Mathematical Finance I

Christoph Belak

Winter Term 2019/20

Christoph Belak

Institut für Mathematik
Technische Universität Berlin
Straße des 17. Juni 136
10623 Berlin

belak@math.tu-berlin.de

Contents

1. Financial Markets, Derivatives, Arbitrage 1
- 1.1. The Mathematics of Financial Markets 1
- 1.2. Basic Financial Securities 2
- 1.3. Absence of Arbitrage and Put-Call Parity 8

2. Finite Financial Market Models 15
- 2.1. Financial Markets and Stochastic Processes 15
- 2.2. Filtrations and the Dynamic Resolution of Information 19
- 2.3. Trading Strategies and the Wealth Process 22

3. Risk Neutral Pricing and Fundamental Theorems 29
- 3.1. Arbitrage and the First Fundamental Theorem 29
- 3.2. Risk Neutral Pricing in Finite Market Models 39
- 3.3. Completeness and the Second Fundamental Theorem 43
- 3.4. Incomplete Markets and Superhedging Duality 46

4. American Options and Optimal Stopping 63
- 4.1. Optimal Stopping Problems and the Snell Envelope 64
- 4.2. Risk Neutral Pricing of American Options 77

5. Risk-Neutral Pricing in the Black-Scholes Model 89
- 5.1. Continuous Time Limit of the CRR Model 89
- 5.2. Brownian Motion and Stochastic Integration 94
- 5.3. Option Pricing in the Black-Scholes Model 118

A. Stopping Times and Hitting Times **143**
A.1. Stopping Times and the Sigma Field of the τ-Past 143
A.2. Hitting Times and Stopped Processes 147

B. Martingale Theory **157**
B.1. Doob's Inequalities 157
B.2. Continuous Martingales of Finite Variation 160

Index **165**

Vocabulary **169**

Chapter 1 Financial Markets, Derivatives, Arbitrage

The aim of this introductory chapter is to give you a first flavor of what **Mathematical Finance** is about and to expose you to a few important notions and concepts which we shall encounter numerous times throughout this course.

1.1 The Mathematics of Financial Markets

Mathematical Finance comes in many shapes and sizes. Typical themes are for example:

- **Valuation/Pricing**: Compute fair prices of financial securities.
- **Risk Management**: Assess risks associated with financial securities.
- **Optimal Investment**: Derive optimal trading strategies.

In this course, our main focus will be on valuation with a short excursion into optimal investment. To facilitate this, we furthermore have to do our fair share of **financial modeling**, i.e. we have to construct and assess mathematical models of financial markets.

A **financial market** is any kind of platform on which institutions and/or individuals come together to trade financial securities. The term platform need not refer to a physical location but may also be an electronic system. Typical examples of financial markets are **exchanges** such as the New York

Stock Exchange or Börse Frankfurt. Typical examples of **securities** traded in financial markets are bonds, stocks, futures, swaps, options, etc.[1]

Modeling financial markets can be quite sophisticated due to a plethora of rules and regulations but also due to the individuality of objectives and attitudes of market participants. Throughout this course, we shall hence restrict to what we call **perfect financial markets**. Such markets are highly idealized in that we assume that

- ▷ securities can be bought and sold in **arbitrary quantities**;
- ▷ **short sales**, i.e. negative positions in securities, are allowed without restrictions or costs;
- ▷ trades do not involve any **transaction costs**;
- ▷ traders are not liable to **taxes**;
- ▷ trades do not **impact** market prices;
- ▷ all market participants have **full information** on all prices;
- ▷ there are no **liquidity** issues;
- ▷ traders are **rational**;
- ▷ ...

Any deviation from these assumptions is called a **market friction**. Much of today's research in Mathematical Finance is concerned with the effects such frictions have on financial markets, trading behavior, and price formation.

1.2 Basic Financial Securities

Throughout this course, we will encounter a number of different securities such as bonds, stocks, and derivative products. The aim of this section is to give you an overview of these securities. Let us start with bonds.

[1] Do not worry if you have never heard of these securities before. The (for our purposes) most important securities will be discussed in the next section.

