid \mathcal{F}_{t-1}\right] = 0$ for all t. Let H be a square-integrable contingent claim, and let $\widehat{W}_t = \mathbb{E}_{\mathbb{P}}[H \mid \mathcal{F}_t]$ be its $\mathbb{P}$-martingale representation. Define the trading strategy $\varphi^* = (W_0^*, \phi^*)$ as

$$\phi_t^* := \frac{\mathbb{E}_{\mathbb{P}}\left[\Delta \widehat{W}_t \Delta X_t \mid \mathcal{F}_{t-1}\right]}{\mathbb{E}_{\mathbb{P}}\left[(\Delta X_t)^2 \mid \mathcal{F}_{t-1}\right]} \mathbb{1}_{\{\mathbb{E}_{\mathbb{P}}[(\Delta X_t)^2 \mid \mathcal{F}_{t-1}] > 0\}}, \quad t = 1, \dots, T,$$

and

$$W_0^* := \widehat{W}_0 = \mathbb{E}_{\mathbb{P}}[H].$$

Then, the trading strategy φ^* is variance-optimal for hedging H, and the mean squared hedging error is given by

$$\mathbb{E}_{\mathbb{P}}\left[\left(W_0^* + \sum_{t=1}^{T} \phi_t^* \Delta X_t - H\right)^2\right] = \sum_{t=1}^{T} \mathbb{E}_{\mathbb{P}}\left[\operatorname{Var}\left(\Delta \widehat{W}_t \mid \mathcal{F}_{t-1}\right) - (\phi_t^*)^2 \operatorname{Var}\left(\Delta X_t \mid \mathcal{F}_{t-1}\right)\right].$$

Proof. Left as Exercise 3.4

□

4 Risk Measures

A drawback of using the variance to measure the hedging error is that the variance is symmetric; that is, positive and negative deviations from being a replicating strategy are treated the same. However, a positive deviation, i.e., when $W_T(\varphi)(\omega) - H(\omega) > 0$, is actually favorable for the bank. Therefore, it would be better to use a measure that focuses on the downside risk. This leads us to the concept of *risk measures*, which we introduce in this section.

4.1 Monotonicity, Cash Invariance, Convexity, and Coherence

Definition 4.1. Let $\mathcal{X}$ denote a class of real-valued random variables on $(\Omega, \mathcal{F})$ representing *financial positions*. We call a map $\rho\colon \mathcal{X} \to \mathbb{R}$ a **risk measure** if it satisfies the following conditions for all $X, Y \in \mathcal{X}$:

i) If $X \leq Y$ almost surely, then $\rho(X) \geq \rho(Y)$ (**monotonicity**).

ii) For all $m \in \mathbb{R}$, $\rho(X + m) = \rho(X) - m$ (**cash invariance**).

If a risk measure ρ satisfies $\rho(0) = 0$, then we call it **normalized**. Additionally, if ρ is **convex**, i.e., if

$$\rho(\lambda X + (1-\lambda)Y) \leq \lambda\rho(X) + (1-\lambda)\rho(Y), \quad \forall\, 0 \leq \lambda \leq 1, \tag{46}$$

then we call ρ a **convex risk measure**.

Remark 4.2.

i) A risk measure is monotone decreasing because the downside risk of a financial position with a larger payoff profile is less than that of a position with a smaller payoff.

ii) Cash invariance means that adding a riskless amount of cash m to a financial position reduces the risk measure by m. This reflects the idea that holding cash reduces risk, while owing cash (debt) increases risk.

iii) Convexity of a risk measure implies that diversification reduces risk. Specifically, combining two financial positions should not increase the risk measure beyond the weighted average of their individual risks.

Definition 4.3. A convex risk measure ρ is called a **coherent risk measure** if it is also positive homogeneous and subadditive. Specifically, ρ satisfies:

i) $\rho(\lambda X) = \lambda\rho(X)$ for all $\lambda \geq 0$ (**positive homogeneity**).

ii) $\rho(X + Y) \leq \rho(X) + \rho(Y)$ (**subadditivity**).

Remark 4.4. Coherence allows financial institutions to split up their financial position into smaller parts, assess the risk of each part, and then bound the total risk by the sum of the risks of the parts. However, positive homogeneity means that the financial risk grows linearly with its size, which may not always hold in real markets due to economies or diseconomies of scale. Therefore, it is sometimes more appropriate to measure risk with non-coherent (convex but not coherent) risk measures.

4.2 Acceptance Sets of Risk Measures

Let $\mathcal{X}$ denote a set of *discounted financial positions* in a financial market $(\Omega, \mathcal{F}, \mathbb{P}, \mathbb{F}, X)$; that is, $\mathcal{X}$ is a set of real-valued measurable functions defined on some set of outcomes Ω. For instance, we could think of $\mathcal{X} = L^p(\Omega, \mathcal{F}, \mathbb{P})$ for $1 \leq p \leq \infty$.
Given a risk measure $\rho\colon \mathcal{X} \to \mathbb{R}$, we define the **acceptance set** of ρ as

$$\mathcal{A}_\rho \coloneqq \{X \in \mathcal{X}\colon \rho(X) \leq 0\}. \tag{47}$$

We call a financial position $X \in \mathcal{A}_\rho$ **acceptable** with respect to the risk measure ρ.
From the perspective of a supervisory agency, the risk $\rho(X)$ of a financial position X can be viewed as the minimal capital requirement that, when added to the position and invested in a risk-free asset, makes the financial position acceptable. Indeed, $\rho(X)$ is the exact amount of capital that one has to add to the financial position X such that, by cash invariance,

$$\rho(X + \rho(X)) = \rho(X) - \rho(X) = 0,$$

i.e., the adjusted position $X + \rho(X)$ is acceptable (since it belongs to $\mathcal{A}_\rho$).
In the following proposition, we show some properties and connections between the acceptance set and the risk measure.

Proposition 4.5. Let ρ be a risk measure on the space of bounded measurable functions $\mathcal{X} = L^\infty(\Omega, \mathcal{F}, \mathbb{P})$. For the acceptance set $\mathcal{A}_\rho$, the following holds true:

i) $\mathcal{A}_\rho$ is non-empty and closed in $\mathcal{X}$ equipped with the supremum norm $\|\cdot\|_\infty$.

ii) $\inf\{m \in \mathbb{R}\colon m\mathbf{1} \in \mathcal{A}_\rho\} > -\infty$.

iii) If $X \in \mathcal{A}_\rho$ and $Y \in \mathcal{X}$ with $Y \geq X$, then $Y \in \mathcal{A}_\rho$.

iv) For all $X \in \mathcal{X}$,

$$\rho(X) = \inf\{m \in \mathbb{R}\colon m + X \in \mathcal{A}_\rho\}. \tag{48}$$

v) ρ is convex if and only if $\mathcal{A}_\rho$ is convex.

vi) ρ is positive homogeneous, i.e., $\rho(\lambda X) = \lambda\rho(X)$ for all $\lambda \geq 0$, if and only if $\mathcal{A}_\rho$ is a cone; that is, $\lambda\mathcal{A}_\rho \subseteq \mathcal{A}_\rho$ for all $\lambda \geq 0$.

Proof. Let $\rho : \mathcal{X} \to \mathbb{R}$ be a monetary risk measure, where $\mathcal{X} = L^\infty(\Omega, \mathcal{F}, \mathbb{P})$. Recall that ρ being monetary means it is monotone and cash-invariant. The acceptance set associated with ρ is defined as

$$\mathcal{A}_\rho := \{X \in \mathcal{X} : \rho(X) \leq 0\}.$$

i) Non-emptiness and Closedness
Non-empty: Consider $X = 0$. If $\rho(0) \le 0$, then $0 \in \mathcal{A}_\rho$. If $\rho(0) > 0$, define $m := \rho(0)$. By cash-invariance,

$$\rho(m\mathbf{1}) = \rho(0) - m = 0,$$

so $m\mathbf{1} \in \mathcal{A}_\rho$. Hence, $\mathcal{A}_\rho$ is non-empty.
Closedness: Let $(X_n)_{n\in\mathbb{N}} \subseteq \mathcal{A}_\rho$ and suppose $X_n \to X$ in $\|\cdot\|_\infty$. Since $\rho(X_n) \le 0$ for all n and ρ is Lipschitz continuous with respect to $\|\cdot\|_\infty$ (shown below), we have $\rho(X_n) \to \rho(X)$. Being a limit of non-positive numbers, $\rho(X) \le 0$, so $X \in \mathcal{A}_\rho$. Thus, $\mathcal{A}_\rho$ is closed.
Lipschitz continuity of ρ: For $X, Y \in \mathcal{X}$, since $X \le Y + \|X - Y\|_\infty \mathbf{1}$, by monotonicity and cash-invariance,

$$\rho(X) \ge \rho(Y + \|X - Y\|_\infty \mathbf{1}) = \rho(Y) - \|X - Y\|_\infty.$$

Reversing the roles of X and Y, we also get

$$\rho(Y) \ge \rho(X) - \|X - Y\|_\infty.$$

Combining these inequalities,

$$|\rho(X) - \rho(Y)| \le \|X - Y\|_\infty.$$

So ρ is 1-Lipschitz continuous.
ii) Boundedness of the Infimum
By cash-invariance,

$$\rho(m\mathbf{1}) = \rho(0) - m.$$

Then

$$m\mathbf{1} \in \mathcal{A}_\rho \iff \rho(m\mathbf{1}) \le 0 \iff \rho(0) - m \le 0 \iff m \ge \rho(0).$$

Hence,

$$\inf\{m \in \mathbb{R} : m\mathbf{1} \in \mathcal{A}_\rho\} = \rho(0) > -\infty.$$

iii) Monotonicity of $\mathcal{A}_\rho$
If $X \in \mathcal{A}_\rho$ and $Y \in \mathcal{X}$ with $Y \ge X$, then by monotonicity,

$$\rho(Y) \le \rho(X) \le 0,$$

so $Y \in \mathcal{A}_\rho$. Thus, $\mathcal{A}_\rho$ is monotone.
iv) Representation of ρ via $\mathcal{A}_\rho$
For any $X \in \mathcal{X}$,

$$\rho(X) = \inf\{m : \rho(X + m\mathbf{1}) \le 0\} = \inf\{m : X + m\mathbf{1} \in \mathcal{A}_\rho\}.$$

Rewriting $X + m\mathbf{1} = m + X$ gives $\rho(X) = \inf\{m \in \mathbb{R} : m + X \in \mathcal{A}_\rho\}$, proving (48).
v) Convexity Equivalence
If ρ is convex, let $X, Y \in \mathcal{A}_\rho$ and $0 \le \lambda \le 1$. Then

$$\rho(\lambda X + (1-\lambda)Y) \le \lambda\rho(X) + (1-\lambda)\rho(Y) \le 0,$$

so $\lambda X + (1-\lambda)Y \in \mathcal{A}_\rho$. Thus $\mathcal{A}_\rho$ is convex. Conversely, assume $\mathcal{A}_\rho$ is convex. For $X, Y \in \mathcal{X}$, set $m_X = \rho(X), m_Y = \rho(Y)$. Then by cash-invariance, $X' := X + m_X\mathbf{1}$ and

$Y' := Y + m_Y\mathbf{1}$ satisfy $\rho(X') = 0$ and $\rho(Y') = 0$, so $X', Y' \in \mathcal{A}_\rho$. By convexity of $\mathcal{A}_\rho$, $Z := \lambda X' + (1-\lambda)Y' \in \mathcal{A}_\rho \implies \rho(Z) \le 0$. But

$$Z - (\lambda m_X + (1-\lambda)m_Y)\mathbf{1} = \lambda X + (1-\lambda)Y.$$

Using cash-invariance,

$$\begin{aligned}\rho(\lambda X + (1-\lambda)Y) &= \rho(Z - (\lambda m_X + (1-\lambda)m_Y)\mathbf{1}) = \rho(Z) + \lambda m_X + (1-\lambda)m_Y \\ &\le \lambda\rho(X) + (1-\lambda)\rho(Y).\end{aligned}$$

Thus ρ is convex.

vi) Positive Homogeneity and Conic Property

If ρ is positive homogeneous, for $\lambda \ge 0$ and $X \in \mathcal{A}_\rho$, $\rho(\lambda X) = \lambda\rho(X) \le 0$, so $\lambda X \in \mathcal{A}_\rho$. Hence $\mathcal{A}_\rho$ is a cone. Conversely, if $\mathcal{A}_\rho$ is a cone, then for $\lambda > 0$,

$$m + \lambda X \in \mathcal{A}_\rho \iff \frac{m}{\lambda} + X \in \mathcal{A}_\rho.$$

Thus $\rho(\lambda X) = \inf\{m : m + \lambda X \in \mathcal{A}_\rho\} = \lambda \inf\{m' : m' + X \in \mathcal{A}_\rho\} = \lambda\rho(X)$.
We get immediatley that $\rho(0) \le 0$. Now, assume that $\rho(0) = a$, for $a < 0$ then by the representation (48) we have $\inf_{m\in}\{m\mathbf{1} \in \mathcal{A}_\rho\} = a$ but since $\mathcal{A}_\rho$ is a cone, we also have $\lambda a\mathbf{1} \in \mathcal{A}_\rho$ for arbitrarily large λ so that arbitrarily large negative constant positions are acceptable, which can not be the case, so that $a = 0$ follows. This shows that positive homogeneity of ρ and the conic property of $\mathcal{A}_\rho$ are equivalent.

□

Conversely, one can start with an acceptance set $\mathcal{A} \subseteq \mathcal{X}$ and define a risk measure based on it. For a position $X \in \mathcal{X}$, define the capital requirement as the minimal capital $m \in \mathbb{R}$ that, when added to the position, makes it acceptable:

$$\rho_{\mathcal{A}}(X) \coloneqq \inf\{m \in \mathbb{R} \colon m + X \in \mathcal{A}\}. \tag{49}$$

Proposition 4.6. Assume that $\mathcal{A}$ is a nonempty subset of $\mathcal{X}$ satisfying:

a) If $X \in \mathcal{A}$ and $Y \in \mathcal{X}$ with $Y \ge X$, then $Y \in \mathcal{A}$ (monotonicity).

b) There exists $m_0 \in \mathbb{R}$ such that $m_0\mathbf{1} \in \mathcal{A}$ (boundedness from below).

Then, the functional $\rho_{\mathcal{A}}$ defined in (49) has the following properties:

i) $\rho_{\mathcal{A}}$ is a risk measure.

ii) If $\mathcal{A}$ is convex, then $\rho_{\mathcal{A}}$ is a convex risk measure.

iii) If $\mathcal{A}$ is a cone, then $\rho_{\mathcal{A}}$ is positive homogeneous. In particular, $\rho_{\mathcal{A}}$ is a coherent risk measure if $\mathcal{A}$ is a convex cone.

iv) The acceptance set of $\rho_{\mathcal{A}}$ satisfies $\mathcal{A}_{\rho_{\mathcal{A}}} = \overline{\mathcal{A}}$, the closure of $\mathcal{A}$ with respect to the supremum norm $\|\cdot\|_\infty$. In particular, $\mathcal{A}_{\rho_{\mathcal{A}}} = \mathcal{A}$ if and only if $\mathcal{A}$ is closed in $\|\cdot\|_\infty$.

Proof. **i) $\rho_{\mathcal{A}}$ is a Risk Measure**
Monotonicity: If $X, Y \in \mathcal{X}$ with $Y \geq X$, then for any $m \in \mathbb{R}$,

$$m + Y \geq m + X.$$

If $m + X \in \mathcal{A}$, then by property (a), $m + Y \in \mathcal{A}$. Therefore,

$$\rho_{\mathcal{A}}(Y) = \inf\{m \in \mathbb{R} \colon m + Y \in \mathcal{A}\} \leq \inf\{m \in \mathbb{R} \colon m + X \in \mathcal{A}\} = \rho_{\mathcal{A}}(X).$$

Cash Invariance: For any $X \in \mathcal{X}$ and $m \in \mathbb{R}$,

$$\begin{aligned} \rho_{\mathcal{A}}(X + m) &= \inf\{k \in \mathbb{R} \colon k + X + m \in \mathcal{A}\} \\ &= \inf\{(k - m) + m + X + m \in \mathcal{A}\} \\ &= \inf\{(k - m) + X + 2m \in \mathcal{A}\} \\ &= \rho_{\mathcal{A}}(X) - m. \end{aligned}$$

Finiteness: Since $\mathcal{A}$ is nonempty and contains $m_0\mathbf{1}$, we have for any $X \in \mathcal{X}$,

$$\begin{aligned} \rho_{\mathcal{A}}(X) &\leq \inf\{m \in \mathbb{R} \colon m + X \geq m_0\mathbf{1}\} \\ &= m_0 - \operatorname{essinf} X, \end{aligned}$$

which is finite since X is bounded.
Similarly, for any $X \in \mathcal{X}$,

$$\rho_{\mathcal{A}}(X) \geq -\|X\|_\infty - \sup\{m \in \mathbb{R} \colon m\mathbf{1} \in \mathcal{A}\}.$$

Thus, $\rho_{\mathcal{A}}(X)$ is finite.
ii) Convexity of $\rho_{\mathcal{A}}$
Suppose $\mathcal{A}$ is convex. Let $X, Y \in \mathcal{X}$ and $m_X, m_Y \in \mathbb{R}$ such that $m_X + X \in \mathcal{A}$ and $m_Y + Y \in \mathcal{A}$. For $0 \leq \lambda \leq 1$, consider

$$\lambda(m_X + X) + (1 - \lambda)(m_Y + Y) = [\lambda m_X + (1 - \lambda)m_Y] + [\lambda X + (1 - \lambda)Y] \in \mathcal{A},$$

since $\mathcal{A}$ is convex. Therefore,

$$\rho_{\mathcal{A}}(\lambda X + (1 - \lambda)Y) \leq \lambda m_X + (1 - \lambda)m_Y.$$

Taking the infimum over all m_X and m_Y, we obtain

$$\rho_{\mathcal{A}}(\lambda X + (1 - \lambda)Y) \leq \lambda \rho_{\mathcal{A}}(X) + (1 - \lambda)\rho_{\mathcal{A}}(Y).$$

Thus, $\rho_{\mathcal{A}}$ is convex.
iii) Positive Homogeneity of $\rho_{\mathcal{A}}$
Suppose $\mathcal{A}$ is a cone. Let $X \in \mathcal{X}$ and $\lambda \geq 0$. Then,

$$\rho_{\mathcal{A}}(\lambda X) = \inf\{m \in \mathbb{R} \colon m + \lambda X \in \mathcal{A}\}.$$

Since $\mathcal{A}$ is a cone, $m + \lambda X \in \mathcal{A}$ if and only if $\lambda^{-1}(m + \lambda X) = \lambda^{-1}m + X \in \mathcal{A}$. Thus,

$$\rho_{\mathcal{A}}(\lambda X) = \inf\{m \in \mathbb{R} \colon \lambda^{-1}m + X \in \mathcal{A}\} = \lambda \rho_{\mathcal{A}}(X).$$

Therefore, $\rho_{\mathcal{A}}$ is positive homogeneous. The case $\lambda = 0$ follows as in the previous proposition.

If $\mathcal{A}$ is a convex cone, then $\rho_{\mathcal{A}}$ is both convex and positive homogeneous, hence a coherent risk measure.
iv) Acceptance Set of $\rho_{\mathcal{A}}$
First, note that

$$\mathcal{A}_{\rho_{\mathcal{A}}} = \{X \in \mathcal{X} \colon \rho_{\mathcal{A}}(X) \leq 0\} = \{X \in \mathcal{X} \colon 0 + X \in \mathcal{A}\} = \mathcal{A}.$$

Therefore, the acceptance set of $\rho_{\mathcal{A}}$ is exactly $\mathcal{A}$.
However, if $\mathcal{A}$ is not closed in $\|\cdot\|_\infty$, then $\rho_{\mathcal{A}}$ might assign negative values to some $X \notin \mathcal{A}$, which would imply that $X \in \mathcal{A}_{\rho_{\mathcal{A}}}$. To resolve this, consider the closure $\overline{\mathcal{A}}$ of $\mathcal{A}$.
Next, we show that $\mathcal{A}_{\rho_{\mathcal{A}}} = \overline{\mathcal{A}}$.
Step 1: $\overline{\mathcal{A}} \subseteq \mathcal{A}_{\rho_{\mathcal{A}}}$.
Let $X \in \overline{\mathcal{A}}$. Then, there exists a sequence $(X_n)_{n\in\mathbb{N}}$ in $\mathcal{A}$ such that $\|X_n - X\|_\infty \to 0$. Since $X_n \in \mathcal{A}$, $\rho_{\mathcal{A}}(X_n) \leq 0$ for all n. By the Lipschitz continuity of $\rho_{\mathcal{A}}$, we have

$$|\rho_{\mathcal{A}}(X_n) - \rho_{\mathcal{A}}(X)| \leq \|X_n - X\|_\infty \to 0,$$

so $\rho_{\mathcal{A}}(X) \leq 0$. Therefore, $X \in \mathcal{A}_{\rho_{\mathcal{A}}}$.
Step 2: $\mathcal{A}_{\rho_{\mathcal{A}}} \subseteq \overline{\mathcal{A}}$.
Let $X \in \mathcal{A}_{\rho_{\mathcal{A}}}$, so $\rho_{\mathcal{A}}(X) \leq 0$. Then, for any $\epsilon > 0$, there exists $m \leq \epsilon$ such that $m+X \in \mathcal{A}$. Therefore,

$$X = (X + m) - m,$$

with $X + m \in \mathcal{A}$ and $|m| \leq \epsilon$. As $\epsilon \to 0$, $X + m \to X$ in $\|\cdot\|_\infty$, so X is a limit point of $\mathcal{A}$. Therefore, $X \in \overline{\mathcal{A}}$. Thus, $\mathcal{A}_{\rho_{\mathcal{A}}} = \overline{\mathcal{A}}$. In particular, $\mathcal{A}_{\rho_{\mathcal{A}}} = \mathcal{A}$ if and only if $\mathcal{A}$ is closed in $\|\cdot\|_\infty$. □

4.3 Examples of Risk Measures

In this section, we present several examples of popular risk measures used in financial risk management.

Example 4.7 (Value-at-Risk (VaR))**.** Consider a probabilistic model $(\Omega, \mathcal{F}, \mathbb{P})$. A financial position X is often considered acceptable if the probability of a loss does not exceed a certain level $\lambda \in (0, 1)$. Specifically, we define the **Value-at-Risk** at confidence level $1 - \lambda$ as:

$$\mathrm{VaR}_{1-\lambda}(X) := \inf\{m \in \mathbb{R} \colon \mathbb{P}(X + m < 0) \leq \lambda\}. \tag{50}$$

Alternatively, VaR can be expressed in terms of the quantile function of X:

$$\mathrm{VaR}_{1-\lambda}(X) = -F_X^{-1}(\lambda),$$

where F_X^{-1} denotes the quantile function of X. The Value-at-Risk $\mathrm{VaR}_{1-\lambda}$ is a positively homogeneous risk measure; however, it is not convex and, therefore, not a coherent risk measure on $\mathcal{L}^0(\Omega, \mathcal{F}, \mathbb{P})$.

Remark 4.8. VaR measures the minimum amount of capital that must be added to the position X so that the probability of a loss exceeding zero is at most λ. While VaR is widely used due to its simplicity and ease of calculation, it does not account for the magnitude of losses beyond the VaR threshold.

Example 4.9 (Worst-Case Risk Measure)**.** We define the **worst-case risk measure** $\rho_{\max}$ as

$$\rho_{\max}(X) = -\operatorname{essinf} X = -\inf_{\omega\in\Omega} X(\omega), \quad \forall X \in \mathcal{X}.$$

The acceptance set of $\rho_{\max}$ is

$$\mathcal{A}_{\rho_{\max}} = \{X \in \mathcal{X} \colon X \geq 0\},$$

i.e., all financial positions that are almost surely non-negative. Thus, $\mathcal{A}_{\rho_{\max}}$ is a convex cone, and therefore, by Proposition 4.5 vi), $\rho_{\max}$ is a coherent risk measure.

Remark 4.10. The risk measure $\rho_{\max}$ is the most conservative among all normalized risk measures ρ, since for any ρ satisfying $\rho(0) = 0$, we have

$$\rho(X) \leq \rho(\operatorname{essinf} X \cdot \mathbf{1}) = \rho(0) - \operatorname{essinf} X = \rho_{\max}(X).$$

Moreover, $\rho_{\max}$ can be represented as the worst-case expected loss over all probability measures in $\mathcal{M}_1$:

$$\rho_{\max}(X) = \sup_{\mathbb{Q}\in\mathcal{M}_1} \mathbb{E}_{\mathbb{Q}}[-X],$$

where $\mathcal{M}_1$ is the set of all probability measures absolutely continuous with respect to $\mathbb{P}$ such that $\mathbb{Q}$ concentrates its mass on the worst outcomes of X.

Example 4.11 (Sharpe Ratio and Mean-Standard Deviation Risk Measure)**.** For an asset with (discounted) payoff $X_1 \in \mathcal{L}^2(\Omega, \mathcal{F}, \mathbb{P})$ and today's price X_0, the **Sharpe Ratio** is defined as:

$$\mathrm{SR}(X_1) := \frac{\mathbb{E}[X_1 - X_0]}{\sigma(X_1)},$$

where $\sigma(X_1)$ is the standard deviation of X_1.
We can define a financial position $X \in \mathcal{L}^2(\Omega, \mathcal{F}, \mathbb{P})$ as acceptable if its Sharpe Ratio is at least a certain threshold $c > 0$. The corresponding acceptance set is:

$$\mathcal{A}_{\rho_{\mathrm{SR}}} = \left\{ X \in \mathcal{L}^2(\Omega, \mathcal{F}, \mathbb{P}) \colon \frac{\mathbb{E}[X] - X_0}{\sigma(X)} \geq c \right\}.$$

The associated risk measure, called the **mean-standard deviation risk measure**, is defined as:

$$\rho_{\mathrm{SR}}(X) = -\mathbb{E}[X] + c\,\sigma(X), \quad \forall X \in \mathcal{L}^2(\Omega, \mathcal{F}, \mathbb{P}).$$

Remark 4.12. While ρ_{SR} is cash-invariant, positively homogeneous, and convex, it is not monotone on all of $\mathcal{L}^2(\Omega, \mathcal{F}, \mathbb{P})$. This means that increasing the payoff X does not necessarily decrease the risk measure $\rho_{\mathrm{SR}}(X)$, which contradicts the monotonicity property required for a risk measure as per Definition 4.1. Nevertheless, the mean-standard deviation risk measure is widely used in portfolio optimization and risk assessment due to its connection with the Markowitz mean-variance framework.

Example 4.13 (Conditional Value-at-Risk (CVaR) or Expected Shortfall)**.** The **Conditional Value-at-Risk** at confidence level $1-\lambda$, also known as **Expected Shortfall**, is defined as:

$$\mathrm{CVaR}_{1-\lambda}(X) = -\frac{1}{\lambda}\int_0^{\lambda} F_X^{-1}(u)\,du,$$

where F_X^{-1} is the quantile function of X. Alternatively, if X is a continuous random variable, CVaR can be expressed as the conditional expectation:

$$\mathrm{CVaR}_{1-\lambda}(X) = -\mathbb{E}\left[X \mid X \leq -\mathrm{VaR}_{1-\lambda}(X)\right].$$

CVaR is a coherent risk measure; it is convex, positively homogeneous, monotonic, and cash-invariant.

Remark 4.14. CVaR provides a more comprehensive assessment of tail risk compared to VaR. It accounts for the magnitude of losses in the tail of the loss distribution.

Example 4.15 (Entropic Risk Measure)**.** The **entropic risk measure** is defined for a financial position X and a risk aversion parameter $\theta > 0$ as:

$$\rho_{\mathrm{ent}}(X) = \frac{1}{\theta}\ln\left(\mathbb{E}\left[e^{-\theta X}\right]\right).$$

The entropic risk measure arises from exponential utility functions and captures the investor's risk aversion. It is convex and cash-invariant but not positively homogeneous, so it is not a coherent risk measure.

Remark 4.16. Despite not being coherent, the entropic risk measure is useful in contexts where the exponential utility function is appropriate, such as in certain insurance and risk-sensitive control applications. It reflects the trade-off between expected return and the variability of outcomes in a way that emphasizes the exponential growth of risk with increasing losses.

Example 4.17 (Mean-Variance Risk Measure)**.** The **mean-variance risk measure** is given by:

$$\rho_{\mathrm{MV}}(X) = -\mathbb{E}[X] + \frac{\lambda}{2}\operatorname{Var}(X),$$

where $\lambda > 0$ is a risk aversion parameter. This risk measure penalizes the variance of the financial position, reflecting the trade-off between expected return and risk.

Remark 4.18. Like the Sharpe Ratio, the mean-variance risk measure is cash-invariant and convex but not monotone. It forms the basis of the Markowitz portfolio optimization framework. While it does not satisfy all the properties of a risk measure as per Definition 4.1, it is instrumental in understanding the balance between expected returns and the variability of those returns.

4.3.1 Summary of Risk Measure Properties

The following table summarizes the key properties of the risk measures discussed:

Risk Measure	**Monotone**	**Cash-Invariant**	**Convex**	**Coherent**
Value-at-Risk:	Yes	Yes	No	No
Worst-Case Risk Measure:	Yes	Yes	Yes	Yes
Sharpe Ratio:	No	Yes	Yes	No
Conditional Value-at-Risk:	Yes	Yes	Yes	Yes
Entropic Risk Measure:	Yes	Yes	Yes	No
Mean-Variance Risk Measure:	No	Yes	Yes	No

4.4 Dual Representation of Convex Risk Measures ♣

An important aspect of convex risk measures is their dual representation in terms of penalty functions and probability measures.

Theorem 4.19 (Dual Representation)**.** Let ρ be a convex risk measure on $\mathcal{X} = L^\infty(\Omega, \mathcal{F}, \mathbb{P})$. Then, there exists a convex, lower semicontinuous function $\alpha\colon \mathcal{M}^1 \to [0, \infty]$, where $\mathcal{M}^1$ is the set of all probability measures on $(\Omega, \mathcal{F})$ absolutely continuous with respect to $\mathbb{P}$, such that

$$\rho(X) = \sup_{\mathbb{Q} \in \mathcal{M}^1} \left(\mathbb{E}_{\mathbb{Q}}[-X] - \alpha(\mathbb{Q}) \right), \quad \forall X \in \mathcal{X}. \tag{51}$$

Remark 4.20. The function α can be interpreted as a penalty term that quantifies the deviation of the alternative measure $\mathbb{Q}$ from the reference measure $\mathbb{P}$. The dual representation expresses the risk measure as the worst-case expected loss over a set of alternative probability measures, adjusted by the penalty.

4.5 Risk Optimal Hedging

In Section 3, we discussed the variance-optimal hedge, which provides us with a hedging strategy $\varphi^* = (W_0^*, \phi^*)$ that minimizes the mean-square hedging error

$$\mathbb{E}_{\mathbb{P}}\left[(W_T(\varphi) - H)^2\right] = \mathbb{E}_{\mathbb{P}}\left[(W_0 + G_T(\phi) - H)^2\right],$$

over all self-financing strategies. This approach is applicable in any incomplete market and, under additional but relatively mild conditions, we can construct such variance-optimal strategies explicitly (see Theorem 3.8).
However, using the mean-square error as a measure of how acceptable or 'good' a hedge is has the drawback that it treats positive and negative deviations of $W_T(\varphi)$ from the contingent claim H symmetrically. In reality, only the case $H > W_T(\varphi)$ is undesirable for the seller of the contingent claim, as it implies a loss.

4.5.1 Risk Optimal Strategies and Indifference Prices

We present an alternative method for hedging a contingent claim within a potentially incomplete financial market $(\Omega, \mathcal{F}, \mathbb{P}, \{\mathcal{F}_t\}_{t=0}^T, X)$. This approach is predicated on the viewpoint of the seller of the contingent claim, e.g., a bank. Suppose a bank's client requests a derivative contract with a payoff function $H \in \mathcal{L}^{0,+}(\mathcal{F}_T)$ and offers at time $t = 0$ the price π to enter this contract. The bank has the discretion to accept or reject this deal, depending on the proposed price π and the risk associated with the financial position $G_T(\phi) - H$ under some self-financing strategy ϕ.
Assuming that the bank evaluates its risk with a risk measure $\rho\colon \mathcal{L}^0(\mathcal{F}_T) \to \mathbb{R}$, the quantity $\rho(G_T(\phi) - H)$ represents the capital that the bank must allocate to the position $G_T(\phi) - H$ to render it acceptable from a risk perspective when trading according to the strategy ϕ.
By denoting the set of all self-financing trading strategies by $\mathcal{H}$, we define the map $\pi\colon \mathcal{L}^0(\mathcal{F}_T) \to \mathbb{R}$ as follows:

$$\pi(H) := \inf_{\phi \in \mathcal{H}} \rho(G_T(\phi) + H), \quad H \in \mathcal{L}^0(\mathcal{F}_T). \tag{52}$$

That π is again a risk measure on $\mathcal{L}^0(\mathcal{F}_T)$ is established in the following proposition.

Proposition 4.21. The function $\pi\colon \mathcal{L}^0(\mathcal{F}_T) \to \mathbb{R}$ defined in (52) is monotonic and cash-invariant. If $\mathcal{H}$ is convex and ρ is convex, then π is a convex risk measure.

Proof. We demonstrate the asserted properties of π.
Monotonicity: Let $H_1, H_2 \in \mathcal{L}^0(\mathcal{F}_T)$ such that $H_1 \leq H_2$. For any admissible strategy $\phi \in \mathcal{H}$, we have

$$G_T(\phi) - H_1 \geq G_T(\phi) - H_2.$$

Since ρ is a risk measure and thus monotonic decreasing (recall that higher financial positions have lower risk), it follows that

$$\rho(G_T(\phi) - H_1) \leq \rho(G_T(\phi) - H_2).$$

Taking the infimum over all $\phi \in \mathcal{H}$ on both sides, we obtain

$$\pi(H_1) \leq \pi(H_2),$$

which shows that π is monotonic.
Cash Invariance: Let $m \in \mathbb{R}$ and $H \in \mathcal{L}^0(\mathcal{F}_T)$. Then,

$$\begin{aligned}
\pi(H+m) &= \inf_{\phi\in\mathcal{H}} \rho(G_T(\phi) - (H+m)) \\
&= \inf_{\phi\in\mathcal{H}} \rho((G_T(\phi) - H) - m) \\
&= \inf_{\phi\in\mathcal{H}} (\rho(G_T(\phi) - H) - m) \\
&= \left(\inf_{\phi\in\mathcal{H}} \rho(G_T(\phi) - H)\right) - m \\
&= \pi(H) - m,
\end{aligned}$$

where we used the cash invariance property of ρ, namely $\rho(X-m) = \rho(X) - m$.
Convexity: Assume that $\mathcal{H}$ is convex and that ρ is convex. For any $H_1, H_2 \in \mathcal{L}^0(\mathcal{F}_T)$ and $\lambda \in [0,1]$, we have

$$\begin{aligned}
\pi(\lambda H_1 + (1-\lambda)H_2) &= \inf_{\phi\in\mathcal{H}} \rho\left(G_T(\phi) - (\lambda H_1 + (1-\lambda)H_2)\right) \\
&= \inf_{\phi\in\mathcal{H}} \rho\left(\lambda(G_T(\phi) - H_1) + (1-\lambda)(G_T(\phi) - H_2)\right).
\end{aligned}$$

Since ρ is convex, we have

$$\rho\left(\lambda(G_T(\phi) - H_1) + (1-\lambda)(G_T(\phi) - H_2)\right) \leq \lambda\rho(G_T(\phi) - H_1) + (1-\lambda)\rho(G_T(\phi) - H_2).$$

Therefore,

$$\pi(\lambda H_1 + (1-\lambda)H_2) \leq \inf_{\phi\in\mathcal{H}} \left(\lambda\rho(G_T(\phi) - H_1) + (1-\lambda)\rho(G_T(\phi) - H_2)\right).$$

Since $\mathcal{H}$ is convex, any $\phi \in \mathcal{H}$ can be represented as $\phi = \lambda\phi_1 + (1-\lambda)\phi_2$ for some $\phi_1, \phi_2 \in \mathcal{H}$. However, the infimum over $\phi \in \mathcal{H}$ of the sum above is less than or equal to $\lambda\pi(H_1) + (1-\lambda)\pi(H_2)$. Therefore,

$$\pi(\lambda H_1 + (1-\lambda)H_2) \leq \lambda\pi(H_1) + (1-\lambda)\pi(H_2),$$

showing that π is convex. □

Following our interpretation, the risk $\pi(H)$ in (52) can be seen as the minimal amount of capital required to supplement the financial position $G_T(\phi) - H$ to make it acceptable from a risk perspective when we hedge optimally. A simple heuristic might suggest that the bank should enter the trade whenever the client's premium offer π is larger than the risk associated with the financial position measured by $\pi(H)$, i.e., whenever $\pi(H) \leq \pi$, and therefore abstain if $\pi(H) > \pi$. However, this overlooks an important option that the bank retains: the option to not engage in the trade at all. In other words, the bank is not obliged to assume a liability of $-H$ at time T. To rectify this and determine the correct threshold price, one must consider the difference between the risk of entering the trade and the risk of not entering it, which is $\pi(H) - \pi(0)$.

Definition 4.22. We define the **risk indifference price** $p(H)$ of the European contingent claim H as

$$p(H) \coloneqq \pi(H) - \pi(0), \tag{53}$$

where π is given by (52).

The indifference price $p(H)$ represents the additional amount of capital (over the baseline risk $\pi(0)$) that the bank requires to accept the liability $-H$ at time T. It is convenient that the indifference price coincides with the "fair price" if the contingent claim is attainable, as shown in the next lemma.

Lemma 4.23. Suppose H is attainable, i.e., there exists a self-financing strategy $\varphi^* = (\pi^*, \phi^*)$ such that $H = \pi^* + G_T(\phi^*)$. Then $p(H) = \pi^*$; that is, the indifference price coincides with the risk-neutral price from Theorem 2.6.

Proof. For any $\phi \in \mathcal{H}$, consider

$$G_T(\phi) - H = G_T(\phi) - (\pi^* + G_T(\phi^*)) = -\pi^* + G_T(\phi - \phi^*).$$

Using the cash invariance property of ρ, we have

$$\rho(G_T(\phi) - H) = \rho(-\pi^* + G_T(\phi - \phi^*)) = \rho(G_T(\phi - \phi^*)) - \pi^*.$$

Taking the infimum over $\phi \in \mathcal{H}$ (note that $\phi - \phi^* \in \mathcal{H}$ if $\mathcal{H}$ is a vector space of strategies), we get

$$\begin{aligned}\pi(H) = \inf_{\phi\in\mathcal{H}} \rho(G_T(\phi) - H) &= \inf_{\phi\in\mathcal{H}} (\rho(G_T(\phi - \phi^*)) - \pi^*) \\ &= \left(\inf_{\psi\in\mathcal{H}} \rho(G_T(\psi))\right) - \pi^* = \pi(0) - \pi^*.\end{aligned}$$

Rearranging, we find

$$p(H) = \pi(H) - \pi(0) = -\pi^*.$$

Since the indifference price is the amount the bank requires to accept the liability $-H$, and the bank can replicate H by investing π^* at time 0, it follows that $p(H) = \pi^*$. □

Remark 4.24. The negative sign in the expression $p(H) = -\pi^*$ arises from our convention that the bank receives $p(H)$ when entering into the contract, and the liability at time T is $-H$. Thus, the initial outlay is $-p(H)$, matching the replication cost π^*.

4.6 Overture on Deep Hedging

5 Expected Utility

5.1 Lotteries

We assume that the set $\mathcal{X}$ is a convex subset of the set of all probability measures on some measurable space $(S, \mathcal{S})$. We write $\mathcal{M}$ instead of $\mathcal{X}$. Our aim is to consider preference relations on the space of lotteries that admit a numerical representation of a special kind.

Definition 5.1. Let $\succ$ be a preference relation on $\mathcal{M}$. A numerical representation U is called a **Von Neumann-Morgenstern representation** if there is a measurable function $u\colon S \to \mathbb{R}$ such that

$$U(\mu) = \int u \, \mathrm{d}\mu, \quad \forall \mu \in \mathcal{M}. \tag{54}$$

It is easy to check that a Von Neumann-Morgenstern representation U is an affine function, i.e. $U(t\mu + (1-t)\nu) = tU(\mu) + (1-t)U(\nu)$, for all $\mu, \nu \in \mathcal{M}$ and $t \in [0,1]$. But, if a numerical representation U of $\succ$ is affine, then it implies two additional properties of $\succ$ (see Proposition 3.3), that we define now.

Definition 5.2. Let $\succ$ be a preference relation on $\mathcal{M}$. It satisfies the **independence axiom** if for all $\mu \succ \nu$ it holds that

$$t\mu + (1-t)\lambda \succ t\nu + (1-t)\lambda, \tag{55}$$

for all $\lambda \in \mathcal{M}$ and $t \in (0,1]$.
The preference relation satisfies the **Archimedean axiom** (also called continuity axiom), if for all $\mu \succ \lambda \succ \nu$, there are $t, s \in (0,1)$ such that

$$t\mu + (1-t)\nu \succ \lambda \succ s\mu + (1-s)\nu. \tag{56}$$

Proposition 5.3. Assume that $\succ$ admits an affine numerical representation. Then $\succ$ satisfies the axioms of Definition 5.2.

Proof. Left as an Exercise. □

The nice thing is that Proposition 5.3 has a converse as well. This is the content of the next theorem.

Theorem 5.4. Suppose that a preference relation $\succ$ on $\mathcal{M}$ satisfies both the independence and Archimedean axioms. Then it has an affine numerical representation, say U. Moreover, for any other affine numerical representation $\overline{U}$, there exist $a > 0$ and $b \in \mathbb{R}$ such that $\overline{U} = aU + b$.

Proof. We leave the proof out of this version. □

Example 5.5. Suppose that $\mathcal{M}$ is the set of all finite mixtures of Dirac measures δ_x, and that an affine representation U exists. Define $u(x) = U(\delta_x)$. Let

$$\mu = \sum t_i \delta_{x_i},$$

where the $t_i \geq 0$ and $\sum t_i = 1$. Affinity of U yields

$$U(\mu) = \sum t_i u(x_i) = \int u \, \mathrm{d}\mu,$$

which is the desired representation. So, in this case, if there exists an affine representation, it is automatically of Von Neumann-Morgenstern type.

In the remainder of this section we assume that the set S is a separable metric space and that $\mathcal{S}$ is its Borel σ-algebra. Recall the definition of weak convergence of probability measures on S: $\mu_n \to \mu$ iff $\int f \, d\mu_n \to \int f \, d\mu$ for all bounded and continuous functions f on S. As a preparation for the final theorem, we have the following lemma.

Lemma 5.6. Consider the space $\mathcal{M}$ of all probability measures on $(S, \mathcal{S})$ endowed with the weak topology. Fix $\mu, \nu \in \mathcal{M}$ and consider $A : t \mapsto t\mu + (1-t)\nu$. Then $A : [0,1] \to \mathcal{M}$ is continuous. If $\succ$ is a continuous preference ordering on $\mathcal{M}$, then it satisfies the Archimedean axiom.

Proof. The first assertion follows from the evident identity $\int f \, d(t\mu + (1-t)\nu) = t \int f \, d\mu + (1-t) \int f \, d\nu$, valid for any bounded and continuous function f on S. Indeed, if $t_n \to t$, one then has for all bounded and continuous functions f on S that $\int f \, dA(t_n) \to \int f \, dA(t)$, which shows that $A(t)$ is the weak limit of the $A(t_n)$.
To prove the second assertion, let $\mu \succ \nu$ and choose $\lambda \in ((\nu, \mu))$. Observe that $t = 1$ is an element of $A^{-1}((\lambda, \rightarrow))$ and that this set is open in $[0,1]$ by the just shown continuity of A. Hence there is also some $t \in (0,1)$ belonging to it, and for this t one has $A(t) = t\mu + (1-t)\nu \succ \lambda$, as required in Definition 5.2 The existence of s in that definition follows similarly. $\square$

Theorem 5.7. Consider the space $\mathcal{M}$ of all probability measures on $(S, \mathcal{S})$ endowed with the weak topology, where S is assumed to be separable. Let $\succ$ be a continuous preference ordering on $\mathcal{M}$, satisfying the independence axiom. Then $\succ$ admits a Von Neumann-Morgenstern representation

$$U(\mu) = \int u \, d\mu, \tag{57}$$

where the function $u : S \to \mathbb{R}$ is bounded, continuous and unique up to affine transformations.

Proof. Consider first the subspace $\mathcal{M}_S$ of simple distributions on S, these are the distributions as in Example 5.5. We conclude from Lemma 5.6 and Theorem 5.4 that $\succ$ restricted to $\mathcal{M}_S$ admits an affine representation, which is, by Example 5.5, automatically of Von Neumann-Morgenstern type.

The function u involved will turn out to be bounded. Suppose that this is not the case, then there is a sequence $(x_n) \subset S$ such that $(u(x_n))$ is increasing and $u(x_n) > n$ (the other possibility $u(x_n) < -n$ can be treated similarly). Put $\mu_n = (1-\sqrt{1/n})\delta_{x_1}+\sqrt{1/n}\delta_{x_n}$. Since $u(x_2) > u(x_1)$, we have $\delta_{x_2} \succ \delta_{x_1}$, so $\delta_{x_1} \in ((\leftarrow, \delta_{x_2}))$. One easily checks that $\mu_n \to \delta_{x_1}$ weakly. Hence, for n big enough, μ_n belongs to any (nonempty) open neighborhood of δ_{x_1}, so eventually we have $\mu_n \in ((\leftarrow, \delta_{x_2}))$. But then $U(\mu_n) \leq u(x_2)$. However, by direct computation, we have $U(\mu_n) > (1 - \sqrt{1/n})u(x_1) + \sqrt{n}$, which yields a contradiction.
We now show that u is continuous. Suppose the contrary, then there is a sequence (x_n) converging to some $x \in S$, whereas $u(x_n)$ doesn't converge to $u(x)$. Assume e.g. that one has $\limsup u(x_n) < u(x)$. Then along a subsequence, again denoted by (x_n), one has $\lim u(x_n) =: a < u(x)$. In particular, there is $m \in \mathbb{N}$ such that $|u(x_n)-a| < \frac{1}{3}(u(x)-a)$, for $n \geq m$; equivalently $\frac{4}{3}a - \frac{1}{3}u(x) < u(x_n) < \frac{2}{3}a + \frac{1}{3}u(x)$, for $n \geq m$. Put $\mu = \frac{1}{2}(\delta_x + \delta_{x_m})$. Then also $U(\delta_x) = u(x) > \frac{2}{3}u(x) + \frac{1}{3}a > \frac{1}{2}(u(x) + u(x_m)) = U(\mu) > \frac{1}{3}u(x) + \frac{2}{3}a > U(\delta_{x_n})$, for $n \geq m$. So, $\delta_x \succ \mu \succ \delta_{x_n}$. This means that δ_{x_n} doesn't belong to the open neighborhood $((\mu, \rightarrow))$ of δ_x, contradicting the fact that $\delta_{x_n} \to \delta_x$ weakly.
We now show that, knowing the function u, Equation (57) defines a numerical representation U of $\succ$. Since u is bounded and continuous, U, as defined in (57), is continuous w.r.t. the weak topology. It is a fact that the set of simple distributions is weak-dense in the set of all probability measures on $(S, \mathcal{S})$. Since we know that U is a numerical representation of $\succ$ on the set of simple distributions, we can argue as in the proof of Theorem B.10 , that U is also a numerical representation on the collection of all probability measures on $(S, \mathcal{S})$.
Finally, u is unique up to affine transformations. This follows from Theorem 5.4 , affine numerical representations are unique up to affine transformations. □

5.2 Risk Aversion

We now delve into the realm of utility and its expectation in the context of portfolio theory. Consider a set $\mathcal{M}$ of probability measures on an interval S of $\mathbb{R}$. Let S represent the Borel σ-algebra on S. For our setup, $\mathcal{M}$ is assumed to be convex and contains all Dirac measures on points in S, hence all simple measures as well.
Typically, the fair price of a lottery $\mu \in \mathcal{M}$ is its expectation, given by

$$m(\mu) := \int_S x\mu(dx)$$

unless stated otherwise, these expectations are assumed to be finite for all $\mu \in \mathcal{M}$.

Remark 5.8. Consider concave functions $u : S \to \mathbb{R}$ satisfying the condition:

$$u(tx + (1 - t)y) \geq tu(x) + (1 - t)u(y)$$

for any $x, y \in S$ and $t \in [0, 1]$. It's worth noting the special properties of such concave functions which will be crucial for our further discussions.
A function $u\colon S \to \mathbb{R}$ exhibits strict concavity if:

$$u(tx + (1 - t)y) > tu(x) + (1 - t)u(y)$$

for distinct $x, y \in S$ and t between 0 and 1.

Risk aversion is a key concept in portfolio theory. It's observed that individuals often prefer a guaranteed outcome over a lottery with the same expected payoff, due to personal preferences.

Definition 5.9 (Monotone and Risk Averse Preference Order)**.** A preference order $\succ$ on $\mathcal{M}$ is called

i) **monotone**, if $x > y$ implies $\delta_x \succ \delta_y$ (with $x, y \in S$).

ii) **risk averse**, if $\delta_{m(\mu)} \succ \mu$, unless μ is degenerate, i.e., $\mu = \delta_{m(\mu)}$.

Proposition 5.10. If a preference order $\succ$ on $\mathcal{M}$ has a Von Neumann-Morgenstern representation given by $U(\mu) = \int u \, \mathrm{d}\mu$ then:

i) the preference order is monotone if and only if u is strictly increasing.

ii) the preference order is risk averse if and only if u is strictly concave.

Proof. (i) Notice that $U(\delta_x) = u(x)$. Then $u(x) > u(y)$ if and only if $U(\delta_x) > U(\delta_y)$ if and only if $\delta_x \succ \delta_y$.

(ii) Suppose that $\succ$ is risk averse. Take $x, y \in S$ and consider $\mu = t\delta_x + (1-t)\delta_y$ for $t \in (0, 1)$. Then $m(\mu) = tx + (1-t)y$. Then the risk averse $\succ$ yields $U(\delta_{m(\mu)}) > U(\mu)$, or $u(tx + (1-t)y) > tu(x) + (1-t)u(y)$. Hence u is strictly concave. Conversely, for strictly concave u, Jensen's inequality gives for any nondegenerate $\mu \in \mathcal{M}$ that $U(\delta_{m(\mu)}) = u(m(\mu)) > \int u \; \mathrm{d}\mu = U(\mu)$.

□

Definition 5.11 (Utility Function)**.** A function $u \colon S \to \mathbb{R}$ is called a utility function if it is strictly increasing, strictly concave, and continuous on S. A preference order $\succ$ on $\mathcal{M}$ is said to have an expected utility representation U if there exists a utility function u such that

$$U(\mu) = \int u d\mu, \quad \text{for all } \mu \in \mathcal{M}.$$

For any utility function u, a unique $c(\mu) \in S$ exists such that playing a lottery μ is indifferent to receiving a certain amount $c(\mu)$ under a given preference order.

Definition 5.12 (Certainty Equivalent and Risk Premium)**.** For a given lottery μ, the number $c(\mu)$ is its **certainty equivalent** and the difference $\rho(\mu) := m(\mu) - c(\mu)$ is termed the **risk premium**.

Notice that always $c(\mu) \leq m(\mu)$ for risk averse $\succ$ and that strict inequality holds for nondegenerate μ. Hence, a risk averse person with utility function u will not pay more than $c(\mu)$ to play a lottery μ. Conversely, the risk premium is the amount of money a

seller of the lottery μ has to pay to a risk averse agent to convince him to exchange the sure amount $m(\mu)$ for the random pay-off of the lottery μ.
In the present context, we consider the following optimization problem. Find, if it exists, a lottery μ^* that is most preferred among all lotteries in a subset of $\mathcal{M}$, equivalently, the one with the highest value of U, where U is of expected utility type.

We specialize to a specific case. Let $(\Omega, \mathcal{F}, \mathbb{P})$ be given and a random variable X defined on it, with values in S, that has a nondegenerate distribution μ. Let $c \in \mathbb{R}$ and consider the convex combination $X_\lambda = \lambda c + (1-\lambda)X$. Note that the distribution function of X_λ is obtained by a location-scale transformation of that of X. Write μ_λ for the distribution of X_λ ($\mu_0 = \mu$). Put

$$f(\lambda) := U(\mu_\lambda) = \int u \, d\mu_\lambda = \mathbb{E}\left[u(X_\lambda)\right].$$

Proposition 5.13. Assume that S is an interval, $X \geq a$ for some $a \in \mathrm{Int}S$, $\mathbb{E}_\mathbb{P}\left[X\right] < \infty$ and $c \in \mathrm{Int}S$.

i) The function $f\colon [0,1] \to \mathbb{R}$ is strictly concave and hence its maximal value is assumed for some unique $\lambda^* \in [0,1]$.

ii) We have $\lambda^* = 1$ if $m(\mu) = \mathbb{E}_\mathbb{P}\left[X\right] \leq c$, and $\lambda^* > 0$ if $c \geq c(\mu)$.

iii) If moreover u is differentiable, then we even have $\lambda^* = 1 \Leftrightarrow \mathbb{E}_\mathbb{P}\left[X\right] \leq c$ and $\lambda^* = 0 \Leftrightarrow c \leq \frac{\mathbb{E}_\mathbb{P}[u'(X)]}{\mathbb{E}_\mathbb{P}[u(X)]}$.

Proof. (i) Since $f(\lambda) = \mathbb{E}\left[u(X_\lambda)\right]$, strict concavity of f follows from exploiting first strict concavity of u and then taking expectations.
(ii) Jensen's inequality yields

$$f(\lambda) \leq u\left(\mathbb{E}\left[X_\lambda\right]\right) = u\left(\mathbb{E}\left[X\right] + \lambda(c - \mathbb{E}\left[X\right])\right),$$

with equality iff $\lambda = 1$. Since u is increasing, the right hand side is non-decreasing in λ if $c \geq \mathbb{E}\left[X\right]$. Under this condition, $\lambda^* = 1$.
Concavity of u yields $u(X_\lambda) \geq (1-\lambda)u(X) + \lambda u(c)$, hence

$$f(\lambda) \geq (1-\lambda)u(c(\mu)) + \lambda u(c),$$

with equality iff $\lambda = 0, 1$. The right hand side is non-decreasing in λ under the condition $c \geq c(\mu)$, in which case $\lambda^* > 0$.
(iii) Assume that u is differentiable. Because f is concave, $\lambda^* = 0$ can only happen if f is decreasing in a neighborhood of zero, so when the right derivative $f'_+(0) \leq 0$. Let us compute this derivative. We have

$$\frac{u(X_\lambda) - u(X)}{\lambda} = \frac{u(X_\lambda) - u(X)}{X_\lambda - X}(c - X).$$

The nonnegative difference quotient on the right, is bounded from above by the derivative of u in the left endpoint of the involved interval. Observe that for all $\lambda \in [0,1]$ one has $X_\lambda = \lambda c + (1-\lambda)X \geq \lambda c + (1-\lambda)a \geq \min\{a, c\}$. Hence the absolute value is bounded by $|u'(c \wedge a)(c - X)|$, which has finite expectation, since $|u'(c \wedge a)| < \infty$ because $c, a \in \mathrm{Int}S$

and $\mathbb{E}[|X|] < \infty$. Taking expectations and letting $\lambda \downarrow 0$, we get by the Dominated Convergence Theorem that the limit is $f'_+(0) = \mathbb{E}[u'(X)(c - X)]$. Hence $f'_+(0) \le 0$ iff $c \le \frac{\mathbb{E}[u'(X)]}{\mathbb{E}[u(X)]}$.
In much the same way, $\lambda^* = 1$ iff f is non-decreasing in a neighborhood of $\lambda = 1$, $f'_-(1) \ge 0$. Working with a difference quotient for $\lambda \uparrow 1$ and using that $X_1 = c$, we get $f'_-(1) = u'(c)(c - \mathbb{E}[X])$. The last assertion now also follows. □

Example 5.14. Consider a risky asset S_1 with price π_1, and a riskless asset with interest rate r ($S_0 = 1 + r$). Suppose that an agent has a C^1 utility function u and a capital (initial wealth) w. Suppose that he builds a portfolio by investing a fraction λ of his capital in the riskless asset and the rest in the risky asset. The value of the portfolio ("at time $t = 1$") is then $\lambda w(1 + r) + (1 - \lambda)w\frac{S_1}{\pi_1}$, and the discounted net gain is

$$\frac{wS_1(1 - \lambda)}{\pi_1(1 + r)} - \pi_1.$$

The previous proposition shows that $\lambda^* = 1$ (all capital invested in the riskless asset) iff $\frac{\mathbb{E}[S_1]}{1+r} \le \pi_1$. Hence such an agent is only willing to invest in the risky asset, when the price is below the expected discounted value. Note that this holds for any risk averse investor, regardless of the special form of the utility function u. Compare this with what happens under the risk-neutral measure.

5.2.1 Arrow-Pratt coefficient

Suppose that one considers a probability measure μ that has finite variance and that is concentrated on a small interval around its mean $m = m(\mu)$. Let u be a C^2 utility function on a neighborhood of this interval and let U be the associated expected utility representation. Look at the following heuristic.
A first order Taylor expansion of u around m gives

$$u(x) \approx u(m) + (x - m)u'(m).$$

With $x = c(\mu)$ one obtains $u(c(\mu)) \approx u(m) + (c(\mu) - m)u'(m)$.
A second order Taylor expansion of u around m gives

$$u(x) \approx u(m) + (x - m)u'(m) + \frac{1}{2}(x - m)^2 u''(m).$$

Taking expectations yields

$$u(c(\mu)) = U(\mu) = \int u d\mu \approx u(m) + \frac{1}{2}\mathrm{Var}(\mu)u''(m).$$

Hence, combining the two approximations, for the risk premium $\rho(\mu) = m - c(\mu)$ we have the approximation

$$\rho(\mu) \approx -\frac{1}{2}\frac{u''(m)}{u'(m)}\mathrm{Var}(\mu).$$

We shall see that, in spite of the rough heuristics, the right hand side of this equation contains a useful quantity.

Definition 5.15. For u, a twice differentiable utility function on some (open) interval S, the quantity

$$\alpha(x) := -\frac{u''(x)}{u'(x)}$$

is called the **Arrow-Pratt coefficient** of absolute risk aversion of u at the level x.

Given that u is strictly concave and strictly increasing, it follows that $\alpha(x) \geq 0$ for all x. Additionally, based on the preceding discussion, for probability measures μ that are tightly clustered around their mean m, the risk premium $\rho(\mu)$ can be approximated as the product of the Arrow-Pratt coefficient of absolute risk aversion at m and half the variance of μ. This factorization highlights that the risk premium depends on both the level of risk aversion at the mean and the variability of the outcomes, with the variance being an intrinsic property of μ that is independent of location.
The Arrow-Pratt coefficients possess the desirable property of remaining unchanged under affine transformations. In the context of Von Neumann-Morgenstern utility representations, where the utility function u is determined uniquely up to such transformations, the Arrow-Pratt coefficient emerges as an inherent characteristic of the preference ordering, rather than a particular numerical representation. This assertion holds provided that u is twice differentiable (i.e., $u \in C^2$) and non-constant; a constant u would imply a degenerate preference order devoid of any interesting structure.
We now proceed to introduce some commonly employed utility functions.

Example 5.16. Let u be such that the Arrow-Pratt function $\alpha(\cdot)$ is a (positive) constant, also denoted by α. Then, by solving a second order linear differential equation, one finds, for some constants $a \in \mathbb{R}$ and $b > 0$,

$$u_{a,b}(x) = a - be^{-\alpha x},$$

which is an affine transformation of $u(x) = 1 - \exp(-\alpha x)$. Note that u is defined on all of $\mathbb{R}$. The functions $u_{a,b}$ are called **CARA functions** (from Constant Absolute Risk Aversion).

Example 5.17. Here we introduce the **HARA** (from Hyperbolic Absolute Risk Aversion) utility functions. For these functions we have that $\alpha(x) = \frac{c}{x}$ for $x > 0$. For convenience we write $c = 1 - \gamma$, and hence $\gamma < 1$. Solving the corresponding differential equation for u yields

$$u_{a,b}(x) = \frac{a\gamma}{\gamma x + b}.$$

for $\gamma \neq 0$ and $ua, b(x) = a \log x + b$ for $\gamma = 0$. Note that $\gamma \geq 1$ is excluded by requiring that u is strictly concave and that for all $\gamma < 1$ it holds that $u'_{a,b}(x) = ax^{\gamma-1}$. The functions ua, b are affine transformations of $u_{1,0}$.

HARA utility functions with $\gamma > 0$ are examples of utility functions $u : [0, \infty) \to R$ satisfying the Inada conditions, i.e. $u \in C^1(0, \infty)$, with $\lim_{x \to 0} u'(x) = \infty$ and $\lim_{x \to \infty} u'(x) = 0$.

There are close connections between utility functions, risk premia and Arrow-Pratt coefficients for different preference orders.

Proposition 5.18. Suppose $u_1, u_2 : S \to R$ are two C^2 utility functions, with corresponding risk premia $\rho_1(\cdot)$, $\rho_2(\cdot)$, certainty equivalents $c_1(\cdot)$, $c_2(\cdot)$ and Arrow-Pratt coefficients $\alpha_1(\cdot)$ and $\alpha_2(\cdot)$. The following are equivalent:

i) $\alpha_1(x) \geq \alpha_2(x), \forall x \in S$.

ii) There exist a strictly increasing concave function F, defined on the range of u_2, such that $u_1 = F \circ u_2$.

iii) $\rho_1(\mu) \geq \rho_2(\mu), \forall \mu \in \mathcal{M}$.