Definition 1.1 (Bond; Zero-Coupon Bond). A **(coupon) bond** with maturity T, face value C, coupons $c_1, \ldots, c_n$, and coupon dates $t_1 < t_2 < \ldots < t_n \leq T$ is a security which guarantees its owner fixed payments of $c_1, \ldots, c_n$ at the dates $t_1, \ldots, t_n$ as well as a payment of C at time T. A **zero-coupon bond** with maturity T is a bond without coupon payments and a face value of 1, i.e. a security which guarantees its owner a payment of 1 at time T. We denote by $B(t,T)$ the price of such a zero-coupon bond at time $t \leq T$. ◇

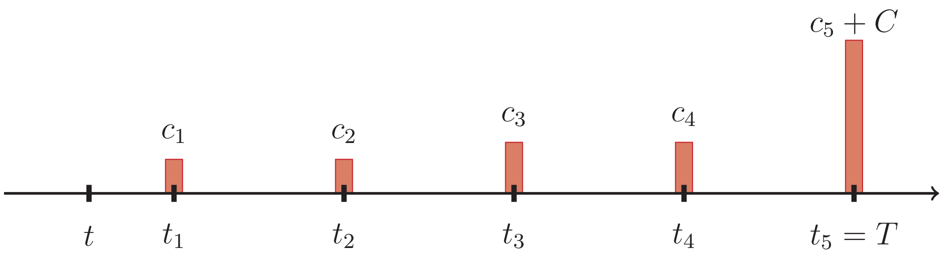

Figure 1.1. Cash flow of a coupon bond.

An example of a bond is the German government bond **Bundesanleihe** which has a maturity of either 10 or 30 years and which pays yearly coupons of currently between 0% and 6.5% of the face value. Bond markets or, more generally, **fixed-income markets** are very large. As of September 2019, the face values of all Bundesanleihen currently available on the market add up to 741.5 billion Euro alone. More information (in German) on the Bundesanleihe and other German government bonds can be found on the

Website of the Deutsche Finanzagentur.

In perfect financial markets, all we really need are zero-coupon bonds as any other bond may be thought of as a portfolio of zero-coupon bonds: The payoff profile of a coupon bond is the same as that of a portfolio consisting of c_i zero-coupon bonds with maturity t_i for each $i = 1, \ldots, n$ and C zero-coupon bonds with maturity T. The price of a bond at time $t \leq t_1$ is thus[2]

$$\sum_{i=1}^{n} c_i B(t, t_i) + C B(t, T).$$

The price $B(t,T)$ of a zero-coupon bond on the other hand is clearly just the time-t value of one unit of money received at time T and hence we can use

[2] You will be asked to give a rigorous proof of this pricing formula in Exercise 3 below.

this price for **discounting** purposes. This point of view can be used to come up with models for the price. Simple yet common **examples** are

$$B(t,T) \triangleq \frac{1}{(1+r)^{T-t}} \qquad \text{or} \qquad B(t,T) \triangleq e^{-r(T-t)}, \qquad t \leq T,$$

where $r \in \mathbb{R}$ is a constant **interest rate**. Note that in these examples a strictly positive interest rate implies $B(t,T) < 1$ whereas a strictly negative interest rate implies $B(t,T) > 1$.

Since we allow short sales, we have the possibility to hold negative positions in the zero-coupon bond. This situation corresponds to a **loan**: A short sale of one zero coupon bond means we receive $B(t,T)$ at time t but are obliged to pay 1 at time T.

Another important class of securities are the **derivative products**.

Definition 1.2 (Derivative; Underlying; Option)**.** A **derivative** is a security whose payoffs are determined by other reference securities, the so-called **underlyings**. A derivative with only positive[3] payoffs is called **option**. ◇

Note that we have already encountered a derivative: The coupon bond may be thought of as an option on zero-coupon bonds. From the point of view of pricing, however, this is a rather boring example as the price of the option (the coupon bond) is just a linear combination of the prices of the underlyings (the zero-coupon bonds) and not much mathematics is needed for this. The situation becomes more sophisticated as soon as the option's payoff depends nonlinearly on the underlyings' prices. The two most important examples of such options are **European call and put options**.