Proof. **i)** $\Rightarrow$ (**ii**): The obvious choice of F is $F(x) = u_1(u_2^{-1}(x))$. Clearly, F is well defined, since u_2 is strictly increasing, and since u_2^{-1} and u_1 are strictly increasing, so is F. To show that F is concave, we compute its second derivative and use that (i) is assumed. Notice that it is sufficient to show that $F''(u_2(x)) \leq 0$, for all $x \in S$. We start with $u_1(x) = F(u_2(x))$ and get

$$
\begin{aligned}
u_1'(x) &= F'(u_2(x))u_2'(x) \\
u_1''(x) &= F''(u_2(x))u_2'(x)^2 + F'(u_2(x))u_2''(x).
\end{aligned}
$$

Solving the second of these two equations for $F''(u_2(x))$ and using the first one yields

$$
\begin{aligned}
F''(u_2(x)) &= \frac{u_1''(x) - u_1'(x)u_2''(x)/u_2'(x)}{u_2'(x)^2} \\
&= \frac{u_1''(x)/u_1'(x) - u_2''(x)/u_2'(x)}{u_2'(x)^2} \\
&= u_2'(x)\left(\alpha_2(x) - \alpha_1(x)\right),
\end{aligned}
$$

by definition of the Arrow-Pratt coefficients. By assumption (i) and the fact that u_1 is increasing, we have $F''(u_2(x)) \leq 0$.
(**ii**) $\Rightarrow$ (**iii**): By Jensen's inequality, applied to the concave function F, it holds that

$$
\begin{aligned}
u_1(c_1(\mu)) &= \int u_1 \, \mathrm{d}\mu \\
&= \int F \circ u_2 \, \mathrm{d}\mu \leq \int F(u_2 \, \mathrm{d}\mu) \\
&= F(u_2(c_2(\mu))) = u_1(c_2(\mu)).
\end{aligned}
$$

Since u_1 is increasing, we must have $c_1(\mu) \leq c_2(\mu)$, from which the result follows, since $\rho_1(\mu) = m(\mu) - c_1(\mu)$ and $\rho_2(\mu) = m(\mu) - c_2(\mu)$.
(**iii**) $\Rightarrow$ (**i**): Suppose that (i) doesn't hold. Then for some x one has $\alpha_1(x) < \alpha_2(x)$, and by continuity of α_1 and α_2, this equality extends to an open neighborhood O of x. By (4.3), which is also valid without assumptions (i) or (ii), we then have $F''(u_2(x)) > 0$ on O. Take now a nondegenerate probability measure μ such that $\mu(O) = 1$. Then strict convexity of $F \circ u_2$ leads to a strict equality in the opposite direction as compared to (4.4), $u_1(c_1(\mu)) > u_1(c_2(\mu))$, from which it follows that $c_1(\mu) > c_2(\mu)$,contradicting assumption (iii). $\square$

5.3 Utility-based shortfall risk

Suppose that a risk-averse investor assesses the downside risk of a financial position $X \in \mathcal{X}$ by taking the expected utility $\mathbb{E}[u(-X^-)]$ derived from the shortfall X^-, or by considering the expected utility $\mathbb{E}[u(X)]$ of the position itself. If the focus is on the downside risk, then it is natural to change the sign and to replace u by the function $l(x) := -u(-x)$. Then l is a strictly convex and increasing function, and the maximization of expected utility is equivalent to minimizing the expected loss $\mathbb{E}[l(-X)]$ or the shortfall risk $\mathbb{E}[l(X^-)]$. In order to unify the discussion of both cases, we do not insist on strict convexity. In particular, l may vanish on $(-\infty, 0]$, and in this case the shortfall risk takes the form $\mathbb{E}[l(X^-)] = \mathbb{E}[l(-X)]$.

Definition 5.19. A function $l \colon \mathbb{R} \to \mathbb{R}$ is called a loss function if it is increasing and not identically constant.

Lemma 5.20. Let $\ell \colon \mathbb{R} \to \mathbb{R}$ be continuous, non-decreasing and convex function (a so called **loss function**), then the following function defines a convex risk measure:

$$\rho(X) := \inf_{w\in\mathbb{R}} \{w + \mathbb{E}[\ell(-X-w)]\}, \quad X \in \mathcal{X}, \tag{58}$$

Proof. Let $X, Y \in \mathcal{X}$ be such that $X \leq Y$. Since ℓ is a non-decreasing function, we have $\mathbb{E}[\ell(-X-w)] \geq \mathbb{E}[\ell(-Y-w)]$ for any $w \in \mathbb{R}$ and thus $\rho(X) \geq \rho(Y)$. Let $m \in \mathbb{R}$, then

$$\rho(X+m) = \inf_{w\in\mathbb{R}} ((w+m) - m + \mathbb{E}[\ell(-X-(w+m))]) = -m + \rho(X).$$

Now, let $\lambda \in [0,1]$, then the convexity of ℓ implies

$$\begin{aligned}
\rho(\lambda X + (1-\lambda)Y) &= \inf_{w\in\mathbb{R}} (w + \mathbb{E}[\ell(-\lambda X - (1-\lambda)Y - w)]) \\
&= \inf_{v,w\in\mathbb{R}} (\lambda w + (1-\lambda)v + \mathbb{E}[\ell(\lambda(-X-w) + (1-\lambda)(-Y-v))] \\
&\leq \inf_{w\in\mathbb{R}} \inf_{v\in\mathbb{R}} \Big(\lambda(v + \mathbb{E}[\ell(-X-v)]) + (1-\lambda)(w + \mathbb{E}[\ell(-Y-w)])\Big)\Big) \\
&= \lambda\rho(X) + (1-\lambda)\rho(Y).
\end{aligned}$$

□

Example 5.21 (Utility-based shortfall risk measures)**.** Consider a utility function u on $\mathbb{R}$, a probability measure $\mathbb{Q} \in \mathcal{M}_1$, and fix some threshold $c \in \mathbb{R}$. Let us call a position X acceptable if its certainty equivalent is at least c, i.e., if its expected utility $\mathbb{E}_{\mathbb{Q}}[u(X)]$ is bounded from below by $u(c)$. Clearly, the set

$$\mathcal{A} := \{X \in X \mid \mathbb{E}_{\mathbb{Q}}[u(X)] \geq u(c)\}.$$

is nonempty, convex, and satisfies (ii) and (iii) from Proposition 4.5. Thus, $\rho_{\mathcal{A}}$ is a convex risk measure, called **utility-based shortfall risk measure**.

Example 5.22 (Entropic risk measure)**.** Let $u\colon \mathbb{R} \to \mathbb{R}$ be a strictly increasing and continuous function. For $X \in \mathcal{X} = \mathcal{L}^0(\Omega, \mathcal{F}, \mathbb{P})$ we consider the certainty equivalent of the law of X under $\mathbb{P}$ given by $c(X) = u^{-1}(\mathbb{E}[u(X)])$ for $X \in \mathcal{X}$. Then we set $\rho(X) = -X$ which is a monotone decreasing function that is cash invariant if and only if u is either linear or of the form $u(x) = a \pm b \exp(\pm\alpha X)$ for some constant $a, b, \alpha \in \mathbb{R}$. Therefore, ρ defined as $\rho(X) = -\mathbb{E}[X]$ or $\rho(X) = \pm\frac{1}{\alpha}\log \mathbb{E}[\exp(\pm\alpha X)]$ give rise to a convex risk measure, called the **entropic risk measure**.

The following lemma sheds light on the connection between utility indifference pricing and risk indifference pricing in the case of exponential utility functions.

Lemma 5.23. Let u be an exponential utility function of the form

$$u(x) = -\frac{1}{\lambda}e^{-\lambda x}, \quad \lambda > 0.$$

Define $\pi_u(H) \in \mathbb{R}$ as the solution to the utility indifference pricing equation:

$$\sup_{\phi\in\mathcal{H}} \mathbb{E}\Big[u(G_T(\phi) - H + \pi_u(H))\Big] = \sup_{\phi\in\mathcal{H}} \mathbb{E}\Big[u(G_T(\phi))\Big]. \tag{59}$$

Then it follows that

$$\pi_u(H) = \rho_{\text{ent}}(G_T(\phi^*) - H) - \rho_{\text{ent}}(G_T(\phi^*)),$$

where ρ_{ent} is the entropic risk measure defined by

$$\rho_{\text{ent}}(X) = \frac{1}{\lambda}\ln\Big(\mathbb{E}[e^{-\lambda X}]\Big).$$

In particular, the utility indifference price $\pi_u(H)$ coincides with the risk indifference price $p(H)$ with respect to the entropic risk measure.

Proof. Starting from the exponential utility function, note that for any $\phi \in \mathcal{H}$:

$$\begin{aligned}\mathbb{E}[u(G_T(\phi) - H + \pi_u(H))] &= \mathbb{E}\left[-\frac{1}{\lambda}e^{-\lambda(G_T(\phi)-H+\pi_u(H))}\right] \\ &= -\frac{1}{\lambda}e^{-\lambda\pi_u(H)}\mathbb{E}\Big[e^{-\lambda(G_T(\phi)-H)}\Big].\end{aligned}$$

Similarly,

$$\mathbb{E}[u(G_T(\phi))] = \mathbb{E}\left[-\frac{1}{\lambda}e^{-\lambda G_T(\phi)}\right] = -\frac{1}{\lambda}\mathbb{E}\Big[e^{-\lambda G_T(\phi)}\Big].$$

By the definition of $\pi_u(H)$ in (63), we have

$$\sup_{\phi\in\mathcal{H}} \mathbb{E}[u(G_T(\phi) - H + \pi_u(H))] = \sup_{\phi\in\mathcal{H}} \mathbb{E}[u(G_T(\phi))].$$

Substitute the exponential forms:

$$\sup_{\phi\in\mathcal{H}} \left(-\tfrac{1}{\lambda}e^{-\lambda\pi_u(H)}\mathbb{E}[e^{-\lambda(G_T(\phi)-H)}]\right) = \sup_{\phi\in\mathcal{H}} \left(-\tfrac{1}{\lambda}\mathbb{E}[e^{-\lambda G_T(\phi)}]\right).$$

Multiplying by $-\lambda$ and simplifying,

$$e^{-\lambda\pi_u(H)} \sup_{\phi\in\mathcal{H}} \mathbb{E}[e^{-\lambda(G_T(\phi)-H)}] = \sup_{\phi\in\mathcal{H}} \mathbb{E}[e^{-\lambda G_T(\phi)}].$$

Taking logarithms on both sides,

$$-\lambda\pi_u(H) + \ln\left(\sup_{\phi\in\mathcal{H}} \mathbb{E}[e^{-\lambda(G_T(\phi)-H)}]\right) = \ln\left(\sup_{\phi\in\mathcal{H}} \mathbb{E}[e^{-\lambda G_T(\phi)}]\right).$$

Rearrange terms:

$$\pi_u(H) = \frac{1}{\lambda}\left(\ln\left(\sup_{\phi\in\mathcal{H}} \mathbb{E}[e^{-\lambda(G_T(\phi)-H)}]\right) - \ln\left(\sup_{\phi\in\mathcal{H}} \mathbb{E}[e^{-\lambda G_T(\phi)}]\right)\right).$$

Now, define $\rho_{\text{ent}}(X) = \frac{1}{\lambda}\ln(\mathbb{E}[e^{-\lambda X}])$. The term $\ln(\sup_\phi \mathbb{E}[e^{-\lambda(G_T(\phi)-H)}])$ represents the optimal entropic evaluation of $G_T(\phi) - H$ over all strategies ϕ. By identifying the optimizing strategy ϕ^*, we have:

$$\rho_{\text{ent}}(G_T(\phi^*) - H) = \frac{1}{\lambda}\ln\left(\mathbb{E}[e^{-\lambda(G_T(\phi^*)-H)}]\right),$$

and similarly for $G_T(\phi^*)$:

$$\rho_{\text{ent}}(G_T(\phi^*)) = \frac{1}{\lambda}\ln\left(\mathbb{E}[e^{-\lambda G_T(\phi^*)}]\right).$$

Thus,

$$\pi_u(H) = \rho_{\text{ent}}(G_T(\phi^*) - H) - \rho_{\text{ent}}(G_T(\phi^*)).$$

By definition, the risk indifference price $p(H)$ with respect to the entropic risk measure also compares the entropic evaluations with and without H. Hence, $\pi_u(H) = p(H)$, showing that the utility indifference price under exponential utility equals the risk indifference price associated with the entropic risk measure. □

5.4 Stochastic dominance ♣

Results in the previous sections were depending on the preference orders, or the utility functions, at hand. In the present section, we will look at preferences that are independent of a particular choice of a utility function belonging to a certain class. The standing assumptions are that we deal with the set $\mathcal{M}$ of all probability measures on (R, B) that admit a finite expectation. As a consequence, for any utility function $u\colon R \to R$, the integrals $\int u \,\mathrm{d}\mu$ are well defined, but may take on the value $-\infty$. This holds, since every concave function has an affine function as a majorant. Indeed, since for some $a, b > 0$, one has $u(x) \leq ax + b$ for all x, it holds that $u(x)_+ \leq a|x| + b$ and hence $\int u_+ \,\mathrm{d}\mu < \infty$.

5.4.1 Uniform order

Definition 5.24. Let $\mu, \nu \in \mathcal{M}$. One says that μ is uniformly preferred over ν, denoted by $\mu \succeq_{\text{uni}} \nu$, if

$$\int u \, d\mu \geq \int u \, d\nu, \text{ for all utility functions } u : R \to R.$$

Remark 5.25. The uniform preference of the above definition is also called second order stochastic dominance. Notice that it is not a weak preference order (see Definition 2.2), since it is not complete. In Section 5.2 we will discuss first order stochastic dominance.

The next theorem gives a number of characterizations of uniform preference, there are many more. The functions f below are defined on all of R.

Theorem 5.26. There is equivalence between the following statements.

i) $\mu \succeq_{\text{uni}} \nu$.

ii) For all increasing concave functions $f : R \to R$, one has $\int f \, d\mu \geq \int f \, d\nu$.

iii) For all $c \in R$, it holds that $\int (c - x)_+ \mu(dx) \leq \int (c - x)_+ \nu(dx)$.

iv) If F_μ and F_ν are the distribution functions of μ and ν respectively, then $\int_{-\infty}^{c} F_\mu(x) \, dx \leq \int_{-\infty}^{c} F_\nu(x) \, dx$, for all $c \in R$.

Proof. The direction i) $\Leftrightarrow$ ii) is obvious. Let us turn to ii) $\Rightarrow$ i). For the converse implication we need a utility function that has finite integral under μ and ν. This can be accomplished as follows. Take a given utility function u and an arbitrary $x_0 \in R$. Modify u on $(-\infty, x_0]$ by replacing u with $x \mapsto u'_+(x_0)(2(x - x_0) - \exp(x - x_0) + 1) + u(x_0)$. Check that the modified function is still a utility function! Moreover, the modified utility function (denoted u again) has finite integral w.r.t. any probability measure with finite expectation. If f is increasing and concave, then $u_\alpha(x) := \alpha f(x) + (1 - \alpha)u(x)$ defines a strictly increasing, strictly concave continuous function, so a utility function, for every $\alpha \in [0, 1)$. Note that the integral $\int u_\alpha \, d\mu$ is now always well defined, possibly taking the value $-\infty$. The assertion follows from

$$\int f \, d\mu = \lim_{\alpha \to 1} \int u_\alpha \, d\mu \geq \lim_{\alpha \to 1} \int u_\alpha \, d\nu = \int f \, d\nu.$$

ii) $\Leftrightarrow$ iii): Clearly ii) $\Rightarrow$ iii). The converse implication basically follows from the fact that every nonnegative convex decreasing function, with limit zero at infinity, is a pointwise limit of positive linear combinations of functions $x \mapsto (c - x)_+$ and that $-f$ is decreasing and convex. More formally, we have that $h = -f$ admits right derivatives $h'_+(x)$ at every point x. The function h' is increasing, right continuous and on any interval $(a, b]$, up to scaling, it is a distribution function of a probability measure. Stated otherwise, there is a measure γ on (R, B) such that $\gamma(a, b] = h'_+(b) - h'_+(a)$, for all $a < b$. Since there exists only countably many discontinuity points of h', we have for $x < b$

$$h(x) = h(b) - \int_{(x,b]} h'(y)\,dy = h(b) - \int_{(x,b]} (h'_+(y) - h'_+(b))\,dy - h'_+(b)(b - x). \tag{60}$$

We first rewrite the integral in (60). Let $B = \{(u, y) : x < y < u \leq b\}$, we have:

$$\begin{aligned}
&\int_{(x,b]} \int \left(h'_+(y) - h'_+(b)\right) \mathrm{d}y \\
&= -\int_{(x,b]} \int \gamma(y, b]\,\mathrm{d}y \\
&\int_{(x,b]} -\int \mathbf{1}_{(y,b]}\,\mathrm{d}\gamma\,\mathrm{d}y \\
&= \int -\int \mathbf{1}_B(u, y)\gamma(\mathrm{d}u)\,\mathrm{d}y \\
&= \int -\int \mathbf{1}_B(u, y)\,\mathrm{d}y\gamma(\mathrm{d}u) \text{ (by Fubini)} \\
&= \int -\mathbf{1}_{(x,b]}(u)(u - x)\gamma(\mathrm{d}u) \\
&= \int -\mathbf{1}_{(-\infty,b]}(u - x)_+\gamma(\mathrm{d}u).
\end{aligned}$$

Hence, going back to (60), we can rewrite $h(x)$ as

$$h(x) = h(b) - h'_+(b)(b - x) + \int 1_{(-\infty,b]}(u - x)_+\gamma(\mathrm{d}u).$$

Let μ be a probability measure on (R, B). Integration of the last expression w.r.t. μ and using Fubini's theorem again, yields

$$\begin{aligned}
\int_{(-\infty,b]} \int h\,\mathrm{d}\mu &= h(b)\mu(-\infty, b] - h'_+(b) \int (b - x)_+\mu(\mathrm{d}x) \\
&\quad + \int_{(-\infty,b]} \int (u - x)_+\mu(dx)\gamma(du) \\
&= h(b)\mu(-\infty, b] - h'_+(b) \int (b - x)_+\mu(dx) \\
&\quad + \int_{(-\infty,b]} \int (u - x)_+\mu(\mathrm{d}x)\gamma(\mathrm{d}u).
\end{aligned}$$

Using condition iii) and the fact that $h'_+ \leq 0$, we have an upper bound for the last displayed expression by replacing μ with ν. It follows that

$$\int_{(-\infty,b]} \int hd\mu \leq \int_{(-\infty,b]} \int hd\nu + h(b)(\mu(-\infty, b] - \nu(-\infty, b]).$$

Since h is lower bounded by an affine function, we have that $\int_{(b,\infty)} hd\mu$ and $\int_{(b,\infty)} hd\nu$ are both finite. Hence we obtain

$$
\begin{aligned}
\int_{(b,\infty)} \int h d\mu \leq & \int_{(b,\infty)} \int h d\nu \\
& + \int_{(b,\infty)} \int h d\mu - \int_{(b,\infty)} \int h d\nu \\
& + h(b)(\mu(-\infty, b] - \nu(-\infty, b]) \\
& = \int_{(b,\infty)} \int h d\nu - \int_{(b,\infty)} (h(b) - h(x))\mu(dx) \\
& + \int_{(b,\infty)} (h(b) - h(x))\nu(dx).
\end{aligned}
$$

We finally show that the last two integrals vanish for $b \to \infty$. Since they are similar, we treat only the first of the two. Fix b_0 and let $b > b_0$. It holds that

$$
0 \leq h(b) - h(x) \leq -h'_+(b_0)(x - b_0) \text{ for } x > b.
$$

Hence

$$
\int_{(b,\infty)} (h(b) - h(x))\mu(dx) \leq -h'_+(b_0) \int (x - b_0) 1_{(b,\infty)}(x) d\mu,
$$

which tends to zero by the Dominated convergence theorem, since $\int |x|\mu(dx)$ is finite. Hence we obtain $\int h d\mu \leq \int h d\nu$, which is equivalent to ii). **iii) $\Leftrightarrow$ iv):** This is just a matter of rewriting, using Fubini's theorem. One has

$$
\begin{aligned}
\int_{-\infty}^{c} F_\mu(y) dy &= \int_{-\infty}^{c} \int 1_{(-\infty,y]}(x)\mu(dx)dy \\
&= \int_{(-\infty,c]} \int dy \mu(dx) \\
&= \int_{(-\infty,c]} (c - x)\mu(dx) = \int (c - x)_+ \mu(dx). \qquad (61)
\end{aligned}
$$

The integral with F_ν can be rewritten in similar terms and the equivalence of iii) and iv) becomes obvious. □

Remark 5.27. It follows from Theorem 5.26 ii), that $\mu \succeq_{\text{uni}} \nu$ implies $m(\mu) \geq m(\nu)$. The integrals w.r.t. the measure μ in assertion iii) of the same theorem in fact determine μ. Indeed, by the computations leading to (61), we see that knowing integrals of $(c - x)_+$ for all c is equivalent to knowing the integrals of F_μ up to c. Taking right derivatives w.r.t. c gives $F_\mu(c)$ and knowing this for all c determines μ. This fact can be used to show that $\succeq_{\text{uni}}$ defines a partial order (left as an exercise).

If μ is the distribution of a random variable $X \in L^1(\Omega, F, P)$ and ν that of $Y \in L^1(\Omega, F, P)$, such that $\mu \succeq_{\text{uni}} \nu \succeq_{\text{uni}} \mu$, then $\mu = \nu$, so X and Y have the same distribution (under P). Yet, X and Y are in general very different as random variables. It may happen that $P(X = Y) = 0$.
When two lotteries with the same mean are compared, we can develop the assertions of Theorem 5.3 a little further.

Proposition 5.28. For all probability measures $\mu, \nu \in \mathcal{M}$ the following are equivalent.

i) $\mu \succeq_{\text{uni}} \nu$ and $m(\mu) = m(\nu)$.

ii) $\int f d\mu \geq \int f d\nu$, for all concave functions f.

iii) $m(\mu) \geq m(\nu)$ and $\int (x-c)_+ \mu(dx) \leq \int (x-c)_+ \nu(dx)$, for all $c \in \mathbb{R}$.

Proof. i) $\Rightarrow$ ii): First we show that the assertion holds true for decreasing concave functions. Such a function is $x \mapsto -(c-x)_-$, for arbitrary $c \in \mathbb{R}$. Since $-(c-x)_- = c - x - (c-x)_+$, the assertion for such a function follows from Theorem 5.26 and the assumptions that $m(\mu) = m(\nu)$ and $\mu \succeq_{\text{uni}} \nu$, because $x \mapsto -(c-x)_+$ is concave and increasing. The proof for arbitrary decreasing concave functions is then similar to the proof of iii) $\Rightarrow$ ii) of Theorem 5.26. The second assertion of Theorem 5.26 also tells us that ii) is true for increasing concave functions, and hence ii) holds for monotone concave functions. If f is concave, but not monotone, then there exists a $x_0 \in \mathbb{R}$, such that

$$f(x) \leq f(x_0), \text{ for all } x \in \mathbb{R}.$$

Let

$$f_1(x) = \begin{cases} f(x) & \text{if } x \leq x_0 \\ f(x_0) & \text{if } x > x_0 \end{cases}$$

and

$$f_2(x) = \begin{cases} f(x_0) & \text{if } x \leq x_0 \\ f(x) & \text{if } x > x_0 \end{cases}.$$

Then f_1 is concave and increasing and f_2 is concave and decreasing. Knowing that the assertions hold true for f_1 and f_2, we obtain the same result for f, because $f(x) = f_1(x) + f_2(x) - f(x_0)$ and integration of the constant is the same for each probability measure.
ii) $\Rightarrow$ iii): Take first $f(x) \equiv x$ to get the first assertion, and then $f(x) \equiv -(x-c)+$, which is concave, to get the second one from ii).
iii) $\Rightarrow$ i): Rewrite the inequality between the integrals in iii) as

$$\int_{(c,\infty)} x\mu(dx) - c + c\mu(-\infty, c] \leq \int_{(c,\infty)} x\nu(dx) - c + c\nu(-\infty, c].$$

Let $c \to -\infty$ and use that both measures have a finite first moment to conclude that $c\mu(-\infty, c]$ and $c\nu(-\infty, c]$ tend to zero as well as $\int_{(c,\infty)} x\mu(dx) \to \int x\mu(dx)$ and $\int_{(c,\infty)} x\nu(dx) \to \int x\nu(dx)$. One then arrives at $\int_{(c,\infty)} x\mu(dx) \leq \int x\nu(dx)$, or $m(\mu) \leq m(\nu)$. Together with the assumption, this gives $m(\mu) = m(\nu)$. To prove $\mu \succeq_{\text{uni}} \nu$ we use the identity $y^+ = y + (-y)^+$ (for $y \in \mathbb{R}$) to get

$$\int (c-x)^+ \mu(dx) = c - m(\mu) + (x-c)^+ \mu(dx).$$

A similar equality holds for ν. Using the assumption and the just proved identity $m(\mu) = m(\nu)$, we arrive at $\int (c-x)^+ \mu(dx) \leq \int (c-x)^+ \nu(dx)$, condition iii) in Theorem 5.26 to get $\mu \succeq_{\text{uni}} \nu$. $\square$

5.4.2 Monotone order

We turn to another concept of stochastic dominance, also called first order stochastic dominance. There are more of these concepts conceivable.

Definition 5.29. Let μ, ν be two probability measures on $(\mathbb{R}, \mathcal{B})$. One says that μ stochastically dominates ν, if for all bounded increasing continuous functions $f : \mathbb{R} \to \mathbb{R}$ it holds that

$$\int f\, d\mu \geq \int f\, d\nu \tag{5.3}$$

In this case one writes $\mu \succeq_{\text{mon}} \nu$.

It is almost trivial to check that $\succeq_{\text{mon}}$ defines a partial order on the space of probability distributions on $(\mathbb{R}, \mathcal{B})$. Below we give an easy characterization of $\mu \succeq_{\text{mon}} \nu$.

Proposition 5.30. Let μ, ν be two probability measures on $(\mathbb{R}, \mathcal{B})$ and let F_μ and F_ν be their distribution functions. The following are equivalent:

i) It holds that $\mu \succeq_{\text{mon}} \nu$.

ii) For all $x \in \mathbb{R}$ one has $F_\mu(x) \leq F_\nu(x)$.

Proof. i) $\Rightarrow$ ii): We'd like to apply the definition of stochastic dominance to the function $u \mapsto 1_{(x,\infty)}(u)$, which is bounded and increasing. The result would then follow. However, this function is not continuous. Therefore, one first uses the functions $u \mapsto (\min\{n(u - x), 1\})^+$ and let $n \to \infty$.
ii) $\Rightarrow$ i): Let f be continuous, bounded, and increasing. We can obtain f (which is measurable) as the pointwise limit of an increasing sequence of simple functions f_n, that are increasing themselves. To see this, we assume for simplicity that $0 \leq f \leq 1$ and we follow the usual approximation scheme, known from measure theory.
Let $n \in \mathbb{N}$ and define $E_{ni} = \{(i-1)2^{-n} < f \leq i2^{-n}\}$ for $i = 1, \ldots, 2^n$. Put

$$f_n = 2^{-n} \sum_{i=1}^{2^n} (i-1) 1_{E_{ni}}.$$

Then we know that $f_n \leq f$ and $f_n \uparrow f$. Using that, for each n, the E_{ni} with $i = 0, \ldots, 2^n$ are disjoint, $S_{i \geq j+1} E_{ni} = \{f > j2^{-n}\}$ and $\{f > 1\} = \emptyset$, we rewrite

$$f_n = 2^{-n} \sum_{i=1}^{2^n} \left(\sum_{j=1}^{i-1} 1 \right) 1_{E_{ni}} = 2^{-n} \sum_{j=1}^{2^n - 1} \sum_{i=j+1}^{2^n} 1_{E_{ni}} = 2^{-n} \sum_{j=1}^{2^n} 1_{\{f > j2^{-n}\}}.$$

Since f is continuous, the sets $\{f > j2^{-n}\}$ are open and since f is increasing, there are real numbers a_{nj} such that $\{f > j2^{-n}\} = (a_{nj}, \infty)$. Hence,

$$\int f_n\, d\mu = 2^{-n} \sum_{j=1}^{2^n} \mu((a_{nj}, \infty)) = 2^{-n} \sum_{j=1}^{2^n} (1 - F_\mu(a_{nj})).$$

It follows from the assumption that $\int f_n\, d\mu \geq \int f_n\, d\nu$. The assertion follows by application of the Monotone Convergence Theorem. $\square$

Portfolio Optimization

Preface to Part III: In this last part of the course "Portfolio Theory", we delve into portfolio optimization in an incomplete market setting, employing the concepts of expected utility and risk measures that we introduced in Part II. Indeed, we set-up portfolio optimization problems using expected utility or risk as preference criteria and optimize over trading strategies the expected utility or risk of a hedged financial position, terminal wealth, consumption or terminal wealth and consumption.

We differentiate between static and dynamic portfolio optimization:

i) Static portfolio optimization involves a one-time trading strategy, typically at the outset. Here, a portfolio is constructed either to hedge a contingent claim or as a "classical investment", where in both cases the agent wants to optimize the expected utility or risk from its terminal financial position without adjusting the portfolio holdings dynamically in time. This form of optimization constitutes a simpler problem than its dynamic counterpart and if one would incorporate trading cost in our consideration, a static trading strategy is cheaper to set up compared to a dynamic trading strategy. However the overall performance, in particular for hedging, might be much worse, as no adjustments to new trading information is possible.

ii) Dynamic portfolio optimization entails a trading strategy that evolves over time, as described in Part I. It allows for trading at multiple points $t_0, t_1, \dots, t_n$, with each subsequent portfolio decision potentially influenced by previous choices. Finding an optimal trading strategy leads to a much more intricate and challenging problem than the static approach. However, the option to adjust the portfolio and react to market and trading events dynamically is very desirable for the agent. In Section 8 we show how dynamic programming and the martingale method can help to decompose the dynamic portfolio optimization into simpler, static optimization or hedging problems.

Note that both the static and dynamic optimization problems often lack closed-form solutions. Nonetheless, under certain conditions, such as specific asset price models and assumptions, and with particular criteria (like CRRA or CARA expected utility), insights into optimal strategies are attainable as we will see in Section 7.2 and Section 7.3 below.

The setting is as follows: Let $(\Omega, \mathcal{F}, \mathbb{P}, \mathbb{F}, X)$ denote a finite discrete-time financial market with $(d+1) \in \mathbb{N}$-tradable assets. Assume that the market is free of arbitrage, but possibly incomplete. Let $u\colon \mathbb{R} \to \mathbb{R}$ be a *utility function*, recall that u is increasing and concave. In this section, we usually think of the following examples:

- **CARA** (Constant Absolute Risk Aversion) utility functions $u_{a,b}$ of the form

$$u_{a,b}(x) = a - be^{-\alpha x}, \quad a \in \mathbb{R}, b > 0, \alpha > 0.$$

- **CRRA** (Constant Relative Risk Aversion) utility functions u_λ of the form

$$u_\lambda(x) := \begin{cases} \frac{x^{1-\lambda}}{1-\lambda}, & x > 0, \\ -\infty, & x \leq 0, \end{cases} \qquad \text{and if } \lambda = 1, \text{ set } u_1(x) := \begin{cases} \ln(x), & x > 0 \\ -\infty, & x \leq 0. \end{cases}$$

Suppose an investor with initial capital W_0 expresses her monetary preference for terminal wealth in terms of expected utility. Then she naturally wants to maximize:

$$\mathbb{E}_{\mathbb{P}}\left[u(W_T(\varphi))\right] = \mathbb{E}_{\mathbb{P}}\left[u(W_0 + G_T(\phi))\right], \tag{62}$$

over all self-financing strategies φ such that the budget constraint $W_0(\varphi) \leq W_0$ is satisfied.

Definition 5.31. We call a self-financing strategy $\varphi^* = (W_0^*, \phi^*)$ a **utility optimal strategy** with respect to utility function u and capital constraint W_0^*, if φ^* maximizes the expected utility of the terminal wealth over all other self-financing strategies $\varphi = (W_0, \phi)$ with $W_0 \leq W_0^*$, i.e.

$$\mathbb{E}_{\mathbb{P}}\left[u(W_0^* + G_T(\varphi))\right] \leq \mathbb{E}_{\mathbb{P}}\left[u(W_0^* + G_T(\varphi^*))\right]$$

for all predictable $\mathbb{R}^d$-valued processes ϕ.

Let us consider as second agent a bank, which is interested in selling a contingent claim $H \in \mathcal{L}^{0,+}(\mathcal{F}_T)$. Since the market is incomplete, there might be an entire interval $\Pi(H)$ of arbitrage-free prices for the claim H. Assume that also the bank's preference is given by the expectec utility of some function u, then the bank faces the two options:

i) Selling the contigent claim for some price π_0 and trading according to a strategy $\varphi = (\pi_0, \phi)$ that maximizes (62) for $W_0 = \pi_0$,

ii) or not selling the claim H and trade according to a strategy with initial capital 0, that maximizes (62) with $W_0 = 0$.

Definition 5.32. We call the price $\pi_u^* \in \mathbb{R}^+$ the **utility indifferent price** of the European contingent claim H, if it satisfies:

$$\sup_{\phi} \mathbb{E}_{\mathbb{P}}\left[u(\pi_u^* + G_T(\phi) - H)\right] = \sup_{\phi} \mathbb{E}_{\mathbb{P}}\left[u(G_T(\phi))\right], \tag{63}$$

and the suprema on both sides are attained. In the following, we denote by $\varphi_H^* = (\pi_H^*, \phi_H^*)$ and $\varphi^* = (0, \phi^*)$ the optimal strategies of the left- and right-hand side of (63), respectively. We refer to $\phi := \phi_H^* - \phi^*$ as the **utility-based hedging strategy**.

Remark 5.33. Following equation (63), the utility indifferent price represents exactly the threshold price at which the bank exhibits no preference, from an expected utility perspective, between participating in and abstaining from selling the claim H. The utility-based hedging strategy describes the adjustment to the bank's optimal portfolio necessitated by the derivative trade.

6 Static Portfolio Optimization

In this section, we consider a *static* portfolio optimization problem, where trading occurs only at two times: today ($t = 0$) and at a future terminal date ($t = T$). The investor wishes to set up a portfolio ϕ today that maximizes their expected utility of terminal wealth $W_T(\phi)$ according to a utility function u. For simplicity, we focus on the pure maximization problem without considering hedging positions, i.e., we set $H = 0$. We will show that this problem has at least one solution if and only if the market $(\Omega, \mathcal{F}, \mathbb{P}, \{\mathcal{F}_t\}_{t=0}^T, X)$ is arbitrage-free.

6.1 Static Utility Optimal Strategies and the Absence of Arbitrage

Let X_0 denote the vector of asset prices at time $t = 0$, and X_T the vector of asset prices at time T. Let r denote the risk-free interest rate (assumed constant), so that the risk-free asset grows to $(1+r)^T$ over the period $[0, T]$.
Define the vector of **discounted net gains** as

$$Y := \frac{X_T}{(1+r)^T} - X_0.$$

Thus, Y represents the discounted profit or loss per unit of each asset over the period $[0, T]$.
Consider an initial investment of W_0, with a portfolio ϕ such that the initial cost is $\phi \cdot X_0 = W_0$. The terminal wealth of the investor is

$$W_T(\phi) = \phi \cdot X_T = (1+r)^T \left(W_0 + \phi \cdot Y\right).$$

Therefore, maximizing the expected utility of terminal wealth $W_T(\phi)$ is equivalent to maximizing the expected utility of $W_0 + \phi \cdot Y$.
Since $(1+r)^T$ is a constant, we can define a new utility function $\tilde{u}$ by

$$\tilde{u}(x) = u\left((1+r)^T x\right),$$

so that

$$u\left(W_T(\phi)\right) = \tilde{u}\left(W_0 + \phi \cdot Y\right).$$

Therefore, the original optimization problem is equivalent to maximizing $\mathbb{E}_{\mathbb{P}}\left[\tilde{u}\left(W_0 + \phi \cdot Y\right)\right]$.
Since W_0 is given and $\tilde{u}$ is strictly increasing, maximizing $\mathbb{E}_{\mathbb{P}}\left[\tilde{u}\left(W_0 + \phi \cdot Y\right)\right]$ over ϕ is equivalent to maximizing $\mathbb{E}_{\mathbb{P}}\left[\tilde{u}\left(\phi \cdot Y\right)\right]$ over ϕ.

Assumption 6.1. Let $u : D \to \mathbb{R}$ be a utility function and Y the vector of discounted net gains. Assume either of the following:

i) $D = \mathbb{R}$ and u is bounded from above, or

ii) $D = [a, \infty)$ for some $a < 0$, and we optimize over the set of φ such that $\varphi \cdot Y \geq a$ a.s. In this case, we also assume that for those φ the expected utility $\mathbb{E}_{\mathbb{P}}\left[u(\varphi \cdot Y)\right]$ is finite.

In both cases above, we define

$$\Xi := \left\{\phi \in \mathbb{R}^d : \phi \cdot Y \in D \text{ a.s.}\right\}.$$

We are thus led to study the following unconstrained optimization problem:
Static Utility Maximization Problem: Let $u : D \to \mathbb{R}$ be a utility function. Maximize

$$\mathbb{E}_{\mathbb{P}}\left[u\left(\phi \cdot Y\right)\right]$$

over all portfolios ϕ such that $\phi \cdot Y \in D$ almost surely. Theorem 6.3 below shows that the maximization problem only makes sense in an arbitrage-free market, just as pricing of portfolios and derivatives. We need a lemma first:

Lemma 6.2. i) Let $h : \mathbb{R}^d \to \mathbb{R} \cup \{-\infty\}$ be a concave and upper semicontinuous function with $h(0) > -\infty$. Then h attains its supremum on $\mathbb{R}^d$ if for all $\phi \neq 0$,

$$\lim_{\alpha\to\infty} h(\alpha\phi) = -\infty.$$

ii) Let $u : D \to \mathbb{R}$ be a utility function, where $D = [a, \infty)$ with $a < 0$. Let $0 \leq b < -a$, and let $X \geq 0$ be a random variable. Then for all $\alpha \in (0, 1]$, the implication

$$\mathbb{E}_{\mathbb{P}}\left[u\left(\alpha X - b\right)\right] < \infty \quad \Rightarrow \quad \mathbb{E}_{\mathbb{P}}\left[u\left(X\right)\right] < \infty$$

holds.

Proof. Proof of (i): Since h is concave and upper semicontinuous, its epigraph is a closed convex set. The condition $\lim_{\alpha\to\infty} h(\alpha\phi) = -\infty$ for all $\phi \neq 0$ implies that h is coercive, i.e., it tends to $-\infty$ at infinity in every direction. Therefore, the level sets $\{\phi \in \mathbb{R}^d : h(\phi) \geq c\}$ are compact for any $c \in \mathbb{R}$. By the Weierstrass theorem, a continuous function on a compact set attains its maximum. However, since h may not be continuous everywhere, but upper semicontinuity suffices to ensure that the supremum is achieved. Proof of (ii): Let $\alpha \in (0, 1]$ and suppose that

$$\mathbb{E}_{\mathbb{P}}\left[u\left(\alpha X - b\right)\right] < \infty.$$

Since u is concave and strictly increasing, we can use the inequality

$$u(X) \leq \frac{1}{\alpha} u\left(\alpha X - b\right) + \left(1 - \frac{1}{\alpha}\right) u\left(\frac{b}{1-\alpha}\right).$$

Note that $\frac{b}{1-\alpha} \in D$ because $0 \leq b < -a$ and $\alpha \in (0, 1]$.
Taking expectations, we get

$$\mathbb{E}_{\mathbb{P}}\left[u\left(X\right)\right] \leq \frac{1}{\alpha}\mathbb{E}_{\mathbb{P}}\left[u\left(\alpha X - b\right)\right] + \left(1 - \frac{1}{\alpha}\right) u\left(\frac{b}{1-\alpha}\right).$$

Since the right-hand side is finite by assumption and $u\left(\frac{b}{1-\alpha}\right)$ is finite, it follows that $\mathbb{E}_{\mathbb{P}}\left[u\left(X\right)\right] < \infty$.

□

Theorem 6.3. Let $u : D \to \mathbb{R}$ be a utility function satisfying Assumption 6.1, and let Y be the vector of discounted net gains. Then a maximizer in the utility maximization problem exists if and only if the market is free of arbitrage.

Proof. We will prove the two implications separately. Suppose the market admits an arbitrage opportunity; that is, there exists a portfolio $\phi_0 \in \mathbb{R}^d$ such that

$$\phi_0 \cdot Y \geq 0 \quad \text{a.s.,} \quad \text{and} \quad \mathbb{P}(\phi_0 \cdot Y > 0) > 0.$$

Consider the sequence of portfolios $\phi_n = n\phi_0$ for $n \in \mathbb{N}$. Then,

$$\phi_n \cdot Y = n(\phi_0 \cdot Y) \geq 0 \quad \text{a.s.,}$$

and $\phi_n \cdot Y$ tends to $+\infty$ on the set $\{\phi_0 \cdot Y > 0\}$.
Since u is strictly increasing and defined on D, which is either $\mathbb{R}$ or $[a, \infty)$ with $a < 0$, and u is bounded above in the first case or satisfies $u(+\infty) = +\infty$ in the second case, it follows that

$$\lim_{n\to\infty} u(\phi_n \cdot Y) = +\infty \quad \text{a.s.}$$

Therefore,

$$\lim_{n\to\infty} \mathbb{E}_{\mathbb{P}}[u(\phi_n \cdot Y)] = +\infty.$$

This contradicts the existence of a maximizer since the expected utility can be made arbitrarily large, implying that no finite maximum exists.
Assume the market is arbitrage-free. We aim to show that the utility maximization problem

$$\max_{\phi\in\Xi} \mathbb{E}_{\mathbb{P}}[u(\phi \cdot Y)]$$

attains its maximum. Consider the function $h : \Xi \to \mathbb{R} \cup \{-\infty\}$ defined by

$$h(\phi) = \mathbb{E}_{\mathbb{P}}[u(\phi \cdot Y)].$$

We will show that h satisfies the conditions of Lemma 6.2, ensuring that h attains its supremum on Ξ. First, note that h is concave in ϕ because u is concave and the expectation operator is linear.
Second, we need to verify that h is upper semicontinuous (u.s.c.) on Ξ. Since u is continuous and $\phi \cdot Y$ is continuous in ϕ, it follows that $u(\phi \cdot Y)$ is continuous in ϕ for almost every $\omega \in \Omega$. By Fatou's lemma and dominated convergence (under appropriate integrability conditions), we can conclude that h is u.s.c.
Third, we must show that for all $\phi \neq 0$,

$$\lim_{\alpha\to\infty} h(\alpha\phi) = -\infty.$$

Suppose not; then there exists a sequence $\alpha_n \to \infty$ such that $h(\alpha_n\phi)$ remains bounded from below. However, since the market is arbitrage-free, we cannot have $\phi \cdot Y$ tending to $+\infty$ with non-negligible probability. Moreover, due to the concavity and behavior of u at infinity (either bounded above or $u(-\infty) = -\infty$), it follows that $h(\alpha\phi) \to -\infty$ as $\alpha \to \infty$.
Finally, since h is concave, u.s.c., and satisfies the condition $\lim_{\alpha\to\infty} h(\alpha\phi) = -\infty$, Lemma 6.2 implies that h attains its maximum on Ξ. Therefore, a maximizer exists.

□

Theorem 6.4. Let $u\colon D \to \mathbb{R}$ be a continuously differentiable utility function satisfying Assumption 6.1. Assume, additionally, that $\mathbb{E}_{\mathbb{P}}\left[|u(\phi \cdot Y)|\right] < \infty$ for all $\phi \in \Xi$. Let the maximizing ϕ^* be an interior point of Ξ. Then $Yu'(\phi^* \cdot Y) \in L^1(\Omega, \mathcal{F}, \mathbb{P})$ and

$$\mathbb{E}_{\mathbb{P}}\left[Yu'(\phi^* \cdot Y)\right] = 0. \tag{64}$$

Proof. If differentiation and expectation commute, one has

$$\nabla_\phi \mathbb{E}_{\mathbb{P}}\left[u(\phi \cdot Y)\right] = \mathbb{E}_{\mathbb{P}}\left[u'(\phi \cdot Y)Y\right],$$

and the result follows by taking $\phi = \phi^*$. Since it is not clear that the commutation is valid, we directly show that the right hand side is zero at $\phi = \phi^*$. Take $\eta \in \mathbb{R}^d$ and $\varepsilon \in (0,1]$. Put $\phi_\varepsilon = \phi^* + \varepsilon\eta$, then $\phi_\varepsilon \in \phi$ for all ε sufficiently small, $\varepsilon < \varepsilon_0$ say. For those ε we put $f(\varepsilon) := u(\phi_\varepsilon \cdot Y)$ and

$$\Delta_\varepsilon := \frac{u(\phi_\varepsilon \cdot Y) - u(\phi^* \cdot Y)}{\varepsilon} = \frac{\eta \cdot Y}{\varepsilon\eta \cdot Y}.$$

Note that $\mathbb{E}_{\mathbb{P}}\left[\Delta_\varepsilon\right] \le 0$, because $\mathbb{E}_{\mathbb{P}}\left[u(\phi^* \cdot Y)\right]$ is maximal. Concavity of u gives that f is concave too. Hence Δ_ε is increasing for $\varepsilon \downarrow 0$, with limit $\eta \cdot Yu'(\phi^* \cdot Y)$. The assumption that $u(\phi \cdot Y) \in L^1(\Omega, \mathcal{F}, \mathbb{P})$ for all $\phi \in \Xi$ implies that $\Delta_{\varepsilon_0} \in L^1(\Omega, \mathcal{F}, \mathbb{P})$. Hence $\Delta_\varepsilon - \Delta_{\varepsilon_0}$ is nonnegative and increasing for $\varepsilon \downarrow 0$, which enables us to apply the Monotone convergence theorem to get

$$0 \ge \mathbb{E}_{\mathbb{P}}\left[\Delta_\varepsilon\right] \uparrow \mathbb{E}_{\mathbb{P}}\left[\eta \cdot Yu'(\phi^* \cdot Y)\right],$$

where the expectation on the right hand side is a finite number. We conclude that $\eta \cdot \mathbb{E}_{\mathbb{P}}\left[Yu'(\phi^* \cdot Y)\right] \le 0$ for all $\eta \in \mathbb{R}^d$. So we can replace η with $-\eta$ in the last inequality and we conclude that the linear map $\eta \mapsto \eta \cdot \mathbb{E}_{\mathbb{P}}\left[Yu'(\phi^* \cdot Y)\right]$ is identically zero. But then we must have $\mathbb{E}_{\mathbb{P}}\left[Yu'(\phi^* \cdot Y)\right] = 0$. □

Proposition 6.5. Under the assumptions of Theorem 6.4, and assuming the market is arbitrage-free, define

$$\frac{d\mathbb{P}^*}{d\mathbb{P}} = \frac{u'(\phi^* \cdot Y)}{\mathbb{E}_{\mathbb{P}}\left[u'(\phi^* \cdot Y)\right]}. \tag{65}$$

Then $\mathbb{P}^*$ is an equivalent martingale measure, i.e., a risk-neutral measure.

Proof. First, we show that $\mathbb{E}\left[u'(\phi^* \cdot Y)\right] < \infty$, so that $\mathbb{P}^*$ is well defined. Define

$$c := \sup\{u'(x) \colon x \in D \text{ and } x \in [-|\phi^*|, |\phi^*|]\}.$$

Consider first the case in which $D = \mathbb{R}$. Then, because u' is decreasing, we have $c = u'(-|\phi^*|)$. If $D = [a, \infty)$, then $c \le \sup\{u'(x) : x \in D\} = u'(a)$. In both cases, we have $c < \infty$. By the Cauchy-Schwarz inequality, we have

$$|\phi^* \cdot Y| \le |\phi^*| \cdot |Y|.$$

Hence, if $|\phi^* \cdot Y| > |\phi^*|$, then $|Y| > 1$. From this, it follows that (we split into the cases $|\phi^* \cdot Y| \le |\phi^*|$ and $|\phi^* \cdot Y| > |\phi^*|$ and use that u' is nonnegative)

$$0 \le u'(\phi^* \cdot Y) = u'(\phi^* \cdot Y)1\{|\phi^* \cdot Y| \le |\phi^*|\} + u'(\phi^* \cdot Y)1\{|\phi^* \cdot Y| > |\phi^*|\}$$

$$
\begin{aligned}
&\leq c1\{|\phi^* \cdot Y| \leq |\phi^*|\} + u'(\phi^* \cdot Y)1\{|Y| > 1\} \\
&\leq c + u'(\phi^* \cdot Y)1\{|Y| > 1\} \\
&\leq c + u'(\phi^* \cdot Y)|Y|1\{|Y| > 1\} \\
&\leq c + u'(\phi^* \cdot Y)|Y|,
\end{aligned}
$$

where the expression on the right-hand side has finite expectation, by Theorem 6.4.
Next, we need to verify that $\mathbb{P}^*$ is a probability measure equivalent to $\mathbb{P}$. Since u' is positive (as u is strictly increasing and continuously differentiable), and $\mathbb{E}_{\mathbb{P}}\left[u'\left(\phi^* \cdot Y\right)\right] < \infty$, the Radon-Nikodym derivative is well-defined and positive, and integrates to 1:

$$
\int_\Omega \frac{d\mathbb{P}^*}{d\mathbb{P}}\, d\mathbb{P} = \mathbb{E}_{\mathbb{P}}\left[\frac{u'\left(\phi^* \cdot Y\right)}{\mathbb{E}_{\mathbb{P}}\left[u'\left(\phi^* \cdot Y\right)\right]}\right] = 1.
$$

Therefore, $\mathbb{P}^*$ is a probability measure equivalent to $\mathbb{P}$.
Next, we need to show that $\mathbb{P}^*$ is a martingale measure, i.e.,

$$
\mathbb{E}_{\mathbb{P}^*}[Y] = 0.
$$

Indeed,

$$
\begin{aligned}
\mathbb{E}_{\mathbb{P}^*}[Y] &= \int_\Omega Y\, d\mathbb{P}^* = \int_\Omega Y \frac{d\mathbb{P}^*}{d\mathbb{P}}\, d\mathbb{P} \\
&= \frac{1}{\mathbb{E}_{\mathbb{P}}\left[u'\left(\phi^* \cdot Y\right)\right]} \mathbb{E}_{\mathbb{P}}\left[Y u'\left(\phi^* \cdot Y\right)\right] \\
&= 0,
\end{aligned}
$$

where the last equality follows from (64).
Therefore, under $\mathbb{P}^*$, the discounted net gains Y have zero expectation, implying that discounted asset prices are martingales under $\mathbb{P}^*$.

□

We will revisit this result in more specific (dynamic) settings in Sections 7.2 and 7.3.

7 Optimal Dynamic Hedging Strategies

This section addresses the problem of portfolio optimization within a dynamic framework. In contrast to the static case, investment decisions are not limited to a single point in time and trades can be sequentially executed at all trading times $0 = t_0 < t_1 < \ldots < t_n = T$.

The dynamic framework affords investors the capability to realign their portfolio allocations at various stages, adapting to the unfolding market dynamics. Nonetheless, the computation of optimal portfolios is challenging due to the high dimensionality of the space of trading strategies. A prevalent method to surmount this complexity is *dynamic programming*, which decomposes the problem into a series of simpler, recursively solved stages. In certain instances, the optimization problem can be explicitly solved, or meaningful insights about the solutions can be deduced.
In Section 7.1, our analysis begins with utility optimal strategies in the dynamic setting. Specifically, we apply martingale methods to derive optimal trading strategies within a specified geometric model, employing the Constant Relative Risk Aversion (CRRA) utility as the criterion. Additionally, we establish a correspondence between the optimal solutions for strategies based on Constant Absolute Risk Aversion (CARA) utility and the concept of relative entropy. Following this, Section 3 explores the principles of variance-optimal hedging. Contrary to maximizing expected utility, this strategy is concerned with minimizing the variance of the residual hedging error. We will also examine the role of risk measures and related hedging strategies.

7.1 Dynamic utility optimal trading strategies

Let us come back to one of our primary objectives of this section, the maximization of the expected utility of terminal wealth $\mathbb{E}_{\mathbb{P}}\left[u(G_T(\varphi))\right]$, $\mathbb{E}_{\mathbb{P}}\left[u(W_T(\varphi))\right]$ or $\mathbb{E}_{\mathbb{P}}\left[u(W_T(\varphi) - H)\right]$. Here, φ represents a trading strategy and H is the pay-off function of a contingent claim sold by our agent. Solving the two utility maximization problems on both sides of equation (63) allows for the derivation of the utility-indifferent price π_u^*.
In this part, we present a martingale method to tackle the optimization problems presented on both sides of

$$\sup_{\phi} \mathbb{E}_{\mathbb{P}}\left[u(\pi_u(H) + G_T(\phi) - H)\right] = \sup_{\phi} \mathbb{E}_{\mathbb{P}}\left[u(G_T(\phi))\right], \tag{66}$$

with $\pi_u(H) \in \mathbb{R}^+$. This can be done using different stochastic optimization approaches. In this section however, we present a version of the so called *martingale method* to construct optimal solutions through the construction of martingale measures. The derivations in this section are based on the following result, that establishes a sufficient condition for the optimality of a self-financing trading strategy in the sense of expected utility.

Theorem 7.1. Assume that the utility function u is differentiable and let $\varphi = (W_0, \phi)$ be a self-financing strategy. Define a probability measure $\mathbb{Q} \sim \mathbb{P}$ by

$$\frac{\mathrm{d}\mathbb{Q}}{\mathrm{d}\mathbb{P}} = \frac{u'(\widetilde{W}_T(\varphi))}{\mathbb{E}_{\mathbb{P}}\left[u'(\widetilde{W}_T(\varphi))\right]}. \tag{67}$$

If $\mathbb{Q}$ is an equivalent martingale measure, then φ constitutes a utility optimal strategy.

Proof. Let $\psi = (W_0', \phi')$ denote another self-financing strategy with $W_0 = W_0'$. Then the concavity of u implies

$$u(\widetilde{W}_T(\psi)) - u(\widetilde{W}_T(\varphi)) \leq u'(\widetilde{W}_T(\varphi))(\widetilde{W}_T(\psi) - \widetilde{W}_T(\varphi)) \; = \mathbb{E}\left[u'(\widetilde{W}_T(\varphi))\right] \frac{\mathrm{d}\mathbb{Q}}{\mathrm{d}\mathbb{P}} G_T(\psi - \varphi).$$

As $G_T(\psi - \varphi)$ is a $\mathbb{Q}$-martingale by assumption and since $G_T(\psi - \varphi)$ is a martingale transform by definition, we deduce that $\mathbb{E}_{\mathbb{Q}}\left[G_T(\psi - \varphi)\right] = 0$ and therefore

$$\mathbb{E}\left[u(\widetilde{W}_T(\psi))\right] \leq \mathbb{E}\left[u(\widetilde{W}_T(\varphi))\right],$$

which concludes the proof. □

Remark 7.2. Note that (67) can be interpreted as a first-order condition for extrema similar to (65) in the static case. Indeed, since the utility function u is strictly concave, any local maxima of $\varphi \mapsto \mathbb{E}\left[u(W_T(\varphi))\right]$ is a global maxima and unique. Following the usual steps to find extremas of a function, we set the first-derivative of the function to zero. If we assume $\widetilde{W}_T(\varphi) = \widetilde{G}_T(\varphi)$ and by using the chain-rule, we obtain something as follows:

$$\frac{\mathrm{d}}{\mathrm{d}\varphi}\mathbb{E}_{\mathbb{P}}\left[u(W_T(\varphi))\right] = e'(u(g(\varphi)))u'(g(\varphi))g'(\varphi),$$

for $e'(\cdot) = \mathbb{E}_{\mathbb{P}}\left[\cdot\right]$, $g'(\varphi) = \widetilde{S}_T - \widetilde{S}_0$ with $g(\varphi) = $ "$\int_0^T \varphi \, \mathrm{d}\widetilde{S}$". We thus obtain:

$$\begin{aligned}
\frac{\mathrm{d}}{\mathrm{d}\varphi}\mathbb{E}_{\mathbb{P}}\left[u(W_T(\varphi))\right] &= \mathbb{E}_{\mathbb{P}}\left[u'(\widetilde{G}_T(\varphi))(\widetilde{S}_T - \widetilde{S}_0)\right] \\
&= \mathbb{E}_{\mathbb{Q}}\left[u'(\widetilde{G}_T(\varphi))\frac{\mathbb{E}_{\mathbb{P}}\left[u'(\tilde{W}_T(\varphi))\right]}{u'(\tilde{W}_T(\varphi))}(\widetilde{S}_T - \widetilde{S}_0)\right] \\
&= \mathbb{E}_{\mathbb{P}}\left[u'(G_T(\varphi))\right]\mathbb{E}_{\mathbb{Q}}\left[\widetilde{S}_T - \widetilde{S}_0\right],
\end{aligned}$$

which equals zero, whenever $(\widetilde{S}_t)_{t\in\mathbf{T}}$ is a $\mathbb{Q}$-martingale, i.e., if the condition in Theorem 7.1 above holds true.

In Section 8 we provide a more general approach, which can be used to solve the utility based hedging problem using dynamic programming to optimize both sides of (66). In the following, we show that in certain cases we can give more information on the utility based hedge.

7.2 Utility indifferent pricing and hedging with CRRA

Let us consider that $(\Omega, \mathcal{F}, \mathbb{P}, \mathbb{F}, X)$ constitutes an arbitrage-free financial market and that $u\colon \mathbb{R} \to \mathbb{R}$ is a differentiable CRRA utility function.
In this section, we assume that the asset price process $X = (S^{(0)}, S^{(1)}, \ldots, S^{(d)})$ has the following geometric form:

$$S_t^{(0)} = S_0^{(0)}\,\mathrm{e}^{rt}, \quad t = 0, 1, \ldots, T,$$

and for $i = 1, 2, \ldots, d$ we set

$$S_t^{(i)} = S_0^{(i)} \, \mathrm{e}^{R_t^{(i)}} = S_0^{(i)} \prod_{j=1}^{t} (1 + \Delta \widehat{R}_j^{(i)}), \quad t = 0, 1, \ldots, T,$$

where $\Delta \widehat{R}_t^{(i)} = \mathrm{e}^{\Delta R_t^{(i)}} - 1$ and $\Delta R_1^{(i)}, \Delta R_2^{(i)}, \ldots, \Delta R_t^{(i)}$ are independent and identically distributed. The discounted price processes for $i = 1, 2, \ldots, d$ then assume the form

$$\widetilde{S}_t^{(i)} = \widetilde{S}_0^{(i)} \prod_{j=1}^{t} (1 + \Delta \widetilde{R}_j^{(i)}), \tag{68}$$

with $\widetilde{R}_0^{(i)} = 0$ and $\Delta \widetilde{R}_j^{(i)} = \frac{1 + \Delta \widehat{R}_j^{(i)}}{1+r} - 1$.

Lemma 7.3. Assume X is an adapted process defined on the discrete-time index set $\mathbf{T} = \{0, 1, \ldots, T\}$, and let $\Delta X_t := X_t - X_{t-1}$ for $t \geq 1$. Consider the difference equation

$$Z = 1 + Z_{-} \bullet X,$$

where $Z_{n-} = Z_{n-1}$ for each n, and $Z_0 = 1$. There exists a unique adapted process Z that solves this equation. The solution Z is called the **stochastic exponential** of X and is denoted by $\mathcal{E}(X)$. Moreover, the stochastic exponential satisfies

$$\mathcal{E}(X)_t = \prod_{j=1}^{t} (1 + \Delta X_j)$$

for all $t \in \mathbf{T}$.

Proof. Suppose Z and Z' are two processes satisfying $Z = 1 + Z_{-} \bullet X$ and $Z' = 1 + Z'_{-} \bullet X$. In particular, at time 0, we have $Z_0 = Z'_0 = 1$.
Assume inductively that $Z_{n-1} = Z'_{n-1}$. Then at time n,

$$Z_n = 1 + Z_{n-1} X_n \quad \text{and} \quad Z'_n = 1 + Z'_{n-1} X_n.$$

By the induction hypothesis, $Z_{n-1} = Z'_{n-1}$, hence

$$Z'_n = 1 + Z'_{n-1} X_n = 1 + Z_{n-1} X_n = Z_n.$$

Since $Z_0 = Z'_0$, an induction shows $Z = Z'$ at all times. Thus, the solution is unique.
Starting from $Z = 1 + Z_{-} \bullet X$, we write it in discrete form:

$$Z_t = 1 + \sum_{j=1}^{t} Z_{j-1} \Delta X_j.$$

Instead of expanding directly, consider increments:

$$Z_t - Z_{t-1} = Z_{t-1} \Delta X_t.$$

Rearrange this to get a recursive formula:

$$Z_t = Z_{t-1} (1 + \Delta X_t).$$

Since $Z_0 = 1$, we can solve this recursion by iteration:

$$Z_1 = Z_0(1 + \Delta X_1) = 1 + \Delta X_1,$$
$$Z_2 = Z_1(1 + \Delta X_2) = (1 + \Delta X_1)(1 + \Delta X_2),$$

and so forth. By induction on t, we find

$$Z_t = \prod_{j=1}^{t}(1 + \Delta X_j).$$

This product representation is precisely the definition of the stochastic exponential $\mathcal{E}(X)$. Hence,

$$\mathcal{E}(X)_t = \prod_{j=1}^{t}(1 + \Delta X_j).$$

The proof is now complete. □

Using Lemma 7.3 we can write the the discounted asset price processes $\widetilde{S}^{(i)}$ in (68) as

$$S_t^{(i)} = \widetilde{S}_0^{(i)} \mathcal{E}(\widetilde{R}^{(i)})_t \text{ for } t = 0, 1, \ldots, T.$$

We denote by $\varphi^* = (\pi^*, \phi^*)$ the solution to the left-hand side of (66) with $\pi_u^* = \pi^*$, which we assume exists and via Theorem 7.1 give rise to an equivalent martingale measure $\mathbb{Q}_H$ defined by

$$\frac{\mathrm{d}\mathbb{Q}_H}{\mathrm{d}\mathbb{P}} := \frac{u'(\widetilde{W}_T(\varphi) - H)}{\mathbb{E}_{\mathbb{P}}\left[u'(\widetilde{W}_T(\varphi) - H)\right]}. \tag{69}$$

Similarly, we let $\hat{\varphi} = (0, \hat{\phi})$ be the solution to the right-hand side of (66), which is also assumed to exists and giving rise to an equivalent martingale measure $\mathbb{Q}_0$ defined by

$$\frac{\mathrm{d}\mathbb{Q}_0}{\mathrm{d}\mathbb{P}} := \frac{u'(\widetilde{W}_T(\varphi))}{\mathbb{E}_{\mathbb{P}}\left[u'(\widetilde{W}_T(\varphi))\right]}. \tag{70}$$

In the following theorem we show how sometimes one can 'reverse engineer' the utility optimal trading strategy ϕ such that a given model $(\widetilde{S}_t)_{t\in\mathbf{T}}$ is a martingal under $\mathbb{Q}_0$. We do this in case of the CRRA utility function and the geometric model above.