Definition 1.3 (European Call Option; European Put Option)**.** A **European Call Option** on an underlying S with maturity T and strike price K grants its owner the right (but not the obligation) to buy the underlying S at time T for the price K. Similarly, a **European Put Option** on an underlying S with maturity T and strike price K grants its owner the right (but not the obligation) to sell the underlying S at time T for the price K. ◇

[3] Here and in the sequel, we use the convention that positive means greater than or equal to zero and strictly positive means strictly greater than zero.

It is important to note that the strike price K is agreed upon at the inception date of these options and that the owner of a call option is not forced to buy the underlying at time T at the price K. If, for example, the price $S(T)$ of the underlying at time T is smaller than the strike price K, it is cheaper for the owner of the option to just buy the underlying at the market. In this scenario, the option has hence become worthless to its owner. On the other hand, if $S(T)$ exceeds the strike price K, the owner of the call can exercise the option to buy the underlying at the price K and immediately sell it at the market for $S(T)$. Thus, in this scenario, the owner of the call makes a profit of $S(T) - K$. Combining these two scenarios, we see that the **payoff of the European call** at time T is

$$C(T) \triangleq (S(T) - K)_+ \triangleq \max\{S(T) - K, 0\}.$$

In particular, we see that the payoff now depends nonlinearly on the price of the underlying. A similar argument shows that the **payoff of the European put** option is

$$P(T) \triangleq (K - S(T))_+ \triangleq \max\{K - S(T), 0\},$$

i.e. the owner of the put option makes a profit if the price of the underlying falls below the strike price K.

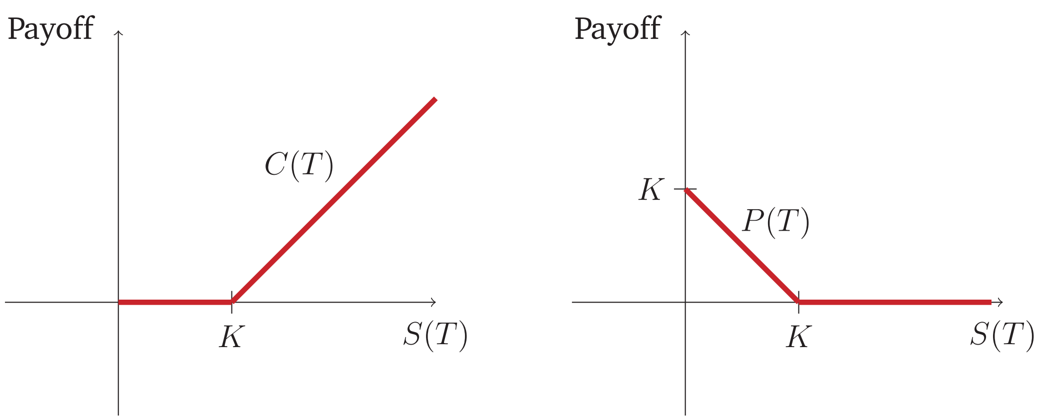

Figure 1.2. Payoff diagrams of European call (left) and put (right) options.

Call options can be used to lock in an upper bound on the price of a security. This makes call options an important **hedging instrument**: Think of an airline company which knows today how much fuel will be needed at a future date T. To lock in an upper bound on the price to be paid for the fuel, the company can simply buy call options as a protective measure against rising prices. Put options on the other hand can be used to hedge against decreasing prices. You may think of a farmer who wants to lock in a minimal price for the agricultural produce. In this sense, call and put options have a **stabilizing effect** on the economy.