Theorem 7.4. Given $\lambda \in (0, \infty)$ as a risk aversion parameter, we define the utility function, $u\colon \mathbb{R} \to \mathbb{R} \cap \{-\infty\}$, as follows: If $\lambda \neq 1$, then

$$u(x) := \begin{cases} \frac{x^{1-\lambda}}{1-\lambda}, & x > 0, \\ -\infty, & x \leq 0. \end{cases} \tag{71}$$

If $\lambda = 1$, we define $u(x)$ differently as:

$$u(x) := \begin{cases} \ln(x), & x > 0 \\ -\infty, & x \leq 0. \end{cases} \tag{72}$$

Further, let $\gamma \in \mathbb{R}^d$ be such that $\gamma^\intercal \Delta \widetilde{R}_1 > -1$ and assume that it solves the following equatioń:

$$\mathbb{E}_{\mathbb{P}}\left[\frac{\Delta\widetilde{R}_1}{(1+\gamma^\intercal\Delta\widetilde{R}_1)^\lambda}\right] = 0. \tag{73}$$

We define the process $(\widetilde{W}_t)_{t\geq 0}$ as

$$\widetilde{W}_t := \widetilde{W}_0\,\mathcal{E}(\gamma^\intercal\widetilde{R})_t, \quad t = 0,1,\dots,T, \tag{74}$$

and for all $i = 1,\dots,d$, we set

$$\phi_t^{(i)} := \frac{\gamma^{(i)}}{\widetilde{S}_{t-1}^{(i)}}\widetilde{W}_{t-1}, \quad t = 0,1,\dots,T, \tag{75}$$

$$\varphi_t^{(0)} := \widetilde{W}_{t-1} - (\phi^{(1)},\dots,\phi^{(d)})_t(\widetilde{S}^{(1)},\dots,\widetilde{S}^{(d)})^\intercal_{t-1}, \quad t = 0,1,\dots,T. \tag{76}$$

Then $\phi = (\phi^{(0)},\phi^{(1)},\dots,\phi^{(d)})$ is a utility optimal strategy for utility function u and intial wealth $W_0 := S_0^{(0)}\widetilde{W}_0$. Moreover, the discounted wealth process $\widetilde{W}((W_0,\phi))$ is equal to $\widetilde{W}$ defined in (74).

Proof. Let $\widetilde{W}_0$ be a discounted initial wealth and let ϕ be as in (75). Then we have

$$\begin{aligned}
\widetilde{W}_t\Big((\widetilde{W}_0,\phi)\Big) &= \widetilde{W}_0 + \sum_{i=1}^d \Big(\phi^{(i)} \bullet \widetilde{S}^{(i)}\Big)_t \\
&= \widetilde{W}_0 + \sum_{i=1}^d \Big(\frac{\gamma^{(i)}}{\widetilde{S}_-^{(i)}}\widetilde{W}_- \bullet \widetilde{S}^{(i)}\Big)_t \\
&= \widetilde{W}_0 + \mathcal{E}(\gamma^\intercal\widetilde{R})_- \sum_{i=1}^d \Big(\frac{\gamma^{(i)}}{\widetilde{S}_-^{(i)}} \bullet (\widetilde{S}_-^{(i)} \bullet \widetilde{R}^{(i)})\Big)_t \\
&= \widetilde{W}_0 + \mathcal{E}(\gamma^\intercal\widetilde{R})_- \sum_{i=1}^d \Big(\gamma^{(i)} \bullet \widetilde{R}^{(i)}\Big)_t \\
&= \widetilde{W}_0 + \Big(\mathcal{E}(\gamma^\intercal\widetilde{R})_- \bullet \gamma^\intercal\widetilde{R}^{(i)}\Big)_t \\
&= \widetilde{R}_0\mathcal{E}(\gamma^\intercal\widetilde{R}),
\end{aligned}$$

which proves that $\widetilde{R}_t^{(i)} = \widetilde{W}_t$ for all $t \in \mathbf{T}$, where $\varphi^{(0)}$ is specified by (76).
Next, we show that $\varphi = (\widetilde{W}_0,\phi)$ is utility optimal for utility u given by (71) or (72).
For this set

$$\alpha := \Big(\mathbb{E}_{\mathbb{P}}\left[(1+\gamma^\intercal\Delta\widetilde{R}_1)^{-\lambda}\right]\Big)^{1/\lambda}, \tag{77}$$

$$Z_t := (\alpha^t\mathcal{E}(\gamma^\intercal\widetilde{R})_t)^{-\lambda}, \quad t = 0,1,\dots,T. \tag{78}$$

The process $(Z_t)_{t\in\mathbf{T}}$ will be our density process, that we will use to define an equivalent martingale measure. Note that for all $t = 1,\dots,T$ we have

$$Z_t = Z_{t-1}\alpha^{-\lambda}\Big(\frac{\mathcal{E}(\gamma^\intercal\widetilde{R})_t}{\mathcal{E}(\gamma^\intercal\widetilde{R})_{t-1}}\Big)^{-\lambda} = Z_{t-1}(\alpha(1+\gamma^\intercal\Delta\widetilde{R}_t))^{-\lambda}.$$

Therefore, $Z_t = \prod_{i=1}^t (\alpha(1+\gamma^\intercal \Delta \widetilde{R}_t))^{-\lambda} = \mathcal{E}(M)_t$ with

$$M_t = \sum_{j=1}^t \Big((\alpha(1+\gamma^\intercal \Delta \widetilde{R}_j))^{-\lambda} - 1 \Big).$$

Now, note that by definition of α and the fact that $(\Delta \widetilde{R}_t)_{t=1,\dots,T}$ are iid, yields

$$\begin{aligned}
\mathbb{E}_{\mathbb{P}}\left[M_t | \mathcal{F}_{t-1}\right] &= M_{t-1} + \mathbb{E}_{\mathbb{P}}\left[(\alpha(1+\gamma^\intercal \Delta \widetilde{R}_t))^{-\lambda} - 1 | \mathcal{F}_t\right] \\
&= \frac{\mathbb{E}_{\mathbb{P}}\left[(1+\gamma^\intercal \Delta \widetilde{R}_t)^{-\lambda}\right]}{\mathbb{E}_{\mathbb{P}}\left[(1+\gamma^\intercal \Delta \widetilde{R}_1)^{-\lambda}\right]} - 1 \\
&= M_{t-1},
\end{aligned}$$

such that $(M_t)_{t\in\mathbf{T}}$ is a martingale. Since $Z = \mathcal{E}(M)$ it thus follows that Z is a martingale as well. We thus define the measure $\mathbb{Q}$ through its density by $\frac{\mathbb{Q}}{\mathbb{P}} = Z_T$ and note that for all $t = 1, \dots, T$ we have

$$\begin{aligned}
\mathbb{E}_{\mathbb{Q}}\left[\Delta \widetilde{R}_t | \mathcal{F}_{t-1}\right] &= \mathbb{E}_{\mathbb{P}}\left[\Delta \widetilde{R}_t \frac{Z_t}{Z_{t-1}} | \mathcal{F}_{t-1}\right] \\
&= \mathbb{E}_{\mathbb{P}}\left[\Delta \widetilde{R}_t \alpha^{-\lambda} (1+\gamma^\intercal \Delta \widetilde{R}_t)^{-\lambda} | \mathcal{F}_{t-1}\right] \\
&= \alpha^{-\lambda} \mathbb{E}_{\mathbb{P}}\left[\Delta \widetilde{R}_t (1+\gamma^\intercal \Delta \widetilde{R}_t)^{-\lambda}\right] \\
&= 0.
\end{aligned}$$

Therefore $(\widetilde{R}_t)_{t\in\mathbf{T}}$ is a martingale with respect to $\mathbb{Q}$ and so is $\widetilde{S}^{(i)} = S_0^{(i)} \mathcal{E}(\widetilde{R}^{(i)})$ for all $i = 1, 2, \dots, d$. Since $\mathbb{Q}$ by definition of Z and the properties of the utility function u satisfies:

$$Z_T = \frac{u'(W_T(\phi))}{\mathbb{E}_{\mathbb{P}}\left[u'(W_T(\phi))\right]},$$

therefore by Theorem 7.1 the trading strategy $\varphi = (W_0, \phi)$ is utility optimal for utility u and initial capital W_0. □

Note that equation (73) represents a system of d equations with d unknowns, $\gamma^{(1)}, \dots, \gamma^{(d)}$. For some concrete models, these equations can be explicitly solved.
Suppose $\Delta \widetilde{R}_t$ is sufficiently small. In this case, the approximation $(1+\gamma^\intercal \Delta \widetilde{R}_1)^{-\lambda} \approx 1 - \lambda \gamma^\intercal \Delta \widetilde{R}_1$ is valid. Consequently, equation (73) simplifies to

$$\mathbb{E}_{\mathbb{P}}\left[\Delta \widetilde{R}_1\right] \approx \lambda \mathbb{E}_{\mathbb{P}}\left[\Delta \widetilde{R}_1 \Delta \widetilde{R}_1^\intercal\right] \gamma.$$

This leads to the following approximation for γ:

$$\gamma \approx \frac{1}{\lambda} \left(\mathbb{E}_{\mathbb{P}}\left[\Delta \widetilde{R}_1 \Delta \widetilde{R}_1^\intercal\right] \right)^{-1} \mathbb{E}_{\mathbb{P}}\left[\Delta \widetilde{R}_1\right],$$

where $\left(\mathbb{E}_{\mathbb{P}}\left[\Delta \widetilde{R}_1 \Delta \widetilde{R}_1^\intercal\right] \right)^{-1}$ is approximately the inverse of the covariance matrix of $\Delta \widetilde{R}_1$.
This result can be interpreted in the context of investment strategy: An investor, whether adhering to a power utility or log utility, tends to maintain a constant relative proportion of wealth in each asset. The degree of risk aversion influences the allocation between the non-risky bank account and the risky assets. Notably, this optimal portfolio strategy is independent of the investment horizon T, a finding that contradicts common financial advice which suggests allocating more to riskier assets for longer investment horizons due to their potential for higher returns.

7.3 Utility indifferent pricing and hedging with CARA

In this section, we focus on utility indifferent pricing and hedging using a CARA utility function of the form $u(x) = 1 - \exp(-\lambda x)$, where $\lambda > 0$ denotes the *risk aversion parameter*. Note that the shifted utility of terminal wealth $u(\widetilde{W}_T(\varphi)) - 1$ depends on the initial capital W_0, only via the constant factor $\mathrm{e}^{-\lambda \widetilde{W}_0}$ and therefore W_0 does not affect the optimality of strategies and the indifferent price when considering expected utility with CARA as criterion. Therefore, without loss of generality, we can assume that $W_0 = 0$. In this section we do not impose any particular model for the asset price process $X = (S^{(0)}, \ldots, S^{(d)})$. We denote the set of all trading strategies by $\mathcal{H}$. Instead, we assume directly that the utility maximization problems on the left and right-hand side of

$$\sup_{\phi \in \mathcal{H}} \mathbb{E}_{\mathbb{P}} \left[u(\pi_u(H) + G_T(\phi) - H) \right] = \sup_{\phi \in \mathcal{H}} \mathbb{E}_{\mathbb{P}} \left[u(G_T(\phi)) \right],$$

have solutions, i.e. the supremum is attained, and that the optimal strategies $\varphi^* = (0, \phi^*)$ and $\varphi^*_H = (\pi^*_H, \phi^*_H)$ are linked via Theorem 7.1 to some equivalent martingale measures $\mathbb{Q}^*$ and $\mathbb{Q}^*_H$, i.e., we assume that

a) $\sup_{\phi \in \mathcal{H}} \mathbb{E} \left[u(\widetilde{G}_T(\phi)) \right] = \mathbb{E} \left[u(\widetilde{G}_T(\phi^*)) \right]$ holds and that $\mathbb{Q}^*$ given by

$$\frac{\mathrm{d}\mathbb{Q}^*}{\mathrm{d}\mathbb{P}} = \frac{u'(\widetilde{G}_T(\phi^*))}{\mathbb{E}_{\mathbb{P}} \left[u'(\widetilde{G}_T(\phi^*)) \right]} \tag{79}$$

is an equivalent martingale measure.

b) $\sup_{(\pi, \phi) \in \mathbb{R} \times \mathcal{H}} \mathbb{E} \left[u(\tilde{\pi} + \widetilde{G}_T(\phi) - H) \right] = \mathbb{E} \left[u(\tilde{\pi}^*_H + G_T(\phi^*_H) - H) \right]$ holds and $\mathbb{Q}^*_H$, given by

$$\frac{\mathrm{d}\mathbb{Q}^*_H}{\mathrm{d}\mathbb{P}} = \frac{u'(\widetilde{W}_T(\varphi^*_H) - H)}{\mathbb{E}_{\mathbb{P}} \left[u'(\widetilde{W}_T(\varphi^*_H) - H) \right]} \tag{80}$$

is an equivalent martingale measure.

Definition 7.5. Let $\mathbb{P} \sim \mathbb{Q}$ be two equivalent probability measures. We define the **relative entropy** or **Kullback-Leibler divergence** $D_{\mathrm{KL}}(\mathbb{Q}, \mathbb{P})$ of $\mathbb{Q}$ and $\mathbb{P}$ as

$$D_{\mathrm{KL}}(\mathbb{Q}, \mathbb{P}) := \mathbb{E}_{\mathbb{Q}} \left[\log(\frac{\mathrm{d}\mathbb{Q}}{\mathrm{d}\mathbb{P}}) \right] = \mathbb{E}_{\mathbb{P}} \left[\frac{\mathrm{d}\mathbb{Q}}{\mathrm{d}\mathbb{P}} \log(\frac{\mathrm{d}\mathbb{Q}}{\mathrm{d}\mathbb{P}}) \right]. \tag{81}$$

Lemma 7.6. For any two probability measures $\mathbb{P}$ and $\mathbb{Q}$ such that $\mathbb{Q} \sim \mathbb{P}$, we have $D_{\mathrm{KL}}(\mathbb{Q}, \mathbb{Q}) = D_{\mathrm{KL}}(\mathbb{P}, \mathbb{P}) = 0$ and $D_{\mathrm{KL}}(\mathbb{Q}, \mathbb{P}) \geq 0$.

Proof. Let $\mathbb{P}$ and $\mathbb{Q}$ be two probability measures such that $\mathbb{Q} \sim \mathbb{P}$. We have

$$D_{\mathrm{KL}}(\mathbb{Q}, \mathbb{Q}) = \mathbb{E}_{\mathbb{Q}} \left[\log \frac{\mathrm{d}\mathbb{Q}}{\mathrm{d}\mathbb{Q}} \right] = \mathbb{E}_{\mathbb{Q}} \left[\log(1) \right] = 0,$$

and anlogously for $\mathbb{P}$. Next, we demonstrate that $D_{\mathrm{KL}}(\mathbb{Q},\mathbb{P}) \geq 0$. By applying Jensen's inequality (Lemma C.2) to the negative logarithm we obtain

$$D_{\mathrm{KL}}(\mathbb{Q},\mathbb{P}) = \mathbb{E}_{\mathbb{Q}}\left[-\log\frac{\mathrm{d}\mathbb{P}}{\mathrm{d}\mathbb{Q}}\right] \geq -\log\mathbb{E}_{\mathbb{Q}}\left[\frac{\mathrm{d}\mathbb{P}}{\mathrm{d}\mathbb{Q}}\right] = -\log\Big(\int\frac{\mathrm{d}\mathbb{P}}{\mathrm{d}\mathbb{Q}}\mathrm{d}\mathbb{Q}\Big) = -\log(1) = 0,$$

where the penultimate equality follows from the fact that $\mathbb{Q}$ and $\mathbb{P}$ are equivalent measures, and thus their Radon-Nikodym derivative integrates to 1. □

Theorem 7.7. Let the equivalent martingale measures $\mathbb{Q}^*$ and $\mathbb{Q}^*_H$ be as in (79) and (80), respectively. Next, define the probability measure $\mathbb{P}_H \sim \mathbb{P}$ by

$$\frac{\mathrm{d}\mathbb{P}_H}{\mathrm{d}\mathbb{P}} = \frac{\exp(\lambda\widetilde{H})}{\mathbb{E}_{\mathbb{P}}\left[\exp(\lambda\widetilde{H})\right]}. \tag{82}$$

Then the following assertions hold true:

i) The measure $\mathbb{Q}^*$ minimzes the entropy $D_{\mathrm{KL}}(\mathbb{Q},\mathbb{P})$ among all equivalent martingale measures $\mathbb{Q} \in \mathcal{P}$. Moreover, $\mathbb{Q}^*$ does not depend on λ and the maximal expected utility $U^* := \mathbb{E}\left[u(\tilde{G}_T(\phi^*))\right]$ is given by

$$U^* = 1 - \exp\Big(-D_{\mathrm{KL}}(\mathbb{Q}^*,\mathbb{P})\Big). \tag{83}$$

ii) The measure $\mathbb{Q}^*_H$ minimizes the relative entropy

$$D_{\mathrm{KL}}(\mathbb{Q},\mathbb{P}_H) = D_{\mathrm{KL}}(\mathbb{Q},\mathbb{P}) - \lambda\mathbb{E}_{\mathbb{Q}}\left[\tilde{H}\right] + \log(\mathbb{E}\left[\exp(\lambda\tilde{H})\right])$$

among all equivalent martingale measures $\mathbb{Q} \in \mathcal{P}$. Moreover, the maximal expected utility $U^*_H := \mathbb{E}\left[u(\tilde{\pi}^*_H + \tilde{G}_T(\phi^*) - \tilde{H})\right]$ is given by

$$U^*_H = 1 - \exp\Big(-D_{\mathrm{KL}}(\mathbb{Q}^*_H,\mathbb{P}) + \lambda\mathbb{E}_{\mathbb{Q}^*_H}\left[\tilde{H}\right] - \lambda\tilde{\pi}^*_H\Big).$$

Proof. We begin with the proof of (i): For $x > 0$ set $g(x) := x\log(x)$ and note that since $g''(x) = \frac{1}{x} > 0$ for all $x > 0$ it is strictly convex on $(0,\infty)$, i.e.,

$$g(y) - g(x) \geq g'(x)(y-x), \quad \text{for all } x, y > 0.$$

For the risk-neutral measure $\mathbb{Q}^*$ we thus obtain

$$g'(\frac{\mathrm{d}\mathbb{Q}^*}{\mathrm{d}\mathbb{P}}) = 1 + \log(\frac{\mathrm{d}\mathbb{Q}^*}{\mathrm{d}\mathbb{P}}) = 1 + \log(\exp(-\lambda\widetilde{G}_T(\phi^*))) - \log(\mathbb{E}\left[\exp(-\lambda\widetilde{G}_T(\phi^*))\right]).$$

Therefore, we obtain for all equivalent martingale measures $\mathbb{Q} \in \mathcal{P}$ the following:

$$
\begin{aligned}
D_{\mathrm{KL}}(\mathbb{Q}, \mathbb{P}) &= \mathbb{E}_{\mathbb{P}}\left[\frac{\mathrm{d}\mathbb{Q}}{\mathrm{d}\mathbb{P}} \log(\frac{\mathrm{d}\mathbb{Q}}{\mathrm{d}\mathbb{P}})\right] && (84)\\
&= \mathbb{E}_{\mathbb{P}}\left[g(\frac{\mathrm{d}\mathbb{Q}}{\mathrm{d}\mathbb{P}})\right] && (85)\\
&\geq \mathbb{E}_{\mathbb{P}}\left[g(\frac{\mathrm{d}\mathbb{Q}^*}{\mathrm{d}\mathbb{P}})\right] + \mathbb{E}_{\mathbb{P}}\left[g'(\frac{\mathrm{d}\mathbb{Q}^*}{\mathrm{d}\mathbb{P}})\Big(\frac{\mathrm{d}\mathbb{Q}}{\mathrm{d}\mathbb{P}} - \frac{\mathrm{d}\mathbb{Q}^*}{\mathrm{d}\mathbb{P}}\Big)\right]\\
&\geq D_{\mathrm{KL}}(\mathbb{Q}^*, \mathbb{P}) + \mathbb{E}_{\mathbb{Q}}\left[\lambda \widetilde{G}_T(\phi^*)\right] - \mathbb{E}_{\mathbb{Q}^*}\left[\lambda \widetilde{G}_T(\phi^*)\right]\\
&= D_{\mathrm{KL}}(\mathbb{Q}^*, \mathbb{P}). && (86)
\end{aligned}
$$

This proves that $\mathbb{Q}^*$ indeed minimzes the relative entropy over all equivalent martingale measures $\mathbb{Q} \in \mathcal{P}$. Moreover, due to the strict convexity of g the equality in (84) holds if and only if $\mathbb{Q} = \mathbb{Q}^*$, which implies that measure $\mathbb{Q}^*$ is the unique equivalent martingale measure minimizing the relative entropy.
To prove the formula for the expected utility U^* in (83), we compute:

$$
\begin{aligned}
D_{\mathrm{KL}}(\mathbb{Q}^*, \mathbb{P}) &= \mathbb{E}_{\mathbb{Q}^*}\left[\log\Big(\frac{\mathrm{d}\mathbb{Q}^*}{\mathrm{d}\mathbb{P}}\Big)\right]\\
&= \log(\frac{1}{\mathbb{E}_{\mathbb{P}}\left[\exp(-\lambda \widetilde{G}_T(\phi^*))\right]}) - \lambda \mathbb{E}_{\mathbb{Q}^*}\left[\widetilde{G}_T(\phi^*)\right]\\
&= -\log\Big(\mathbb{E}_{\mathbb{P}}\left[\exp(-\lambda \widetilde{G}_T(\phi^*))\right]\Big),
\end{aligned}
$$

which gives

$$
\begin{aligned}
\mathbb{E}_{\mathbb{P}}\left[u(\widetilde{G}_T(\phi^*))\right] &= 1 - \mathbb{E}_{\mathbb{P}}\left[\exp\Big(-\lambda \widetilde{G}_T(\phi^*)\Big)\right]\\
&= 1 - \exp\Big(-D_{\mathrm{KL}}(\mathbb{Q}^*, \mathbb{P})\Big).
\end{aligned}
$$

Next, we show (ii). First, note that

$$
\begin{aligned}
D_{\mathrm{KL}}(\mathbb{Q}, \mathbb{P}_H) &= \mathbb{E}_{\mathbb{Q}}\left[\log(\frac{\mathrm{d}\mathbb{Q}}{\mathrm{d}\mathbb{P}_H})\right]\\
&= \mathbb{E}_{\mathbb{Q}}\left[\log(\frac{\mathrm{d}\mathbb{Q}}{\mathrm{d}\mathbb{P}}) - \log(\frac{\mathrm{d}\mathbb{P}_H}{\mathrm{d}\mathbb{P}})\right]\\
&= D_{\mathrm{KL}}(\mathbb{Q}, \mathbb{P}) - \lambda \mathbb{E}_{\mathbb{Q}}\left[\tilde{H}\right] + \log\Big(\mathbb{E}_{\mathbb{P}}\left[\exp(-\lambda \tilde{H})\right]\Big).
\end{aligned}
$$

Note that

$$
\begin{aligned}
\frac{\mathrm{d}\mathbb{Q}_H^*}{\mathrm{d}\mathbb{P}_H} &= \frac{\mathrm{d}\mathbb{Q}_H^*}{\mathrm{d}\mathbb{P}} \frac{\mathrm{d}\mathbb{P}}{\mathrm{d}\mathbb{P}_H} = \frac{\mathrm{e}^{-\lambda(\widetilde{W}_T(\varphi_H^*) - \tilde{H})}}{\mathbb{E}_{\mathbb{P}}\left[\mathrm{e}^{-\lambda(\widetilde{W}_T(\varphi_H^*) - \tilde{H})}\right]} \frac{\mathbb{E}_{\mathbb{P}}\left[\exp(\lambda \tilde{H})\right]}{\exp(\lambda \tilde{H})}\\
&= K u'(\widetilde{W}_T(\varphi_H^*)),
\end{aligned}
$$

for some positive constant K. It follows now from part (i), that $\mathbb{Q}_H^*$ minimzes the relative entropy with respect to $\mathbb{P}_H$ among all equivalent martingale measures $\mathbb{Q} \in \mathcal{P}$ and the formula for the expected utility follows from that as well, since

$$
\begin{aligned}
\mathbb{E}_{\mathbb{P}}\left[u(\widetilde{W}_T(\varphi_H^*)\right] &= 1 - \mathbb{E}_{\mathbb{P}}\left[\exp(-\lambda(\widetilde{W}_T(\varphi_H^*) - \tilde{H}))\right]\\
&= 1 - \mathrm{e}^{-\lambda \pi_H^*} \exp(-D_{\mathrm{KL}}(\mathbb{Q}^*, \mathbb{P}_H)).
\end{aligned}
$$

This proves the assertions. □

Recall that given the two utility optimal strategies φ^* and φ^*_H we can define the **utility based hedging strategy** $\varphi \coloneqq \varphi^*_H - \varphi^*$, as this is the strategy that adjusts the banks utility optimal strategy according to the contingent claim trade. We have the following:

Theorem 7.8. In the situation of Theorem 7.7, the following holds true:

i) The utility-based hedging strategy φ and the measure $\mathbb{Q}^*_H$ depend on the premium π^*_H only through the bank account part.

ii) The (discounted) utility indifferent price $\tilde{\pi}_u(H)$ is given by

$$\begin{aligned}\tilde{\pi}_u(H) &= \frac{1}{\lambda}\log\Big(\mathbb{E}\left[\mathrm{e}^{\lambda\tilde{H}}\right]\Big) + D_{\mathrm{KL}}(\mathbb{Q}^*,\mathbb{P}) - D_{\mathrm{KL}}(\mathbb{Q}^*_H,\mathbb{P}_H)\Big) \\ &= \frac{1}{\lambda}\log\Big(\mathbb{E}_{\mathbb{Q}^*}\Big[\exp\Big(-\lambda(\tilde{G}_T(\varphi) - \tilde{H})\Big)\Big]\Big).\end{aligned} \tag{87}$$

Proof. (i) We already observed that the inital investment enters the shifted utility only as multiplicative constant. The same holds for the premium π^*_H. Therefore, it does not affect the optimaility of a trading strategy and the equivalent martingale measure $\mathbb{Q}^*_H$.
(ii) By part (i) and (ii) of Theorem 7.7, the utility indifferent price π^*_u satisfies the following equation:

$$1 - \exp(-D_{\mathrm{KL}}(\mathbb{Q}^*,\mathbb{P})) = 1 - \mathbb{E}\left[\mathrm{e}^{\lambda\tilde{H}}\right]\exp(-D_{\mathrm{KL}}(\mathbb{Q}^*_H,\mathbb{P}_H))\exp(-\lambda\tilde{\pi}^*_H).$$

Therefore, $\mathrm{e}^{\lambda\tilde{\pi}^*_H} = \mathbb{E}\left[\exp(\lambda\widetilde{H})\right]\exp(-D_{\mathrm{KL}}(\mathbb{Q}_H,\mathbb{P}_H) + D_{\mathrm{KL}}(\mathbb{Q}^*,\mathbb{P}))$, which is equivalent to

$$\begin{aligned}\tilde{\pi}^*_H &= \frac{1}{\lambda}\Big(\log(\mathbb{E}_{\mathbb{P}}\left[\mathrm{e}^{\lambda\widetilde{H}}\right]) - D_{\mathrm{KL}}(\mathbb{Q}^*_H,\mathbb{P}_H) + D_{\mathrm{KL}}(\mathbb{Q}^*,\mathbb{P})\Big) \\ &= \frac{1}{\lambda}\left(\lambda\mathbb{E}_{\mathbb{Q}^*_H}\left[\widetilde{H}\right] - D_{\mathrm{KL}}(\mathbb{Q}^*_H,\mathbb{P}_H) + D_{\mathrm{KL}}(\mathbb{Q}^*,\mathbb{P})\right).\end{aligned}$$

But $\mathbb{E}_{\mathbb{P}}\left[\exp\Big(-\lambda\varphi^*\bullet\widetilde{S}_T\Big)\right] = \mathbb{E}_{\mathbb{P}}\left[\exp\Big(-\lambda(\pi^*_u(H) + \varphi^*_H\bullet\widetilde{S}_T - \tilde{H})\Big)\right]$ due to the fact that $U^* = U^*_H(\pi^*_u(H))$ and hence $\exp\Big(\mathrm{e}^{-\lambda\pi^*_u(H)}\Big) = \mathbb{E}_{\mathbb{Q}^*}\left[\exp\Big(-\lambda(-\tilde{H} + \varphi\bullet\widetilde{S}_T)\Big)\right]$, by $\varphi^*_H = \varphi^* + \varphi$ and the fact that

$$\frac{\mathrm{d}\mathbb{Q}^*}{\mathrm{d}\mathbb{P}} = \frac{\exp\Big(-\lambda\varphi^*\bullet\widetilde{S}_T\Big)}{\mathbb{E}_{\mathbb{P}}\left[\exp\Big(-\lambda\varphi^*\bullet\widetilde{S}_T\Big)\right]}.$$

This concludes the proof. □

The quantities above all depent on the risk-aversion parameter $\lambda > 0$. To reflect that, we henceforth denote the respective quantities as u_λ, $U^*(\lambda)$, $U^*_H(\pi,\lambda)$, $\varphi^*(\lambda)$, $\varphi^*_H(\lambda)$, φ_λ and lastly $\pi^*_{u_\lambda}(H)$, the utility indifferent price. In the following theorem we treat the extreme cases that $\lambda \to 0$ and $\lambda \to \infty$. We need the following lemma:

Lemma 7.9. With the notation from above we have:

i) $\varphi^*(\lambda) = \lambda^{-1}\varphi^*(1)$ and $U^*(\lambda) = U^*(1)$ for any $\lambda > 0$,

ii) $\lambda \mapsto \pi^*_{u_\lambda)}(H)$ is increasing on $\mathbb{R}^+$.

Proof. (i) Note that for some normalizing constant $c > 0$, we have:

$$\frac{\mathrm{d}\mathbb{Q}^*}{\mathrm{d}\mathbb{P}} = cu_1'\Big(\tilde{W}_T(\varphi^*(1))\Big) = c\exp(-\varphi(1)^\intercal \widetilde{S}_T) = \frac{c}{\lambda}u_\lambda'(\widetilde{W}_T(\frac{1}{\lambda}\varphi^*(1))).$$

Since $\mathbb{Q}^* \in \mathcal{P}$, we conclude that $\lambda^{-1}\varphi^*(1)$ is an optimal portfolio for u_λ by Theorem 7.1. The second assertion follows from part (i) in Theorem 7.7. For (ii), by Theorem 7.8

$$\begin{aligned}\pi^*_{u_\lambda}(H) &= \mathbb{E}_{\mathbb{Q}^*_H}\left[\tilde{H}\right] + \frac{1}{\lambda}\Big(D_{\mathrm{KL}}(\mathbb{Q}^*,\mathbb{P}) - D_{\mathrm{KL}}(\mathbb{Q}^*_H,\mathbb{P})\Big)\\ &= \inf_{\mathbb{Q}\in\mathcal{P}}\left(\mathbb{E}_{\mathbb{Q}}\left[\tilde{H}\right] + \frac{1}{\lambda}\Big(D_{\mathrm{KL}}(\mathbb{Q}^*,\mathbb{P}) - D_{\mathrm{KL}}(\mathbb{Q},\mathbb{P})\Big)\right).\end{aligned}$$

As $\mathbb{Q}^*$ is minimal, we find that $D_{\mathrm{KL}}(\mathbb{Q}^*,\mathbb{P}) - D_{\mathrm{KL}}(\mathbb{Q},\mathbb{P}) \le 0$ for any $\mathbb{Q} \in \mathcal{P}$ and therefore $\pi^*_{u_\lambda}(H)$ is increasing in λ. □

For large risk aversion we have the following result:

Theorem 7.10. Let the utility indifferent price be denoted by $\pi^*_{u_\lambda}(H)$. Then the following holds true:

i) $\pi^*_{u_\lambda}(H) \to \sup \Pi(H)$ as $\lambda \to \infty$, i.e., the utility indifferent price converges to the upper price bound of arbitrage-free prices.

ii) The utility based hedging strategy $\varphi(\lambda)$ is a so called **asymptotic superhedge** for H, i.e., it satisfies

$$\liminf_{\lambda\to\infty} \widetilde{W}_T(\varphi(\lambda)) \ge H \quad \mathbb{P}\text{-a.s.}, \tag{88}$$

where $\widetilde{W}_T(\varphi(\lambda)) = \pi^*_{u_\lambda}(H) + (\varphi(\lambda) \bullet \widetilde{S})_T$ is the terminal value of the utility based hedging strategy.

iii) If the cheapest superhedge ψ is unique, then $W_t(\varphi(\lambda)) \to W_t(\psi)$ for all $t \in \mathbf{T}$ as $\lambda \to \infty$.

Proof. We provide the proof for the case of a finite probability space $(\Omega, \mathcal{F}, \mathbb{P})$. Observe that the inequality $D_{\mathrm{KL}}(\mathbb{Q}^*,\mathbb{P}) - D_{\mathrm{KL}}(\mathbb{Q}^*_H,\mathbb{P}) \le 0$ and (87) imply $\hat{\pi}(\lambda) \le \mathbb{E}_{\mathbb{Q}^*_H}\left[\widetilde{H}\right] \le \hat{\pi}_{u_\lambda}$ for $\lambda > 0$. We address the first two assertions by contradiction.
Suppose $\sup_{\lambda>0}\hat{\pi}(\lambda) \le \hat{\pi}_{u_\lambda} - 2\epsilon$ for some $\epsilon > 0$. For a fixed λ, there exists an $\omega_\lambda \in \Omega$ such that

$$W_T(\varphi(\lambda))(\omega_\lambda) := \hat{\pi}(\lambda) + (\varphi(\lambda) \bullet \hat{S})_T(\omega_\lambda) \le \widetilde{H}(\omega_\lambda) - \epsilon.$$

Otherwise, if this were not the case, $\hat{\pi}(\lambda) + \epsilon < \hat{\pi}_{u_\lambda}$ would represent the discounted initial value of a superhedge. We then conclude

$$\begin{aligned}U(\lambda) &= U_H(\pi(\lambda), \lambda)\\ &= \mathbb{E}_{\mathbb{P}}\left[u_\lambda(\widetilde{W}_T(\varphi(\lambda)) - \widetilde{H})\right]\\ &= 1 - \mathbb{E}_{\mathbb{P}}\left[\exp\Big(-\lambda(\hat{\pi}(\lambda) + (\varphi(\lambda)\bullet\hat{S})_T - \widetilde{H})\Big)\right]\\ &\le 1 - \mathbb{P}(\{\omega_\lambda\})\exp(\lambda\epsilon),\end{aligned}$$

where the right-hand side converges to $-\infty$ as $\lambda \to \infty$, since the finiteness of Ω implies $\min_{\omega\in\Omega} \mathbb{P}(\{\omega\}) > 0$. However, this contradicts $U(\lambda) \geq 1 - e^0 = 0$.
Next, suppose that (88) does not hold. Then there exist an $\omega \in \Omega$ and an $\epsilon > 0$ such that for some sequence of arbitrarily large λ,

$$\widetilde{W}_T(\varphi(\lambda))(\omega) \leq \widetilde{H}(\omega) - \epsilon,$$

which leads to the contradiction outlined in the proof of the first statement.
To prove the last assertion, consider a sequence $(\lambda_k)_{k\in\mathbb{N}} \subseteq \mathbb{R}_+$ with $\lambda_k \to \infty$. We aim to show that $\lim_{k\to\infty} \widetilde{W}_T(\varphi(\lambda_k)) = \widetilde{W}_t(\varphi_u)$ for $t = 0, 1, \ldots, T$. Starting with the terminal trading time $t = T$,
by (88), the sequence $(\widetilde{W}_T(\varphi(\lambda_k))(\omega))_{k\in\mathbb{N}}$ is bounded from below for any $\omega \in \Omega$. Since $\mathbb{E}_{\mathbb{Q}^*}\left[\widetilde{W}_T(\varphi(\lambda_k))\right] = \widetilde{W}_0(\varphi(\lambda_k)) = \hat{\pi}(\lambda_k) \leq \hat{\pi}_u$ and $\min_{\omega\in\Omega} \mathbb{P}(\{\omega\}) > 0$, these sequences are also bounded from above. Consequently, there exists a convergent subsequence of $(\widetilde{W}_T(\varphi(\lambda_k)))_{k\in\mathbb{N}}$. Any such convergent subsequence, denoted again by $(\widetilde{W}_T(\varphi(\lambda_k)))_{k\in\mathbb{N}}$, converges to $\widetilde{W}_T(\varphi)$ for some self-financing strategy φ. By (88), we see that $\widetilde{W}_T(\varphi) = \lim_{k\to\infty} \widetilde{W}_T(\varphi(\lambda_k)) \geq \widetilde{H}$, indicating that φ is a superhedge. Moreover, the finiteness of Ω leads to

$$\widetilde{W}_0(\varphi) = \mathbb{E}_{\mathbb{Q}^*}\left[\widetilde{W}_T(\varphi)\right] = \lim_{k\to\infty} \mathbb{E}_{\mathbb{Q}^*}\left[\widetilde{W}_T(\varphi(\lambda_k))\right] = \lim_{k\to\infty} \hat{\pi}(\lambda_k) \leq \hat{\pi}_u.$$

Therefore, φ is the cheapest superhedge, and $\widetilde{W}_T(\varphi) = \widetilde{W}_T(\varphi_u)$.
For $t < T$, pointwise convergence as $k \to \infty$ and the finiteness of Ω imply

$$\widetilde{W}_t(\varphi(\lambda_k)) = \mathbb{E}_{\mathbb{Q}^*}\left[\widetilde{W}_T(\varphi(\lambda_k))|\mathcal{F}_t\right] \to \mathbb{E}_{\mathbb{Q}^*}\left[\widetilde{W}_T(\varphi_u)|\mathcal{F}_t\right] = \widetilde{W}_t(\varphi_u).$$

This completes the proof of the last statement. □

Note that for small λ, i.e., as $\lambda \to 0$, we can use the following heuristic argument: We assume that the utility-based hedge converges for $\lambda \to 0$ as

$$\varphi(\lambda) = \eta + \mathcal{O}(\lambda),$$

where η is some limiting hedge and $\mathcal{O}(\lambda)$ represents a strategy such that $\mathcal{O}(\lambda)/\lambda$ is bounded by a constant independent of λ. For $\varphi^*_H(\lambda)$, we have

$$\varphi^*_H(\lambda) = \varphi^*(\lambda) + \varphi(\lambda) = \frac{1}{\lambda}\varphi^*(\lambda) + \eta + \mathcal{O}(\lambda).$$

Regarding the utility indifferent price, we use a linear approximation

$$\pi^*_{u_\lambda}(H) = \pi^*_{u_0}(H) + \lambda\delta + \mathcal{O}(\lambda^2),$$

for $\pi^*_{u_0}(H), \delta \in \mathbb{R}$. Considering the expected utility maximization problem for strategies of the above form and applying the Taylor approximation of the exponential function, we get:

$$\begin{aligned}
\mathbb{E}_{\mathbb{P}}\left[u_\lambda(\tilde{\pi} + \widetilde{G}_T(\varphi) - \widetilde{H})\right] &= 1 - \mathbb{E}_{\mathbb{P}}\left[\exp\Big(-\varphi^*(1) \bullet \widetilde{S}_T - \lambda(\tilde{\pi} + (\eta + \mathcal{O}(\lambda)) \bullet \widetilde{S}_T - \widetilde{H})\Big)\right] \\
&= 1 - K\mathbb{E}_{\mathbb{Q}^*}\left[\exp\Big(-\lambda(\tilde{\pi} + (\eta + \mathcal{O}(\lambda)) \bullet \widetilde{S}_T - \widetilde{H})\Big)\right] \\
&= 1 - K - \lambda K\mathbb{E}_{\mathbb{Q}^*}\left[(\tilde{\pi} + (\eta + \mathcal{O}(\lambda)) \bullet \widetilde{S}_T - \widetilde{H})\right] \\
&\quad + \frac{\lambda^2}{2}K\mathbb{E}_{\mathbb{Q}^*}\left[\Big(\tilde{\pi} + \eta \bullet \widetilde{S}_T - \widetilde{H}\Big)^2\right] + \mathcal{O}(\lambda^3) \\
&= 1 - K + \lambda K\Big(\tilde{\pi} - \mathbb{E}_{\mathbb{Q}^*}\left[\widetilde{H}\right]\Big) \\
&\quad - \frac{\lambda^2}{2}K\mathbb{E}_{\mathbb{Q}^*}\left[\Big(\tilde{\pi} + \eta \bullet \widetilde{S}_T - \widetilde{H}\Big)^2\right] + \mathcal{O}(\lambda^3),
\end{aligned}$$

with $K = \mathbb{E}_{\mathbb{P}}\left[\exp\left(-\varphi^*(1)\bullet\widetilde{S}_T\right)\right]$, assuming that $\widetilde{S}$ is a $\mathbb{Q}^*$-martingale. Ignoring the $\mathcal{O}(\lambda^3)$ term, we focus on minimizing

$$\mathbb{E}_{\mathbb{Q}^*}\left[\left(\tilde{\pi}^* + \eta\bullet\widetilde{S}_T - \widetilde{H}\right)^2\right], \tag{89}$$

over all strategies η. This problem is known as variance optimal hedging, which we investigate in the subsequent section. Assume the minimum value is

$$(\tilde{\pi} - \widetilde{W}_0)^2 + \varepsilon^2,$$

then

$$U_H^*(\lambda) = 1 - K + \lambda K(\tilde{\pi} - \widetilde{W}_0) - \frac{\lambda^2}{2}K\left((\tilde{\pi} - \widetilde{W}_0)^2 + \varepsilon^2\right) + \mathcal{O}(\lambda^3),$$

provided that the linear expansion of the optimal hedge is valid. The maximal utility investment for the plain investment problem is

$$U^*(\lambda) = U^*(1) = 1 - \mathbb{E}\left[\exp(-\varphi^*(1)\bullet\widetilde{S}_T)\right] = 1 - K.$$

The utility indifferent price $\tilde{\pi}^*_{u_\lambda}(H)$ satisfies:

$$K(\tilde{\pi}(0) - \widetilde{W}_0) - \frac{\lambda}{2}K\left(2\delta + (\tilde{\pi}(0) - \widetilde{W}_0)^2 + \varepsilon^2\right) + \mathcal{O}(\lambda^2),$$

implying that $\tilde{\pi}^*_{u_0}(H) = \widetilde{W}_0 = \mathbb{E}_{\mathbb{Q}^*}\left[\widetilde{H}\right]$, and $\delta = \varepsilon^2/2$. Thus, as $\lambda \to 0$, the discounted utility indifferent price converges to the expectation of the discounted payoff relative to the minimal entropy martingale measure $\mathbb{Q}^*$. In the first order, there is a risk premium linearly dependent on λ and contingent on the approximation quality of the claim by a self-financing strategy in a variance-optimal sense. In particular, the utility-based hedge in the leading term is equivalent to the variance-optimal hedge of the claim relative to the measure $\mathbb{Q}^*$.

These results establish a connection between different approaches to valuing and hedging derivatives. For exponential utility, a continuum of prices and hedging strategies emerges. Superreplication, which typically requires a high option premium, is at one end of this spectrum, while at the other end, we have formulations closely related to variance-optimal hedging, which will be explored in the following section.

8 Optimal Control Problems in Portfolio Theory

In this section, we delve into stochastic optimal control problems within the realm of portfolio theory. Our focus is on elucidating three pivotal consumption-investment problems (Problems 1-3 below) and exploring two primary methods for addressing these challenges:

i) *The Dynamic Programming Method*: This approach simplifies complex dynamic portfolio optimization problems into more manageable static optimization problems. It then tackles the original problem through a recursive solution strategy. The proofs of key results are provided as exercises or can be understood in detail from Appendix A.5 of Peter Spreij's lecture notes.

ii) *The Martingale Method*: This method bifurcates the dynamic portfolio optimization problem into two parts: a static optimization problem and a hedging problem. We then use a Lagrange Multiplier technique so solve the static optimization problem.

8.1 Optimal Consumption-Investment Problems

Consider a financial market in finite-discrete time, represented as $(\Omega, \mathcal{F}, \mathbb{P}, \mathbb{F}, X)$ with trading times $\mathbf{T} = \{0, 1, 2, \ldots, T\}$ and an asset price vector $X = (S^{(0)}, S^{(1)}, \ldots, S^{(d)})$ for $d \in \mathbb{N}$. The riskless asset $S^{(0)}$ evolves as $S^{(0)}_{t+1} = S^{(0)}_t(1 + r_t)$, where $(r_t)_{t\in\mathbf{T}}$ is the locally riskless rate. Risky assets follow $S^{(i)}_{t+1} = S^{(i)}_t(1 + R^{(i)}_t)$ for return sequences $(R^{(i)}_t)_{t\in\mathbf{T}}$.
Define the utility function u, denoting investor preferences. The first objective is to maximize the expected utility from terminal wealth under budget restriction:

Problem 1. Expected Utility from Terminal Wealth

$$\begin{cases} \max_{(W_0,\phi)} \mathbb{E}_{\mathbb{P}}\left[u(W_0 + G_T(\phi))\right] \\ \text{s.t.} \quad W_0 = w \in \mathbb{R} \quad \text{and} \quad \phi \text{ is predictable.} \end{cases}$$

Addressing the maximization of expected utility from consumption and terminal wealth requires defining a consumption-investment plan:

Definition 8.1. A **consumption-investment plan** is a pair $(K_t, \varphi_t)_{t\in\mathbf{T}}$ consisting of a nonnegative adapted process $(K_t)_{t\in\mathbf{T}}$ and a trading strategy $(\varphi_t)_{t\in\mathbf{T}}$. It is called **self-financing** if

$$W_t(\varphi) = K_t + \varphi_{t+1}^\intercal X_t, \quad t = 0, \ldots, T-1, \tag{90}$$

and **admissible** if $K_T \leq W_T$ $\mathbb{P}$-almost surely.

The self-financing condition (90) is expressed as:

$$\Delta W_t(\varphi) = \varphi_t^\intercal \Delta S_t - K_{t-1}, \quad t = 1, \ldots, T,$$

$$\text{or in discounted terms:} \Delta \widetilde{W}_t(\varphi) = \varphi_t^\intercal \Delta \widetilde{S}_t - \gamma_{t-1}, \quad t = 1, \ldots, T,$$

where $\gamma_{t-1} \triangleq \frac{K_{t-1}}{S^{(0)}_t}$ is the discounted consumption. The wealth evolution is then given by

$$\widetilde{W}_t(\varphi) = \widetilde{W}_0 + \sum_{s=1}^{t} \left(\varphi_t^\intercal \Delta \widetilde{S}_t - \gamma_{s-1}\right).$$

We explore two additional optimization problems related to expected utility from consumption:

Problem 2. Expected utility from consumption

$$\begin{cases} \max_{(K,\varphi)} \mathbb{E}_{\mathbb{P}} \left[\sum_{t=1}^{T} \beta_t u(K_t)\right], \\ \text{s.t. } W_0 = w \text{ and } \beta_t \in (0,1], \\ \text{and } (K,\varphi) \text{ is an admissible consumption-investment plan.} \end{cases} \tag{91}$$

Here, β_t are discount factors and u is a utility function, assumed constant across all times t, but can be time-dependent.

Consider two utility functions u_c and u_p for the next problem:

Problem 3. Expected utility from consumption and terminal wealth

$$\begin{cases} \max_{(K,\varphi)} \mathbb{E}_{\mathbb{P}} \left[\sum_{t=1}^{T} \beta_t u_c(K_t) + \beta_T u_p(W_T(\varphi) - K_T)\right], \\ \text{s.t. } \varphi \text{ is a self-financing strategy}, \beta_t \in (0,1], \\ \text{and } K \text{ is a non-negative admissible consumption plan.} \end{cases} \tag{92}$$

8.2 Stochastic Optimal Control

We frame **Problems 1–3** from Section 8.1 in the context of *stochastic optimal control problems* and address them using *dynamic programming* and the *martingale method* in the forthcoming sections. Note here that the maximization problems on both sides of the utility based hedging problem (63) can be considered as **Problem 1** (with slight adjustments).

Let $(V_t)_{t\in\mathbf{T}}$ be a given process (think of the wealth process of some strategy) that depend on the choice of a $\mathbb{R}^m$-valued **control sequence** $(\alpha_t)_{t\in\mathbf{T}}$ (think of a trading strategy), that is adapted to a filtration $(\mathcal{F}_t)_{t\in\mathbf{T}}$. More precisely, we assume that $(V_t)_{t\in\mathbf{T}}$ is of the following form:

$$V_{t+1} = G_t(V_t, \alpha_t, \varepsilon_t), \quad t = 0, \ldots, T-1, \tag{93}$$

with initial value V_0 and given random quantities $\varepsilon_0, \varepsilon_1, \ldots, \varepsilon_{T-1}$ (think of random returns of asset prices) that are assumed to be *iid* for now. If the ε_t for $t = 0, 1, \ldots, T-1$ are k-dimensional, then the functions G_t in (93) are defined on some appropriate subsets of $\mathbb{R}^d \times \mathbb{R}^m \times \mathbb{R}^k$, and are also assumed to be jointly measurable in their arguments.

As filtration we take the family of σ-algebras $(\mathcal{F}_t)_{t\in\mathbf{T}}$ given by

$$\mathcal{F}_t = \sigma(V_0, \varepsilon_0, \ldots, \varepsilon_{t-1}),$$

and note that the processes $(V_t)_{t\in\mathbf{T}}$ and $(\alpha_t)_{t\in\mathbf{T}}$ are both $(\mathcal{F}_t)_{t\in\mathbf{T}}$-adapted.

We further assume that the m-dimensional random variables α_t are supposed to be of the form

$$\alpha_t = \alpha(t, V_0, \ldots, V_t), \tag{94}$$

for certain measurable functions $\alpha(t,\cdot)\colon (\mathbb{R}^d)^{t+1} \to \mathbb{R}^m$, for which we use the notation $\alpha(t,\cdot) \in \mathcal{B}((\mathbb{R}^d)^{t+1}, \mathbb{R}^m)$. If $\alpha_t = \alpha(t, V_t)$ for all $t \in \mathbf{T}$ then we say that $(\alpha_t)_{t\in\mathbf{T}}$ is a **Markov control** and in this case we see that the process $(V_t)_{t\in\mathbf{T}}$ given by the recursion (93) is Markovian.

Lemma 8.2. The process X is Markov with respect to the filtration specified above, if α_t depends on X_t only, i.e. $\alpha_t = \alpha(t, X_t)$ for all $t \in \mathbf{T}$ for some measurable functions $\alpha(t,\cdot)\colon \mathbb{R}^d \to \mathbb{R}^m$.

Proof. Left as an Exercise. $\square$

8.2.1 The Control Problem

Next, we define a general stochastic optimal control problem that includes the investment-consumption control problems as given by **Problems 1-3** as special cases:

Problem 4. Stochastic optimal control problem Let $g_0, \dots, g_{T-1}\colon \mathbb{R}^d \times \mathbb{R}^m \to \mathbb{R}$, as well as $g_T\colon \mathbb{R}^d \to \mathbb{R}$ be measurable functions. We refer to $g_0, \dots, g_T$ as the **reward functions** (think of them as utility functions). The **optimal control problem** is to maximize over all controls $\alpha = (\alpha_0, \dots, \alpha_{T-1})$ the additive over time criterion given by the expectation

$$J(\alpha) \coloneqq \mathbb{E}_{\mathbb{P}}\left[\sum_{t=0}^{T-1} g_t(V_t, \alpha_t) + g_T(V_T)\right], \tag{95}$$

with each α_t for $t = 0, 1, \dots, T-1$ as in (94).

Remark 8.3. Under assumption (94), **Problem 4** is equivalent to finding measurable functions $\alpha(t,\cdot)$ maximizing $J(\alpha(\cdot,\cdot))$ in (95), where for $\alpha = (\alpha_0, \dots, \alpha_{T-1})$ we also write $J(\alpha)$! Usually, the functions $\alpha_0, \dots, \alpha_{T-1}$ have to satisfy certain admissibility constraints. These will be clear in the appropriate context and not always explicitly mentioned. For instance, it is tacitly assumed that all random quantities involved are such that the expectations exist.

Definition 8.4. A sequence of admissible functions $\alpha^* = (\alpha_0^*, \dots, \alpha_{T-1}^*)$ is called an **optimal control** for (95) if $J(\alpha^*) = \sup_\alpha J(\alpha)$ holds, where the supremum is taken over all admissible control sequences $\alpha = (\alpha_0, \dots, \alpha_{T-1})$ with $\alpha(t,\cdot) \in \mathcal{B}((\mathbb{R}^d)^{t+1}, \mathbb{R}^m)$ for $t = 0, \dots, T-1$.

Definition 8.5. Let $(\alpha_t)_{t\in\mathbf{T}}$ be a sequence of functions and let the process $(V_t^\alpha)_{t\in\mathbf{T}}$ be defined by (93) and (94), where the notation expresses the dependence of $(V_t)_{t\in\mathbf{T}}$ on $(\alpha_t)_{t\in\mathbf{T}}$. Note that $V_0^\alpha = V_0$. Set $J_T(\alpha) = g_T(V_T^\alpha)$, and for $t = 0, 1, \dots, T-1$:

$$J_t(\alpha) \coloneqq \mathbb{E}_{\mathbb{P}}\left[\sum_{s=t}^{T-1} g_s(V_s^\alpha, \alpha_s) + g_T(V_T^\alpha) \mid \mathcal{F}_t\right].$$

The so defined adapted process $(J_t(\alpha))_{t\in\mathbf{T}}$ is called the **expected future reward process** of the optimal control problem (95).

Note, that $\mathbb{E}_{\mathbb{P}}[J_0(\alpha)] = J(\alpha)$ and for $t = 0, 1, \dots, T-2$ the expected future reward process satisfies:

$$\begin{aligned} J_t(\alpha) &= g_t(V_t^\alpha, \alpha_t) + \mathbb{E}\left[\sum_{s=t+1}^{T-1} g_s(V_s^\alpha, \alpha_s) + g_T(V_T^\alpha) \mid \mathcal{F}_t\right] \\ &= g_t(V_t^\alpha, \alpha_t) + \mathbb{E}\left[\mathbb{E}\left[\sum_{s=t+1}^{T-1} g_s(V_s^\alpha, \alpha_s) + g_T(V_T^\alpha) \mid \mathcal{F}_{t+1}\right] \mid \mathcal{F}_t\right] \\ &= g_t(V_t^\alpha, \alpha_t) + \mathbb{E}\left[J_{t+1}(\alpha) \mid \mathcal{F}_t\right]. \end{aligned}$$

8.3 Dynamic Programming

The idea to obtain the control maximizing (95) is based on the *dynamic programming principle*, which states that if a process is optimal over the entire sequence of periods $t = 0, 1, \dots, T$, then is has to be also optimal over each single period, see Proposition 8.8 below. In the previous section's context, this means that in order to determine the optimal control sequence $\alpha_0, \dots, \alpha_T$ we aim to sequentially optimize over the individual controls α_t for $t = 0, 1, \dots, T$. In the Theorem 8.6 below, we show that under the asssumptions of Markov controls, i.e. if $\alpha_t = \alpha(t, V_t)$ for some measurable functions $u(t, \cdot)$ for all $t \in \mathbf{T}$, this idea yields an optimal control. Let us denote the class of Markov control sequences by $\mathcal{M}$ and define for certain given measurable functions $v_0, \dots, v_T \colon \mathbb{R}^d \to \mathbb{R}$ the following sequence of functions:

$$\hat{v}_{t+1}(x, y) = \mathbb{E}_{\mathbb{P}}\left[v_{t+1}(G_t(x, y, \varepsilon_t))\right], \quad t = 0, \dots, T-1. \tag{96}$$

Note that the functions $\hat{v}_{t+1} \colon \mathbb{R}^d \times \mathbb{R}^m \to \mathbb{R}$ are measurable in x and y. The main theorem of this section is the following:

Theorem 8.6. Define recursively the functions $(v_t)_{t \in \mathbf{T}}$ by

$$v_T(x) = g_T(x), \tag{97}$$

$$v_t(x) = \sup_{y \in \mathbb{R}^m} \{g_t(x, y) + \hat{v}_{t+1}(x, y)\}, \quad t = 0, \dots, T-1. \tag{98}$$

Assume that the v_t are measurable functions. Then the following hold true:

i) For any control sequence $\alpha \in \mathcal{M}$ of functions $(\alpha(t, \cdot))_{t \in \mathbf{T}}$ one has $v_t(V_t^\alpha) \geq J_t(\alpha)$ $\mathbb{P}$-almost surely and $\mathbb{E}_{\mathbb{P}}[v_0(V_0)] \geq J(\alpha)$.

ii) Let $\alpha^* \in \mathcal{M}$. Then α^* is optimal iff the supremum in (98) is attained for $y = \alpha_t^*(x)$. If this happens, then $v_t(V_t^{\alpha^*}) = J_t(\alpha^*)$ and

$$\sup_\alpha J(\alpha) = J(\alpha^*) = \mathbb{E}_{\mathbb{P}}[v_0(X_0)].$$

Remark 8.7. Theorem 8.6 depends on $\hat{v}_{t+1}(x, y) = \mathbb{E}\left[v_{t+1}(G_t(x, y, \varepsilon_t)) \mid \mathcal{F}_t\right]$, which is due to the independence of V_0 and the ε_t. This assumption might be dropped.

Proposition 8.8 (Optimality Principle)**.** Let $\alpha^* = (\alpha_0^*, \dots, \alpha_{T-1}^*)$ be the optimal control of (95). Then the sequence $(\alpha_t^*, \dots, \alpha_{T-1}^*)$ is optimal for the corresponding optimization problem over the time set $\{t, \dots, T\}$, when starting in $V_t = V_t^{\alpha^*}$. In this case, the optimal value is equal to $\mathbb{E}_{\mathbb{P}}\left[v_t(V_t^{\alpha^*})\right]$.

Proof. Left as an Exercise. □

Algorithm: *Dynamic Programming*
Suppose that the suprema in Equation (98) are attained for all $t \in \mathbf{T}$. Then define

$$\begin{aligned} v_T(x) &= g_T(x) \\ \alpha_{T-1}^*(x) &= \arg\sup\{g_{T-1}(x,y) + \hat{v}_T(x,y)\}, \end{aligned}$$

and by backward recursion for $t \in \{0, \dots, T-1\}$,

$$\begin{aligned} v_t(x) &= \sup\{g_t(x,y) + \hat{v}_{t+1}(x,y)\} \\ &= g_t(x, \alpha_t^*(x)) + \hat{v}_{t+1}(x, \alpha_t^*(x)) \\ \alpha_{t-1}^*(x) &= \arg\sup\{g_{t-1}(x,y) + \hat{v}_t(x,y)\}. \end{aligned}$$

This yields the sequence of functions $v_T, \alpha_{T-1}^*, v_{T-1}, \alpha_{T-2}^*, \dots, v_1, \alpha_0^*, v_0$ where the α_t^* constitute the optimal sequence α^* and $\mathbb{E}_{\mathbb{P}}[v_0(V_0)] = J(\alpha^*)$. The functions v_t are called the **optimal value functions**.

From here on, we drop the assumption that the $\varepsilon_1, \dots, \varepsilon_{T-1}$ are independent. One can still define random functions $\hat{v}_{t+1}$ as above, but now we alter the definition in (96) to

$$\hat{v}_{t+1}(x,y) \coloneqq \mathbb{E}\left[v_{t+1}(G_t(x,y,\varepsilon_t)) \mid \mathcal{F}_t\right]. \tag{99}$$

Then the $\hat{v}_{t+1}(x,y)$ are in general not deterministic anymore, but become $\mathcal{F}_t$-measurable random variables. However, many of the above results continue to hold, but we have to reprove some of the results (see also Appendix of Peter Spreij's lecture notes).
The following proposition uses the concept of essential supremum. We will consider $\mathcal{B}(\mathbb{R}^d) \times \mathcal{F}_t$-measurable functions v_t, meaning that the mapping $(x,\omega) \mapsto v_t(x,\omega)$ is $\mathcal{B}(\mathbb{R}^d) \times \mathcal{F}_t$-measurable. As usual, dependence on ω is often suppressed and then we write $v_t(x)$ for the random variable $\omega \mapsto v_t(x,\omega)$.

Proposition 8.9. Suppose that $\mathcal{B}(\mathbb{R}^d) \times \mathcal{F}_t$-measurable functions v_t (for $t = 0, \dots, T$) $\mathbb{P}$-almost surely satisfy

$$\begin{aligned} v_T(x) &= g_T(x), \\ v_t(x) &= \text{ess sup}\{g_t(x,y) + \mathbb{E}\left[v_{t+1}(G_t(x,y,\varepsilon_t)) | \mathcal{F}_t\right] : y \in \mathcal{F}_t\}, \quad t \le T-1. \end{aligned}$$

Then, for any admissible control α it holds that $v_t(V_t^\alpha) \ge J_t(\alpha)$ for $t \in \mathbf{T}$ and the optimum is given by $\mathcal{F}_t$ measurable random variables $\alpha_t^*(x)$ for which the supremum is attained.

8.3.1 Maximizing expected utility from terminal wealth using dynamic programming

We align **Problem 1** with the general stochastic optimal control problem as in (95). Recall from Section 8.1 the aim to maximize expected utility from terminal wealth using utility function u:

$$\begin{cases} \max_{(W_0,\phi)} \mathbb{E}_{\mathbb{P}}\left[u(W_0 + G_T(\phi))\right], \\ \text{s.t. } W_0 = w \text{ and } \phi \text{ is predictable.} \end{cases} \qquad \text{(Problem 1)}$$

In the context of (95), we set $g_t = 0$ for $t = 0, 1, \ldots, T-1$ and $g_T = u$, with V_t as the wealth at time t for a self-financing strategy $\varphi = (w, \phi)$.
We consider $\mathbb{R}^d$-valued predictable processes $(\phi_t)_{t\in\mathbf{T}}$ satisfying $\phi_0 = 0$ and $W_T((w,\phi)) \geq 0$.
For relative portfolio holdings, define control variables $\alpha_t = (\alpha_t^{(0)}, \alpha_t^{(1)}, \ldots, \alpha_t^{(d)})$ as

$$\alpha_t^{(i)} = \frac{\phi_{t+1}^{(i)} S_t^{(i)}}{V_t} \quad \text{for } i = 0, 1, \ldots, d, \tag{100}$$

and note that $\alpha_t^{(0)} = 1 - \sum_{i=1}^d \alpha_t^{(i)}$ with $\alpha_t^{(i)} \in (0,1)$.
The self-financing property leads to

$$V_t = \varphi_t^{(0)} S_t^{(0)} + \sum_{i=1}^d \varphi_t^{(i)} S_t^{(i)} = \varphi_{t+1}^{(0)} S_t^{(0)} + \sum_{i=1}^d \varphi_{t+1}^{(i)} S_t^{(i)}.$$

Given the predictable nature of ϕ and the evolution of $S_{t+1}^{(i)} = S_t^{(i)}(1 + R_t^{(i)})$, we derive

$$\begin{aligned} V_{t+1} &= V_t + \varphi_{t+1}^{(0)} \Delta S_t^{(0)} + \sum_{i=1}^d \varphi_{t+1}^{(i)} \Delta S_t^{(i)} \\ &= V_t + \varphi_{t+1}^{(0)} S_t^{(0)} r_t + \sum_{i=1}^d \varphi_{t+1}^{(i)} S_t^{(i)} R_t^{(i)} \\ &= V_t + V_t\Big(\alpha_t^{(0)} r_t + \sum_{i=1}^d \alpha_t^{(i)} R_t^{(i)}\Big). \end{aligned} \tag{101}$$

Comparing (93) with (101), we find that $V_{t+1} = G_t(V_t, \alpha_t, \xi_t)$, with $\xi_t = R_t$, V_t and α_t as defined above and such that

$$G_t(V_t, \alpha_t, \xi_t) = V_t\Big(\alpha_t^{(0)}(1 + r_t) + \sum_{i=1}^d \alpha_t^{(i)}(1 + \xi_t^{(i)})\Big). \tag{102}$$

For maximizing expected utility from terminal wealth only, we have

$$g_t(V_t, \alpha_t) = 0 \text{ for } t < T, \text{ and } g_T(V_T, \alpha_T) = u(V_T),$$

leading to the dynamic programming algorithm:

$$\begin{cases} v_T(v) &= g_T(v) = u(v), \\ v_t(v) &= \max_{\alpha_t} \mathbb{E}_{\mathbb{P}}\left[v_{t+1}(G_t(V_t, \alpha_t, R_t)) \mid V_t = v\right], \quad t = 0, 1, \ldots, T-1, \end{cases} \tag{103}$$

with no maximization over α_T. The optimal control sequence u^*, if existent, is derived accordingly. Indeed, let us consider as example the problem in the binomial model of Section 2.5.

Example 8.10 (Maximizing expected utility out of terminal wealth in the Binomial model)**.** Recall the Dynamic Programming algorithm (103), namely

$$\begin{cases} v_T(v) = u(v) \text{ and, for } t < T, \\ v_t(v) = \max_{\alpha_t} \mathbb{E}_{\mathbb{P}} \left[v_{t+1}(G_t(V_t, \alpha_t, \xi_t)) \mid V_t = v \right] \end{cases}$$

where the dynamics $G_t(V_t, \alpha_t, \xi_t)$ were specified as above in (102) namely

$$G_t(V_t, \alpha_t, \xi_t) = V_t\Big(\alpha_t^{(0)}(1 + r_t) + \sum_{i=1}^{d} \alpha_t^{(i)}(1 + \xi_t^{(i)})\Big).$$

The only particular aspect for the binomial model here is that $d = 1$ and, letting $\alpha_t = \alpha_t^{(1)} = \frac{\phi_{t+1}^{(1)} S_t}{V_t}$, one has $\alpha_t^{(0)} = (1 - \alpha_t)$. With prohibition of short selling, i.e., requiring $\phi_t^{(1)} > 0$, we obtain $\alpha_t \in (0, 1)$. If $r_t = 0$, we may then write

$$G_t(V_t, \alpha_t, \xi_t) = V_t\Big(1 + \alpha_t R_t\Big).$$

An investment strategy is here given by the two-dimensional vector $\alpha_t = (\alpha_t^{(0)}, \alpha_t^{(1)})$ where $\alpha_t^{(0)}$ and $\alpha_t^{(1)}$ denote the number of units of the non-risky and the risky assets respectively held in the portfolio in period t. Recall also that the process α_t is supposed to be predictable, i.e. α_t is taken to be $\mathcal{F}_{t-1}$-measurable. We now describe the backwards recursion, from period $T-1$ to 0, to determine the optimal investment strategy $(\alpha_t^*)_{t=0,\ldots,T-1}$.

Take period $T - 1$: we have to decide for $\alpha_{T-1}^{(0)}, \alpha_{T-1}^{(1)}$. Le $S_{T-1} = S_0 U^n D^{T-n-1}$. Recalling that we had assumed $r = 0$, we have to impose that for all outcomes, i.e. independently on whether prices go up or down,

$$\alpha_T^{(1)} S_T + \alpha_T^{(0)} = V_T = V^* \tag{104}$$

This implies that the following system of equations in $\alpha_T^{(1)}, \alpha_T^{(0)}$ has to be satisfied:

$$\alpha_T^{(1)} S_{T-1} U + \alpha_T^{(0)} = v \left(\frac{p}{q}\right)^{n+1} \left(\frac{1-p}{1-q}\right)^{T-n-1} \quad \text{(prices go up)}$$

$$\alpha_T^{(1)} S_{T-1} D + \alpha_T^{(0)} = v \left(\frac{p}{q}\right)^{n} \left(\frac{1-p}{1-q}\right)^{T-n} \quad \text{(prices go down)},$$

where $p = \mathbb{P}(R_t = U)$ and $q = \frac{1+r-D}{U-D}$, from which we obtain

$$\alpha_T^{(1)} S_{T-1}(U - D) = v(1+r)^T \left(\left(\frac{p}{q}\right)^{n} \left(\frac{1-p}{1-q}\right)^{T-n-1} - \left(\frac{p}{q}\right)^{n+1} \left(\frac{1-p}{1-q}\right)^{T-n} \right).$$

Consequently, we have

$$\alpha_T^{(1)} = \frac{v(1+r)^T}{S_{T-1}(U-D)} \left(\frac{pq}{(p-q)(1-q)}\right)^{n} \left(\frac{1-p}{1-q}\right)^{T-n-1},$$

and therefore

$$\alpha_{T-1}^{(1)} = \frac{v(pq)^{n/T}(1-p)^{T-n-1/T}}{(p-q)S_0U^n(U-D)(1-q)^{T-n-1}}$$

$$\alpha_{T-1}^{(0)} = v\left(\frac{q}{p}\right)^n\left(\frac{1-p}{1-q}\right)^{T-n-1}\frac{\big(U(1-p)q-D(1-q)p\big)}{(U-D)q(1-q)}$$

Recalling the self-financing condition, the fact that $q = \frac{1+r-D}{U-D}$, and putting

$$C(p,q) := \frac{U-D}{p-q} + (Dq-Dp) + (Dpq-Upq),$$

for the optimal value induced in period $T-1$ we thus obtain

$$V_{T-1}^* = \alpha_{T-1}^{(1)}S_{T-1} + \alpha_{T-1}^{(0)} = v\left(\frac{pq}{(1-q)}\right)^n\left(\frac{1-p}{1-q}\right)^{T-n-1}C(p,q). \tag{105}$$

Notice now that V_{T-1}^* has the same structure as V_T^* so that the calculations for the following period $T-2$ proceed in exactly the same way as for $T-1$ and so forth, which allows to straightforwardly complete the backwards recursion.
In the generic period $t \leq T$, with $n \leq t$ up-movement up to time t, the condition (104) becomes

$$\alpha_t^{(1)}S_t + \alpha_t^{(0)} = V_t^* \tag{106}$$

and one obtains

$$\alpha_t^{(1)} = \frac{v(pq)^{n/t}(1-p)^{t-n-1/t}}{(p-q)S_0U^n(U-D)(1-q)^{t-n-1}}$$

$$\alpha_t^{(0)} = v\left(\frac{q}{p}\right)^n\left(\frac{1-p}{1-q}\right)^{t-n-1}\frac{\big(U(1-p)q-D(1-q)p\big)}{(U-D)q(1-q)}$$

with the optimal wealth

$$V_t^* = v\left(\frac{p}{q}\right)^n\left(\frac{1-p}{1-q}\right)^{t-n}. \tag{107}$$

Notice also that the ratio of the wealth invested in the risky asset is

$$\pi_t = \frac{\alpha_{t+1}^{(1)}}{\alpha_t^{(1)}S_t} = \frac{1}{(U-D)q(1-q)(p-q)},$$

and it is independent of t and S_t. It follows that the optimal strategy requires investing in the risky asset, in each period and in every state, the same fraction of wealth.

8.3.2 Maximizing expected utility from consumption using dynamic programming

We align **Problem 2** with the general stochastic optimal control problem as follows:

$$\begin{cases} \max_{(K,\varphi)} \mathbb{E}_{\mathbb{P}} \left[\sum_{t=1}^{T} \beta_t u(K_t)\right] \\ \text{s.t. } W_0 = w \text{ and } \alpha_t \in (0,1], \\ \text{and } (K,\varphi) \text{ is an admissible consumption-investment.} \end{cases} \tag{Problem 2}$$

Incorporating a consumption process into the optimal control setting, for a self-financing consumption-investment plan (K, α), we have:

$$\begin{aligned} V_{t+1} &= V_t + \phi_{t+1}^{(0)} \Delta S_t^{(0)} + \sum_{i=1}^{d} \phi_{t+1}^{(i)} \Delta S_t^{(i)} - K_t \\ &= V_t \left(\alpha_t^{(0)} r_t + \sum_{i=1}^{d} \alpha_t^{(i)} R_t^{(i)} \right) - K_t \\ &= V_t \left((1+r_t) + \sum_{i=1}^{d} \alpha_t^{(i)} \left(R_t^{(i)} - (1+r_t) \right) \right) - K_t(1+r_t), \end{aligned}$$

where $\alpha^{(0)}, \ldots, \alpha^{(d)}$ denote the relative portfolio holdings $\alpha^{(i)} = \frac{\phi_{t+1}^{(i)} S_t^i}{V_t}$ as before. Therefore, we identify:

$$V_{t+1} = G_t(V_t, \alpha_t, \xi_t),$$

with

$$G_t(V_t, \alpha_t, \xi_t) = V_t \left((1+r_t) + \sum_{i=1}^{d} \alpha_t^{(i)} \left(\xi_{t+1}^{(i)} - (1+r_t) \right) \right) - K_t(1+r_t).$$

Here, the control sequence is $\alpha_t = (\alpha_t^{(1)}, \ldots, \alpha_t^{(d)}, K_t)$, controlling both relative portfolio holdings and consumption at each time $t = 0, 1, \ldots, T-1$. The utility is solely on consumption:

$$g_t(V_t, \alpha_t) = u(K_t), \quad t = 0, \ldots, T,$$

with the wealth V_t entering only through the constraint $K_T \leq V_T$.
The dynamic programming algorithm is thus defined as:

$$\begin{cases} v_T(v) = \begin{cases} u(v) & \text{if } K_T = v, \\ -\infty & \text{if } K_T > v, \end{cases} \\ v_t(v) = \max_{\alpha_t} \Big(u(K_t) + \beta \mathbb{E}_{\mathbb{P}} \left[v_{t+1}(G_t(V_t, \alpha_t, \xi_t)) \mid V_t = v \right] \Big), \end{cases}$$

where $u(K) = -\infty$ for $K < 0$ ensures nonnegativity of K, and we set:

$$J_t \triangleq \max_{K_t, \ldots, K_T \ (K_T \leq V_T)} \mathbb{E}_{\mathbb{P}} \left[\sum_{s=t}^{T} \beta_{s-t} u(K_s) \mid V_t = v \right],$$

with the condition $K_T \leq V_T$ enforced by setting $u(K_T) = -\infty$ for $K_T > V_T$.

8.3.3 Maximizing expected utility from consumption & terminal wealth using dynamic programming

In the previous case of expected utility only from consumption, all of V_T was consumed at the terminal time $t = T$. Here we generalize this problem by leaving $V_T - K_T$ for future investment and maximizing expected utility from consumption of K_t for $t = 0, \dots, T$, as well as from the terminal wealth $V_T - K_T$.
Letting $u_c(\cdot)$ and $u_p(\cdot)$ denote the utility functions for consumption and terminal wealth respectively, then the problem is given by:

$$\begin{cases} \max_{(K,\varphi)} \mathbb{E}_{\mathbb{P}} \left[\sum_{t=1}^{T} \beta_t u_c(K_t) + \beta_T u_p(W_T(\varphi) - K_T) \right], \\ \text{s.t. } \varphi \text{ is a self-financing strategy,} \\ \text{and } K \text{ is a non-negative admissible consumption plan.} \end{cases} \qquad \text{(Problem 3)}$$

For this problem, we can now adapt the previous solution method: we identify

$$V_{t+1} = G_t(V_t, \alpha_t, \xi_t)$$

with $\xi_t = R_t$ and

$$G_t(V_t, \alpha_t, \xi_t) = V_t\Big((1 + r_t) + \sum_{i=1}^{d} \alpha_t^{(i)} \Big(R_t^{(i)} - (1 + r_t)\Big)\Big) - K_t(1 + r_t),$$

and with the control sequence α given by

$$\alpha_t = (\alpha_t^{(1)}, \dots, \alpha_t^{(d)}, K_t)$$

where $\alpha_t^{(i)} = \frac{\phi_{t+1}^{(i)} S_t^{(i)}}{V_t}$ as before. In the present case, however, we have

$$g_t(V_t, \alpha_t) = u_c(K_t), \quad t = 0, \dots, T - 1,$$

and terminal utility

$$g_T(V_T, \alpha_T) = u_p(V_T - K_T)$$

The dynamic programming algorithm becomes

$$\begin{aligned} v_T(x) &= \max_{0 \le K \le x} [u_c(K) + u_p(x - K)] \\ v_t(x) &= \max_{\alpha_t} \Big(u_c(K_t) + \beta_t \mathbb{E}_{\mathbb{P}} [v_{t+1}(G_t(V_t, \alpha_t, \xi_t)) \mid V_t = v] \Big) \end{aligned}$$

and notice that the only difference with the pure consumption-investment problem in Section 8.3.2 is the form of $v_T(x)$ (in the previous case it is essentially given by $v_T(x) = u(x) = u_c(x)$).

8.4 The Martingale Method

The martingale method serves as a second approach, alongside the dynamic programming algorithm, for addressing stochastic optimal control problems in portfolio theory. This method hinges on the following key observation regarding hedging contingent claims in an arbitrage-free market: to hedge a contingent claim H, an initial investment $W_0 = \mathbb{E}_{\mathbb{Q}}\left[\widetilde{H}\right]$ is needed for some EMM $\mathbb{Q} \in \mathcal{P}$, to construct a trading strategy $\varphi = (W_0, \phi)$ such that $W_0 + G_t(\phi) = \mathbb{E}_{\mathbb{Q}}\left[\widetilde{H} \mid \mathcal{F}_t\right]$ for $t \in \mathbf{T}$, i.e. to finance a replicating strategy. Finding a replicating trading strategy ϕ is equivalent to representing the martingale $(\mathbb{E}_{\mathbb{Q}}\left[\widetilde{H} \mid \mathcal{F}_t\right])_{t\in\mathbf{T}}$ as a discrete stochastic integral.

In scenarios where we focus on consumption-investment **Problems 1-3** instead of hedging a contingent claim H, we can view the optimal reachable wealth from an initial investment as a claim to replicate using the controls. The martingale method thus involves the following three steps:

i) Deriving the set of all *reachable terminal portfolio wealth* W_T;

ii) Computing the optimal reachable wealth W_T^*;

iii) Determining a self-financing strategy ϕ^* such that $W_T(\phi^*) = W_T^*$.

We call this method the *martingale method*, which effectively decomposes the dynamic portfolio optimization problem into a static problem of computing the optimal reachable wealth and a hedging problem, where the optimal reachable wealth is considered as the contingent claim.

8.4.1 Reachable Portfolio Values

In the first step, we aim to determine the set of reachable portfolio values with a fixed initial investment v:

$$\mathcal{V} \triangleq \{W_T \colon W_T = W_T(\phi) \text{ for some self-financing and predictable } \phi \text{ such that } W_0 = v\}.$$

Here, $W_T(\phi)$ denotes the terminal portfolio wealth resulting from a self-financing trading strategy ϕ that starts with initial wealth $W_0 = v$.
Since the market is arbitrage-free, the Fundamental Theorem of Asset Pricing tells us that discounted asset prices are martingales under some equivalent martingale measure (EMM) $\mathbb{Q}$. By Doob's system Theorem 1.16, it follows that for any self-financing strategy ϕ, the discounted terminal wealth $\widetilde{W}_T = W_T(\phi) \cdot (S_T^{(0)})^{-1}$ satisfies:

$$v = W_0 = \mathbb{E}_{\mathbb{Q}}\left[\widetilde{W}_T\right],$$

for any $\mathbb{Q}$ in the set $\mathcal{P}$ of EMMs. Therefore, the initial wealth v equals the expected discounted terminal wealth under any EMM.

Complete Market Case In a complete market with a unique equivalent martingale measure $\mathbb{Q}$, the set of reachable terminal wealths is:

$$\mathcal{V} = \left\{W_T \colon \mathbb{E}_{\mathbb{Q}}\left[\widetilde{W}_T\right] = v\right\}, \tag{108}$$

where $\widetilde{W}_T = W_T \cdot (S_T^{(0)})^{-1}$.

Incomplete Market Case In an incomplete market with multiple EMMs, the set of reachable terminal wealth is characterized by:

$$\mathcal{V} = \left\{ W_T \colon \mathbb{E}_{\mathbb{Q}}\left[\widetilde{W_T}\right] = v \text{ for some } \mathbb{Q} \in \mathcal{P} \right\}, \tag{109}$$

where $\mathcal{P}$ denotes the convex set of all EMMs.

8.4.2 Optimal Reachable Portfolio Values Using the Lagrange Multiplier Technique

To determine an optimal W^* in the set of reachable portfolio values $\mathcal{V}$, we employ the Lagrange multiplier technique, both in complete and incomplete market settings.

The Complete Market Case In the complete market, $\mathcal{V}$ is given by equation (108). The optimal terminal wealth W^* maximizes the expected utility:

$$\max_{W_T \in \mathcal{V}} \mathbb{E}_{\mathbb{P}}\left[u(W_T)\right].$$

This can be reformulated using the Lagrangian approach. Let $L := \frac{d\mathbb{Q}}{d\mathbb{P}}$ denote the Radon-Nikodym derivative of $\mathbb{Q}$ with respect to $\mathbb{P}$, and let λ be the Lagrange multiplier associated with the budget constraint. The optimization problem becomes:

$$\max_{W_T} \mathbb{E}_{\mathbb{P}}\left[u(W_T) - \lambda (S_T^{(0)})^{-1} L W_T\right] + \lambda v.$$

First-Order Condition and Determining λ Taking the derivative with respect to W_T and setting it to zero, we get:

$$u'(W_T) = \lambda (S_T^{(0)})^{-1} L \quad \Longrightarrow \quad W_T = I\left(\lambda (S_T^{(0)})^{-1} L\right),$$

where $I(y) = (u')^{-1}(y)$ is the inverse of the marginal utility function. The Lagrange multiplier λ is found by satisfying the budget constraint:

$$v = \mathbb{E}_{\mathbb{Q}}\left[(S_T^{(0)})^{-1} W_T\right] = \mathbb{E}_{\mathbb{P}}\left[L (S_T^{(0)})^{-1} W_T\right] = \mathbb{E}_{\mathbb{P}}\left[L (S_T^{(0)})^{-1} I\left(\lambda (S_T^{(0)})^{-1} L\right)\right].$$

Define the function:

$$V(\lambda) := \mathbb{E}_{\mathbb{P}}\left[L (S_T^{(0)})^{-1} I\left(\lambda (S_T^{(0)})^{-1} L\right)\right].$$

If $V(\lambda)$ is invertible, we can solve for $\lambda = V^{-1}(v)$, and thus find the optimal terminal wealth:

$$W^* = I\left(V^{-1}(v) (S_T^{(0)})^{-1} L\right).$$

The optimal expected utility is then $\mathbb{E}_{\mathbb{P}}\left[u(W^*)\right] = \log v - \mathbb{E}_{\mathbb{P}}\left[\log\left(L (S_T^{(0)})^{-1}\right)\right]$.

Example 8.11. For the logarithmic utility function $u(W_T) = \log W_T$, we have $u'(W_T) = \frac{1}{W_T}$ and $I(y) = \frac{1}{y}$. Then:

$$W^* = \frac{1}{\lambda (S_T^{(0)})^{-1} L} = \frac{v}{L (S_T^{(0)})^{-1}},$$

since the budget constraint simplifies to $v = \frac{1}{\lambda}$ implying $\lambda = \frac{1}{v}$.

Example 8.12. Consider a binomial model with a Bernoulli process $B_T \sim \text{Bin}(T, p)$ representing the number of up movements. Using $u(W_T) = \log W_T$, we have:

$$W^* = v \left(\frac{p}{q}\right)^{B_T} \left(\frac{1-p}{1-q}\right)^{T-B_T},$$

where q is the risk-neutral probability of an up move, and $L = \left(\frac{q}{p}\right)^{B_T} \left(\frac{1-q}{1-p}\right)^{T-B_T}$. The optimal expected utility is:

$$\mathbb{E}_{\mathbb{P}}\left[u(W^*)\right] = \log v - T\left[p \log\left(\frac{q}{p}\right) + (1-p)\log\left(\frac{1-q}{1-p}\right)\right].$$

The Incomplete Market Case with Finitely Many Extremal EMMs In an incomplete market with a finite number J of extremal equivalent martingale measures $\mathbb{Q}_j$, we proceed similarly.
Let $L_j := \frac{\mathrm{d}\mathbb{Q}_j}{\mathrm{d}\mathbb{P}}$ and consider a vector of Lagrange multipliers $\lambda = (\lambda_1, \ldots, \lambda_J)$. The optimization problem becomes:

$$\max_{W_T} \mathbb{E}_{\mathbb{P}}\left[u(W_T) - \left(\sum_{j=1}^{J} \lambda_j L_j (S_T^{(0)})^{-1}\right) W_T\right] + \sum_{j=1}^{J} \lambda_j v.$$

First-Order Condition Differentiating with respect to W_T:

$$u'(W_T) = \sum_{j=1}^{J} \lambda_j L_j (S_T^{(0)})^{-1} \quad \Longrightarrow \quad W_T = I\left(\sum_{j=1}^{J} \lambda_j L_j (S_T^{(0)})^{-1}\right).$$

Budget Equations The multipliers λ_j are determined by the budget constraints, which require that W_T is attainable under each $\mathbb{Q}_j$:

$$v = \mathbb{E}_{\mathbb{Q}_j}\left[(S_T^{(0)})^{-1} W_T\right] = \mathbb{E}_{\mathbb{P}}\left[L_j (S_T^{(0)})^{-1} W_T\right] = \mathbb{E}_{\mathbb{P}}\left[L_j (S_T^{(0)})^{-1} I\left(\sum_{k=1}^{J} \lambda_k L_k (S_T^{(0)})^{-1}\right)\right], \quad \forall j = 1, \ldots, J.$$

This forms a system of J equations that can be solved for λ_j.

8.4.3 Maximizing Expected Utility from Consumption Using the Martingale Method

To maximize utility from consumption over multiple periods, we adapt the three-step martingale method used for maximizing utility from terminal wealth. The steps are modified as follows:

i) Determine the set of attainable consumption processes.

ii) Determine the optimal attainable consumption process.

iii) Determine an investment strategy that allows consuming according to the optimal process.

Step i): Determining the Set of Attainable Consumption Processes We begin by establishing how consumption affects the portfolio's evolution.

Lemma 8.13. Given an initial wealth $v \geq 0$, a consumption process $(K_t)_{t=0}^T$, and a self-financing trading strategy ϕ, the portfolio value W_t at time t is given by:

$$W_t = v + \tilde{G}_t - \sum_{s=0}^{t-1} \frac{K_s}{S_s^{(0)}}, \quad t = 1, \ldots, T. \tag{110}$$

Proof. Using the self-financing condition in the presence of consumption, the portfolio value evolves according to:

$$W_t = W_{t-1} + \sum_{i=1}^{N} \phi_t^i \left(S_t^i - S_{t-1}^i \right) - K_{t-1}.$$

Expressing in terms of discounted prices:

$$\tilde{W}_t = \tilde{W}_{t-1} + \sum_{i=1}^{N} \phi_t^i \left(\tilde{S}_t^i - \tilde{S}_{t-1}^i \right) - \frac{K_{t-1}}{S_{t-1}^{(0)}}.$$

By iterating this equation from $t = 1$ to $t = T$ and summing up, we arrive at equation (110). □

Definition 8.14. A consumption process $(K_t)_{t=0}^T$ is called **attainable** if there exists a self-financing trading strategy ϕ with (K, ϕ) admissible such that $K_T = V_T$.

Under an equivalent martingale measure $\mathbb{Q}$, the discounted gains process $\tilde{G}_t$ is a martingale starting at zero. At maturity T, the terminal portfolio value is:

$$W_T = v + \tilde{G}_T - \sum_{s=0}^{T-1} \frac{K_s}{S_s^{(0)}}.$$

Since $W_T = K_T$ we have:

$$v + \tilde{G}_T = \sum_{s=0}^{T} \frac{K_s}{S_s^{(0)}}.$$

Taking expectations under $\mathbb{Q}$ and noting that $\tilde{G}_T$ is a $\mathbb{Q}$-martingale with zero initial value, the budget constraint is:

$$v = \mathbb{E}_{\mathbb{Q}} \left[\sum_{t=0}^{T} \frac{K_t}{S_t^{(0)}} \right].$$

Proposition 8.15. In a complete financial market with unique martingale measure $\mathbb{Q}$, for a given initial wealth $v \geq 0$, a consumption process $(K_t)_{t=0}^T$ is attainable if and only if:

$$v = \mathbb{E}_{\mathbb{Q}} \left[\sum_{t=0}^{T} \frac{K_t}{S_t^{(0)}} \right].$$

Whereas in an incomplete market with multiple martingale measures $\mathbb{Q}_j$, $j = 1, \ldots, J$

the consumption process is attainable if and only if:

$$v = \mathbb{E}^{\mathbb{Q}_j}\left[\sum_{t=0}^{T} \frac{K_t}{S_t^{(0)}}\right], \quad \forall j = 1, \dots, J.$$

Step ii): Determining the Optimal Attainable Consumption Process We aim to maximize the expected utility from consumption:

$$\max_{K} \sum_{t=0}^{T} \mathbb{E}_{\mathbb{P}}\left[\beta_t u(K_t)\right],$$

subject to the budget constraints derived above.
In the incomplete market with multiple equivalent martingale measures, we face ambiguity in the budget constraints. To ensure attainability under all extremal martingale measures $\mathbb{Q}_j$, $j = 1, \dots, J$, we impose:

$$v = \mathbb{E}_{\mathbb{Q}_j}\left[\sum_{t=0}^{T} \frac{K_t}{S_t^{(0)}}\right], \quad \forall j = 1, \dots, J.$$

Optimization Problem:

$$\begin{cases} \max\limits_{K} \sum\limits_{t=0}^{T} \mathbb{E}_{\mathbb{P}}\left[\beta_t u(K_t)\right] \\ \text{s.t. } \sum\limits_{t=0}^{T} \mathbb{E}_{\mathbb{Q}_j}\left[\frac{K_t}{S_t^{(0)}}\right] = v, \quad \forall j = 1, \dots, J. \end{cases} \tag{111}$$

Remarks:

- The decision variables are the consumption amounts K_t at each time t.
- The investment strategy ϕ_t does not appear explicitly in the optimization problem.

To solve this problem, we express the constraints in terms of expectations under the real-world measure $\mathbb{P}$.
Let

- $L_j = \frac{d\mathbb{Q}_j}{d\mathbb{P}}$ be the Radon-Nikodym derivative of $\mathbb{Q}_j$ with respect to $\mathbb{P}$.
- $N_t^j = \frac{1}{S_t^{(0)}}\mathbb{E}_{\mathbb{P}}\left[L_j \mid \mathcal{F}_t\right]$.

Then

$$\sum_{t=0}^{T} \mathbb{E}_{\mathbb{Q}_j}\left[\frac{K_t}{S_t^{(0)}}\right] = \sum_{t=0}^{T} \mathbb{E}_{\mathbb{P}}\left[L_j \cdot \frac{K_t}{S_t^{(0)}}\right] = \sum_{t=0}^{T} \mathbb{E}_{\mathbb{P}}\left[K_t N_t^j\right].$$

Then, rewriting the optmization problem yields:

$$\begin{cases} \max\limits_{K} \sum\limits_{t=0}^{T} \mathbb{E}_{\mathbb{P}}\left[\beta_t u(K_t)\right] \\ \text{s.t. } \sum\limits_{t=0}^{T} \mathbb{E}_{\mathbb{P}}\left[K_t N_t^j\right] = v, \quad \forall j = 1, \dots, J. \end{cases} \tag{112}$$

Lagrangian Formulation We introduce Lagrange multipliers λ_j for the budget constraints:

$$\mathcal{L}(K, \lambda) = \sum_{t=0}^{T} \mathbb{E}_{\mathbb{P}}\left[\beta_t u(K_t)\right] - \sum_{j=1}^{J} \lambda_j \left(\sum_{t=0}^{T} \mathbb{E}_{\mathbb{P}}\left[K_t N_t^j\right] - v \right).$$

Simplifying:

$$\mathcal{L}(K, \lambda) = \sum_{t=0}^{T} \mathbb{E}_{\mathbb{P}}\left[\beta_t u(K_t) - \sum_{j=1}^{J} \lambda_j N_t^j K_t\right] + \left(\sum_{j=1}^{J} \lambda_j \right) v.$$

First-Order Condition For each t, we set the derivative of the Lagrangian with respect to K_t to zero:

$$\frac{\partial \mathcal{L}}{\partial K_t} = 0 \quad \Longrightarrow \quad \beta_t u'(K_t) = \sum_{j=1}^{J} \lambda_j N_t^j, \quad \text{a.s.}$$

Solving for K_t

$$K_t = I\left(\frac{\sum_{j=1}^{J} \lambda_j N_t^j}{\beta_t} \right),$$

where $I(y) = (u')^{-1}(y)$ is the inverse of the marginal utility function.

Determining the Lagrange Multipliers λ_j Substitute K_t back into the budget constraints:

$$\sum_{t=0}^{T} \mathbb{E}_{\mathbb{P}}\left[N_t^j \cdot I\left(\frac{\sum_{k=1}^{J} \lambda_k N_t^k}{\beta_t} \right) \right] = v, \quad \forall j = 1, \ldots, J.$$

This forms a system of J equations to solve for λ_j.

Optimal Value Function The optimal expected utility is:

$$J(v) = \sum_{t=0}^{T} \mathbb{E}_{\mathbb{P}}\left[\beta_t u\left(I\left(\frac{\sum_{j=1}^{J} \lambda_j N_t^j}{\beta_t} \right) \right) \right].$$

Step iii): Determining the Investment Strategy Finally, we need to find an investment strategy ϕ that finances the optimal consumption process $(K_t)_{t=0}^{T}$.

Approach

- Use the martingale representation theorem to express the discounted gains process $\tilde{G}_t$ in terms of the martingale $\mathbb{E}_{\mathbb{Q}}\left[\cdot \mid \mathcal{F}_t\right]$.
- Since the process $M_t = v + \tilde{G}_t - \sum_{s=0}^{t-1} \dfrac{K_s}{S_s^{(0)}}$ is a $\mathbb{Q}$-martingale, we can represent $\tilde{G}_t$ as a stochastic integral with respect to the discounted asset prices.
- Determine ϕ_t by solving the hedging problem for the contingent claim corresponding to the cumulative consumption.

The specific method to determine ϕ_t depends on the market model. In discrete-time models like the binomial model, we can use backward induction. In continuous-time models, we may apply Itô's lemma and stochastic calculus techniques.
By adapting the martingale method to the consumption setting, we decompose the problem of maximizing expected utility from consumption into:

1. **Static Optimization**: Finding the optimal consumption process by solving an optimization problem with budget constraints expressed in terms of expectations under equivalent martingale measures.

2. **Dynamic Hedging**: Determining an investment strategy that finances the optimal consumption process, leveraging the martingale representation of the gains process.

8.4.4 Maximizing expected utility from consumption & terminal wealth using the martingale method

The approach here differs from the previous case in that we no longer require $K_T = V_T$ (as per Definition 6 of attainability).
In this context, attainability implies that for an incomplete market with J extremal martingale measures:

$$v = \mathbb{E}_{\mathbb{Q}_j}\left[\frac{K_0}{S_0^{(0)}} + \cdots + \frac{K_{T-1}}{S_{T-1}^{(0)}} + \frac{V_T}{S_T^{(0)}}\right], \quad j = 1, \ldots, J.$$

To maximize expected utility from consumption and terminal wealth, the problem becomes:

$$\begin{aligned} &\max_{K,V} \mathbb{E}_{\mathbb{P}}\left[\sum_{t=0}^{T} \beta_t u_c(K_t) + \beta_T u_p(V_T - K_T)\right] \\ &\text{subject to} \quad \mathbb{E}_{\mathbb{P}}\left[\sum_{t=0}^{T-1} K_t N_t + V_T N_T\right] = v; \quad j = 1, \ldots, J. \end{aligned}$$

Applying the Lagrange multiplier technique, we need to maximize:

$$\max_{K,V} \mathbb{E}_{\mathbb{P}}\left[\sum_{t=0}^{T} \beta_t u_c(K_t) + \beta_T u_p(V_T - K_T) - \sum_{j=1}^{J} \lambda_j \left(\sum_{t=0}^{T-1} K_t N_t + V_T N_T\right)\right].$$

The necessary condition for K_t and V_T is:

$$\begin{aligned} \beta_t u_c'(K_t) &= \sum_{j=1}^{J} \lambda_j N_t, \quad t = 0, \ldots, T-1, \\ \beta_T u_c'(K_T) &= \beta_T u_p'(V_T - K_T), \\ \beta_T u_p'(V_T - K_T) &= \sum_{j=1}^{J} \lambda_j N_T. \end{aligned}$$

Using the inverses $I_c(\cdot)$ and $I_p(\cdot)$ of $u_c'(\cdot)$ and $u_p'(\cdot)$, respectively:

$$\begin{aligned} K_t &= I_c\left(\frac{\sum_{j=1}^{J} \lambda_j N_t}{\beta_t}\right), \quad t = 0, \ldots, T, \\ V_T &= I_p\left(\frac{\sum_{j=1}^{J} \lambda_j N_T}{\beta_T}\right) + I_c(\cdot). \end{aligned}$$

The Lagrange multipliers λ_j satisfy the budget equations:

$$\mathbb{E}_{\mathbb{P}}\left[\sum_{t=0}^{T} N_t^j I_c\left(\frac{\sum_{j=1}^{J}\lambda_j N_t^j}{\beta^t}\right) + N_T^j I_p\left(\frac{\sum_{j=1}^{J}\lambda_j N_T^j}{\beta^T}\right)\right] = v, \quad j = 1, \ldots, J.$$

The optimal value is then:

$$J(v) = \mathbb{E}_{\mathbb{P}}\left[\sum_{t=0}^{T} \beta^t u_c\left(I_c\left(\frac{\sum_{j=1}^{J}\lambda_j N_t^j}{\beta^t}\right)\right) + \beta^T u_p\left(I_p\left(\frac{\sum_{j=1}^{J}\lambda_j N_T^j}{\beta^T}\right)\right)\right].$$

As before, in the last step, determining the investment strategy of a self-financing portfolio that achieves a terminal value of V_T, paralleling the method used for expected utility from terminal wealth.

Appendix

A Stochastic Calculus in Discrete-Time

A.1 Fubini's Theorem

The following theorem is a version of the stochastic Fubini theorem:

Theorem A.1 (Stochastic Fubini)**.** For $i = 0, 1$ let $(E_i, \mathcal{E}_i)$ be measurable spaces and $U_i : (\Omega, \mathcal{F}) \to (E_i, \mathcal{E}_i)$ measurable. Let $\mathcal{F}_0 = \sigma(U_0)$ and U_1 be independent of U_0. Then

$$E[f(U_0, U_1)|\mathcal{F}_0](\omega) = E[f(U_0(\omega), U_1)] =: h(U_0(\omega))$$

for all non-negative measurable functions f on $E_0 \times E_1$.

Proof. The right hand side of the preceding equation is $\mathcal{F}_0$-measurable so by definition of the conditional expectation we only need to show that $E[Zf(U_0, U_1)] = E[Zh(U_0)]$ for all $\mathcal{F}_0$-measurable random variables Z. To this end, notice first that Z allows a representation of the form $Z = g(U_0)$. Thus, Fubini's theorem yields

$$\begin{aligned} E[Zf(U_0, U_1)] &= E[g(U_0)f(U_0, U_1)] \\ &= \int_{E_1} \int_{E_0} g(u_0)f(u_0, u1)\mu_0(du_0)\mu_1(du_1) \\ &= \int_{E_0} \int_{E_1} g(u_0)f(u_0, u_1)\mu_1(du_1)\mu_0(du_0) \\ &= \int_{E_0} \left(\int_{E_1} f(u_0, u_1)\mu_1(du_1) \right) \mu_0(du_0) \\ &= E[g(U_0)h(U_0)] \\ &= E[Zh(U_0)]. \end{aligned}$$

□

A.2 Martingales

Definition A.2. We call a stochastic process $(M_t)_{t \in \mathbf{T}}$ on a filtered probability space $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$ a $\mathbb{P}$-*martingale* if M is $\mathbb{F}$-adapted, satisfies $\mathbb{E}_{\mathbb{P}}[|M_t|] < \infty$ for all t, and if

$$M_s = \mathbb{E}_{\mathbb{P}}[M_t|\mathcal{F}_s] \quad \text{for} \quad 0 \le s \le t \le T. \tag{113}$$

A.3 Martingale transforms

Definition A.3. Let $(M_t)_{t \in \mathbf{T}}$ denote a martingale on a filtered probability space $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$ and let $(H_t)_{t \in \mathbf{T}}$ be an $\mathbb{F}$-predictable process, i.e., H_t is $\mathcal{F}_{t-1}$-measurable for all $t \in \mathbf{T}$ with $t > 0$, and H_0 is $\mathcal{F}_0$-measurable. The *martingale transform* $(H \bullet M)$ of M by H is defined as

$$(H \bullet M)_t := \sum_{s=1}^{t} H_s \Delta M_s, \quad \text{with } \Delta M_s := M_s - M_{s-1}.$$

A.4 Itô's formula

In the discrete-time setting, Itô's formula is a consequence of Taylor's theorem.

Theorem A.4 (Itô's formula in discrete-time)**.** Let $(M_t)_{t\in\mathbf{T}}$ be a martingale on a filtered probability space $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$ and let $f\colon \mathbb{R} \to \mathbb{R}$ be a twice continuously differentiable function. Then the process $(f(M_t))_{t\in\mathbf{T}}$ is given by

$$f(M_t) = f(M_0) + \sum_{s=1}^{t} f'(M_{s-1})\Delta M_s + \frac{1}{2}\sum_{s=1}^{t} f''(M_{s-1})(\Delta M_s)^2.$$

A.5 Girsanov's theorem

In the discrete-time setting, Girsanov's theorem can be stated as follows:

Theorem A.5 (Girsanov's theorem in discrete-time)**.** Let $(M_t)_{t\in\mathbf{T}}$ be a martingale on a filtered probability space $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$ and let $(H_t)_{t\in\mathbf{T}}$ be an $\mathbb{F}$-predictable process. Define the process $(Z_t)_{t\in\mathbf{T}}$ by

$$Z_t = \exp\left(\sum_{s=1}^{t} H_s \Delta M_s - \frac{1}{2}\sum_{s=1}^{t} H_s^2\right).$$

Suppose that $\mathbb{E}_{\mathbb{P}}[Z_T] = 1$. Then there exists a probability measure $\mathbb{Q}$ on $(\Omega, \mathcal{F}_T)$ equivalent to $\mathbb{P}$ such that

$$\frac{d\mathbb{Q}}{d\mathbb{P}} = Z_T.$$

Furthermore, the process $(\widetilde{M}_t)_{t\in\mathbf{T}}$ defined by

$$\widetilde{M}_t = M_t - \sum_{s=1}^{t} H_s,$$

is a martingale with respect to the measure $\mathbb{Q}$ and the filtration $\mathbb{F}$.

In other words, Girsanov's theorem in discrete-time states that under certain conditions, a change of measure can be achieved by applying a martingale transform, resulting in a new martingale under the new measure.

B Utility Theory

Utility Theory is a central theme in economics and financial decision-making. Its origins can be traced back to the attempts by early economists to understand and quantify human preferences, satisfaction, and the choices that individuals make under conditions of uncertainty.

In essence, utility represents a measure of relative satisfaction or preference that an individual derives from consuming goods, making investments, or generally, facing uncertain

outcomes. Instead of merely focusing on monetary value or profit, *Utility Theory* considers the inherent value or satisfaction one gains from an action or decision, making it especially crucial in the realm of finance, where risk and uncertainty are omnipresent.
In this section, we will delve deep into the foundational concepts of Utility Theory. We will explore its mathematical formulations, discuss its relevance in portfolio optimization, and understand how investors and financial institutions leverage utility functions to make informed decisions under uncertainty.

B.1 Preference Relations

In a market, agents have preferences for certain commodities over others, such as favoring apples over pears. These preferences are articulated through preference relations. Commodities also encompass risky assets or contingent claims with uncertain future pay-offs. While complete markets assign unique arbitrage-free prices to such claims based on the risk-neutral measure, incomplete markets present a range of arbitrage-free prices. Preference relations and their quantitative equivalent, utility functions, help select among these prices. Utility functions depict an agent's risk attitude and can yield different prices based on the risk quantification method used. They can also guide choices between portfolios with identical prices in complete markets. The detailed exploration of utility functions will be covered in Section 5.2, while this section focuses on preference relations.

Let $\mathcal{X}$ be a non-empty set representing commodities, securities, or more generally, possible choices an economic agent can make. A binary relation R on $\mathcal{X}$ can be represented as a subset of $\mathcal{X} \times \mathcal{X}$, where xRy means $(x, y) \in R$. In the following, we will denote binary relations by $\succ$, $\succeq$, $\prec$, and $\preceq$.

Definition B.1. A **strict preference relation** or preference order on $\mathcal{X}$ is a binary relation $\succ$ satisfying:

i) **Asymmetry**: If $x \succ y$, then $y \not\succ x$.

ii) **Negative Transitivity**: If $x \succ y$ and $z \in \mathcal{X}$, then either $x \succ z$ or $z \succ y$.

A **weak preference relation** on $\mathcal{X}$ is a binary relation $\succeq$ satisfying:

i) **Completeness**: For all $x, y \in \mathcal{X}$, either $x \succeq y$ or $y \succeq x$.

ii) **Transitivity**: If $x \succeq y$ and $y \succeq z$, then $x \succeq z$.

Proposition B.2. Strict and weak preference relations interrelate as follows:

i) If $\succ$ is a strict preference relation on $\mathcal{X}$, then defining $x \succeq y$ by the condition $y \not\succ x$ yields a weak preference relation on $\mathcal{X}$.

ii) Conversely, if $\succeq$ is a weak preference relation, then $x \succ y$ defined by $x \succeq y$ and $y \not\succeq x$ yields a strict preference relation.

Proof. This is left as an exercise. □

For strict and weak preference relations $\succ$ and $\succeq$, it is assumed they interrelate as in the previous proposition. An **indifference relation** $\sim$ is defined by $x \sim y$ if and only if $x \succeq y$ and $y \succeq x$. This indifference relation is an *equivalence relation* and we can also express $x \succ y$ as $x \succeq y \wedge x \not\sim y$.
For notational convenience, we sometimes use reverse preference relations. So, instead of $x \succ y$, we might write $y \prec x$, and likewise, $y \preceq x$ can represent $x \succeq y$.

B.2 Numerical representations

Abstract preference orders can be equated with numerical representations using the standard order $\leq$ on $\mathbb{R}$.

Definition B.3. A function $U : \mathcal{X} \to \mathbb{R}$ provides a **numerical representation** of a preference relation $\succ$ if $x \succ y$ is equivalent to $U(x) > U(y)$.

Alternatively, a numerical representation can be defined by:

$$x \succeq y \iff U(x) \geq U(y)$$

Note that any strictly increasing transformation of a numerical representation U results in another numerical representation, making them non-unique. When these numerical representations and preference relations have specific additional properties, they can be expressed as utility functions. This will be discussed further in Section 5.2.

Definition B.4. For a preference relation $\succ$ on $\mathcal{X}$, a subset $\mathcal{Z}$ of $\mathcal{X}$ is termed **order dense** in $\mathcal{X}$ if, for any $x, y \in \mathcal{X}$ with $x \succ y$, there exists some $z \in \mathcal{Z}$ such that $x \succeq z \succeq y$.

We have the following results:

Theorem B.5. A preference relation $\succ$ on $\mathcal{X}$ admits a numerical representation if and only if $\mathcal{X}$ contains a countable order dense subset.

Proof. Let $\mathcal{Z}$ be a countable order dense subset of $\mathcal{X}$. Choose a probability measure μ on $\mathcal{Z}$ with $\mu(z) = \mu(\{z\}) > 0$ for all $z \in \mathcal{Z}$. Then we put

$$U(x) := \sum_{z: x \succ z} \mu(z) - \sum_{z: z \succ x} \mu(z). \tag{114}$$

Notice that the existence of such a probability distribution is guaranteed by countability of $\mathcal{Z}$, that also makes $U(x)$ well-defined in terms of the given summations. By construction, we have $x \succ y$ iff $U(x) > U(y)$. To see this, we first compute for $x \succeq y$ the difference

$$U(x) - U(y) = \sum_{z: x \succ z \succeq y} \mu(z) + \sum_{z: x \succeq z \succ y} \mu(z).$$

If $x \succ y$, then there is $z_0 \in \mathcal{Z}$ such that $x \succeq z_0 \succeq y$. By negative transitivity, we also have $z_0 \succ y$ or $x \succ z_0$. Hence, we have $x \succ z_0 \succeq y$ or $x \succeq z_0 \succ y$, and we see that at least one of the two sums in the display is strictly positive, which yields $U(x) > U(y)$. If $x \succeq y$, the right-hand side of the displayed formula is still well-defined, has nonnegative terms (possibly zero) and hence $U(x) \geq U(y)$. It then follows by contraposition that $U(x) > U(y)$ implies $x \succ y$. We conclude that U as in (114) is a numerical representation of $\succ$.

Conversely, we assume that a numerical representation is given. We also assume that X is uncountable, otherwise there is nothing to prove. Let $J := \{[a, b] : a, b \in \mathbb{Q}, a < b, U^{-1}([a, b]) \neq \emptyset\}$. Then, for every $I \in J$, there exists $z_I \in \mathcal{X}$ with $U(z_I) \in I$. Put $A := \{z_I : I \in J\}$ and observe that A is countable.

The set A is almost the set $\mathcal{Z}$ we are after. A naive approach could be as follows. Suppose $y \succ x$, then $U(y) > U(x)$ and there are rational a and b such that $U(x) < a < b < U(y)$. The problem arises that it is not guaranteed that $U^{-1}([a, b])$ is non-void.

To remedy this, we will enlarge the set A with certain elements of A^c and consider thereto first the set $C := \{(x, y) \in A^c \times A^c : y \succ x \text{ and } \forall z \in A : x \succeq z \text{ or } z \succeq y\}$. Let $x, y) \in C$, but suppose that there exists $z \in \mathcal{X} \setminus A$ such that $y \succ z \succ x$. Then we can also find rational a and b such that $U(x) < a < U(z) < b < U(y)$ and therefore $I := [a, b] \in J$. By definition of A, we can then find $z_I \in A$ that then also has the property $U(x) < a \leq U(z_I) \leq b < U(y)$ and hence $y \succ z_I \succ x$. This contradicts $(x, y) \in C$. We conclude that if $(x, y) \in C$, then for all $z \in \mathcal{X}$ it holds that $x \succeq z$ or $z \succeq y$.

This implies the following observation. If $(x, y) \in C$ and $(x', y') \in C$, such that $U(x) \neq U(x')$ or $U(y) \neq U(y')$, then $(U(x), U(y)) \cap (U(x'), U(y')) = \emptyset$. We argue as follows. The situation $x \sim x'$ and $y \sim y'$ is ruled out by assumption. Therefore, assume w.l.o.g. that $x \not\sim x'$. Since $(x, y) \in C$, we must have $x \succeq x'$ or $x' \succeq y$, which implies that either $U(x) \geq U(x')$ or $U(x') \geq U(y)$. In the latter case, we are done. Let then the former inequality hold. Since also $(x', y') \in C$, we have $x' \succeq x$ or $x \succeq y'$. The first of these possibilities cannot happen, since we ruled out $x \sim x'$, and therefore the second one holds, and we obtain

$$U(x) \geq U(y'),$$

from which the conclusion follows as well. Knowing that the intervals $(U(x), U(y))$ with $(x, y) \in C$ are disjoint, we conclude that there are only countably many of them and it follows that the collection of these intervals can be written as a collection of intervals

$$(U(x), U(y)),$$

where x and y run through a countable subset of $\mathcal{X}$, B say. We put $\mathcal{Z} = A \cup B$, a countable set as well, and we will see that it is order dense. Take $x, y \in \mathcal{X} \setminus \mathcal{Z}$ with $y \succ x$. If there is $z \in A$ such that $y \succ z \succ x$, we are done. If such a z doesn't exist, then $(x, y) \in C$, in which case we have for instance $U(x) = U(z)$ for some $z \in B$. But then $y \succ z \succeq x$. $\square$

Not every preference relation admits a numerical representation. An example is the **lexicographical order** on $[0,1] \times [0,1]$, defined by $x \succ y$ if $x_1 > y_1$ or ($x_1 = y_1$ and $x_2 > y_2$).

Example B.6. Consider $\mathcal{X} = [0,1] \times [0,1]$ endowed with the lexicographical order $\succ$. Suppose $\succ$ admits a numerical representation U. Given $(\alpha, 1) \succ q(\alpha, 0)$, we have $d(\alpha) := U(\alpha, 1) - U(\alpha, 0) > 0$ for all $\alpha \in [0,1]$. Define

$$A_n := \{\alpha \in [0,1] \colon d(\alpha) > \frac{1}{n}\}.$$

Thus, $[0,1] = \bigcup_n A_n$. Since $[0,1]$ is uncountable, there must be a set A_m with infinitely many elements. In this set, we can choose for any positive integer N, real numbers $\alpha_0 < \ldots < \alpha_N$. Note that $U(\alpha_{i+1}, 0) > U(\alpha_i, 1)$, and so we get $U(\alpha_{i+1}, 0) - U(\alpha_i, 0) > d(\alpha_i) > \frac{1}{m}$. Hence we get

$$\begin{aligned} U(1,1) - U(0,0) = U(1,1) - U(\alpha_N, 0) + \sum_{i=0}^{N-1} (U(\alpha_{i+1}, 0) - U(\alpha_i, 0)) + U(\alpha_0, 0) - U(0,0) \\ > \frac{N}{m}. \end{aligned}$$

Letting $N \to \infty$ yields $U(1,1) - U(0,0) = \infty$, which is excluded.

Now, let's define various types of preference intervals. The first two are:

$$\begin{aligned} ((x, \to)) &:= \{y \in X : y \succ x\}, \\ ((\leftarrow, x)) &:= \{y \in X : x \succ y\}. \end{aligned}$$

Additionally, we use

$$((x, y)) \text{ for } ((\leftarrow, y)) \cap ((x, \to)).$$

Furthermore, we have

$$\begin{aligned} [[x, \to)) &= \{y \in X : y \succeq x\}, \\ ((\leftarrow, x]] &= \{y \in X : x \succeq y\} \text{ and so on.} \end{aligned}$$

Expressing the negative transitivity of $\succ$ using intervals, we get

$$((\leftarrow, x)) \cup ((y, \to)) = X \text{ if } x \succ y.$$

Definition B.7. Let $\mathcal{X}$ be a topological space. A **continuous preference relation** is any preference relation $\succ$ such that for every $x \in \mathcal{X}$, the sets $((x, \to))$ and $((\leftarrow, x))$ are open.

Assuming $\succ$ admits a numerical representation U, due to the identity

$$((x, \to)) = U^{-1}(U(x), \infty),$$

we see that $\succ$ is continuous if U is continuous. However, there are preference relations that are not continuous. Take for example the lexicographical order on $\mathcal{X} = [0,1] \times [0,1]$. The set $\{(y_1, y_2) \in \mathcal{X} \colon (y_1, y_2) \succ (\frac{1}{2}, \frac{1}{2})\}$ is not open in the standard topology.

Proposition B.8. Let $\mathcal{X}$ be a Hausdorff space, and let $\mathcal{X} \times \mathcal{X}$ be equipped with the product topology. Then, the following are equivalent:

i) $\succ$ is continuous.

ii) The set $\{(x, y) \colon y \succ x\}$ is open.

iii) The set $\{(x, y) \colon y \succeq x\}$ is closed.

Proof. First we show (**i**) $\Rightarrow$ (**ii**): Let $(x_0, y_0) \in M := \{(x, y) : y \succ x\}$. We show that there are open subsets U and V of $\mathcal{X}$ such that $(x_0, y_0) \in U \times V \subset M$. Suppose first that the preference interval $((x_0, y_0)) \neq \emptyset$. Pick a z from this preference interval, then $y_0 \succ z \succ x_0$. The sets $U := ((\leftarrow, z))$ and $V := ((z, \rightarrow))$ are open and contain x_0 and y_0 respectively. Moreover, one quickly sees that $U \times V \subset M$. If the preference interval $((x_0, y_0))$ is empty, we choose $U = ((\leftarrow, y_0))$ and $V = ((x_0, \rightarrow))$. Take $(x, y) \in U \times V$. Then $y_0 \succ x$ and $y \succ x_0$. To show that $y \succ x$, we assume the contrary. By negative transitivity we must have $y_0 \succ y$. But then $y \in ((x_0, y_0))$, which was empty. Contradiction.
(**ii**) $\Rightarrow$ (**iii**): It follows from (ii) that also $\{(x, y) : x \succ y\}$ is open. But its complement is just $\{(x, y) : y \succeq x\}$.
(**iii**) $\Rightarrow$ (**i**): Since X is Hausdorff, every singleton $\{x\}$ is closed and so $\{x\} \times \mathcal{X}$ is closed in the product topology. By assumption, then also $\{x\} \times \{y : y \succeq x\} = \{x\} \times X \cap \{(u, v) : v \succeq u\}$ is closed. But then, the set $\{y : y \succeq x\}$ is closed in X since a product set is closed iff all factors are closed, and so $\{y : x \succ y\}$ is open. In a similar way one proves that $\{y : y \succ x\}$ is open. □

Proposition B.9. Let $\mathcal{X}$ be a connected topological space endowed with a continuous preference order $\succ$. If $\mathcal{X}$ is dense, then $\mathcal{X}$ is also order dense. If $\mathcal{X}$ is separable, then $\succ$ admits a numerical representation.

Proof. First we rule out the trivial situation in which all elements of $\mathcal{X}$ are indifferent. So, we can take $x, y \in \mathcal{X}$ with $y \succ x$. Observe that $y \in ((x, \rightarrow))$ and $x \in ((\leftarrow, y))$, so both open preference intervals are non-empty. Moreover, their union is $\mathcal{X}$, because of negative transitivity. Then we must have that $((x, \rightarrow)) \cap ((\leftarrow, y)) \neq \emptyset$, because $\mathcal{X}$ is connected. The intersection is open as well, so it must contain a z from $\mathcal{Z}$, since $\mathcal{Z}$ is dense. Then $y \succ z \succ x$, and so $\mathcal{Z}$ is order dense. If $\mathcal{X}$ is separable, there exists a countable dense and thus order dense subset. Now, apply Theorem B.5 to conclude the proof. □

Without proof we give the following result:

Theorem B.10. Let $\mathcal{X}$ be a connected and separable topological space, endowed with a continuous preference order $\succ$. Then $\succ$ admits a continuous numerical representation.

C Analysis

Proposition C.1. For $p \in [0, \infty]$, the dimension of the linear space $L^p(\Omega, \mathcal{F}, \mathbb{P})$ is given by

$$\dim L^p(\Omega, \mathcal{F}, \mathbb{P}) = \sup\{n \in \mathbb{N} \mid \exists \text{ a partition } A_1, \ldots, A_n \text{ of } \Omega \text{ with } A_i \in \mathcal{F} \text{ and } \mathbb{P}[A_i] > 0\}. \tag{115}$$

Moreover, $n := \dim L^p(\Omega, \mathcal{F}, \mathbb{P}) < \infty$ if and only if there exists a partition of Ω into n atoms of $(\Omega, \mathcal{F}, \mathbb{P})$.

Proof. Suppose there is a partition $A_1, \ldots, A_n$ of Ω such that $A_i \in \mathcal{F}$ and $\mathbb{P}[A_i] > 0$ for all i. Consider the corresponding indicator functions $1_{A_1}, \ldots, 1_{A_n}$, which are non-zero elements of $L^p(\Omega, \mathcal{F}, \mathbb{P})$. These indicator functions are linearly independent. Hence, $\dim L^p \geq n$.
It therefore remains to consider the case where the right-hand side of (115) is finite, say n_0. Let $A_1, \ldots, A_{n_0}$ be a partition of Ω that attains this maximum, meaning n_0 is the largest integer for which such a partition exists. Each A_i must be an atom of $(\Omega, \mathcal{F}, \mathbb{P})$; otherwise, we could refine the partition further and increase n_0.
By definition of an atom, any $Z \in L^p(\Omega, \mathcal{F}, \mathbb{P})$ is essentially constant on each A_i. Denote this constant by z_i. Then we have

$$Z = \sum_{i=1}^{n_0} z_i 1_{A_i} \quad \mathbb{P}\text{-a.s.}$$

Thus, every element of L^p can be represented as a linear combination of the n_0 indicator functions $1_{A_1}, \ldots, 1_{A_{n_0}}$. Since these are linearly independent, we conclude $\dim L^p = n_0$. □

Lemma C.2 (Jensen's Inequality)**.** Let $u\colon S \to \mathbb{R}$ be concave with finite $m(\mu)$ and $\int u \, \mathrm{d}\mu$. Then, $\int u \, \mathrm{d}\mu \leq u(m(\mu))$. If u is strictly concave and μ is not degenerate, then $\int u \, \mathrm{d}\mu < u(m(\mu))$.

C.1 Hahn-Banach Separation Theorem

The following theorem is the famous geometric version of Hahn-Banach's theorem. Usually, it is called the *Hahn-Banach separation theorem.* It generalizes the *Hyperplane separation theorem* on Euclidean spaces to general topological vector spaces.

Theorem C.3 (Hahn-Banach Separation Theorem)**.** Let A and B be non-empty convex subsets of a real locally convex topological vector space X. If $\operatorname{Int} A \neq \varnothing$ and $B \cap \operatorname{Int} A = \varnothing$, then there exists a continuous linear functional f on X such that

$$\sup f(A) \leq \inf f(B), \quad \text{and} \quad f(a) < \inf f(B) \text{ for all } a \in \operatorname{Int} A,$$

where such an f is necessarily non-zero.