Another purpose of call and put options is **speculation**. Suppose you possess one unit of money at time t and you are sure that the time-T stock price $S(T)$ of Volkswagen will be larger than today's price $S(t)$. One possibility to make a profit could be to use all your cash to buy $1/S(t)$ shares of Volkswagen stock today (this will cost you exactly 1 unit of money) and sell them at time T. Your profit at time T is then

$$\frac{1}{S(t)}\big(S(T) - S(t)\big).$$

The second possibility is to use all your cash to buy $1/C(t)$ call options with maturity T and strike $K = S(t)$ at the current option price $C(t)$. Assuming $S(T) > S(t)$, your profits are

$$\frac{1}{C(t)}\big(S(T) - K\big) = \frac{1}{C(t)}\big(S(T) - S(t)\big),$$

which is typically significantly larger than what you would receive by just trading in the stock since $C(t)$ is usually much smaller than $S(t)$.

The downside of this strategy is its riskiness. If $S(T) \leq S(t)$ you lose all your money if you follow the option strategy, whereas you are left with $S(T)/S(t) < 1$ at time T if you invest only in the stock. In this sense, options can also have a **destabilizing** effect on financial markets.

Exercise 1 (Option Trading and Payoff Diagrams)**.** Suppose that $K_1 > K_2$. Draw payoff diagrams for the following strategies:

(i) **Bearish Spread**: Buy a European call with strike K_1 and sell a European call with strike K_2.

(ii) **Strangle**: Buy a European call with strike K_1 and a European put with Strike K_2.

How does the market have to evolve for you to make profits with these two strategies? ◇

Exercise 2 (Financial Engineering)**.** Construct portfolios consisting of European calls and puts such that

(i) you make a profit if the time-T price $S(T)$ of the underlying does not differ by much from today's price $S(t)$;

(ii) your profit first grows linearly as a function of $S(T)$ if $S(T)$ is greater than $S(t)$ and constant if $S(T)$ is significantly larger than $S(t)$. ◇

When call and put options are exercised, the underlying is usually not delivered physically, but the owner receives the payoff $C(T)$ or $P(T)$ directly in monetary units. An exception from this common practice are commodity options if the owner of the option is really interested in buying/selling the underlying commodity as in the airline and farmer examples above.

The term **European** in Definition 1.3 indicates that the options can only be exercised at maturity T. There are also variants of these options which can be exercised at any time prior to maturity, in which case we refer to them as **American** options.

Definition 1.4 (American Call Option; American Put Option)**.** An **American call option** on an underlying S with maturity T and strike price K grants its owner the right to buy the underlying S at any time $t \leq T$ for the price K. Similarly, an **American put option** on an underlying S with maturity T and strike price K grants its owner the right to sell the underlying S at any time $t \leq T$ for the price K. ◇

Since American options can always be exercised at time T, it is clear that these options are more valuable than their European counterparts.[4] There are also options with less common exercise rights such as **Bermudan** or **Canary options**, for which the possible exercise dates are somewhere in between European and American options – just like their geographic locations.

[4] Again, this statement can be proved rigorously and you will be asked to do so in Exercise 4 below.

Call and put options are the most common types of options which is why they are also referred to as **vanilla options** whereas all other options are referred to as **exotic options**. We will encounter a few exotic options later but content ourselves with vanillas for now.

1.3 Absence of Arbitrage and Put-Call Parity

A natural question is whether we can say anything about the prices of vanilla options at time $t < T$. On the one hand, these options are traded in the market, so their prices should be determined by **supply and demand**. On the other hand, the values of these options are determined by the underlyings' prices, so there is hope that we can express option prices in terms of the prices of their underlyings. Settling the pricing question is involved and we will spend most of this course to come up with a good answer. Nevertheless, there are some particular situations in which determining prices is easy (as in the linear case; recall that bonds are options on zero-coupon bonds) or in which we can at least derive price bounds.

Consider for example a European call option. Intuitively, we expect that the price $C(t)$ of this option at time $t < T$ should not exceed $S(t)$ as buying the option only gives us the right to buy the underlying for K, whereas we can always just buy the underlying right away for $S(t)$. Can we come up with a formal argument to prove this statement? One possibility is to argue by contradiction, i.e. we suppose that $C(t) > S(t)$ and see where it takes us. As the price of the call seems unrealistically high compared to the underlying, we could sell the option short and buy the underlying, resulting in a profit of $C(t) - S(t) > 0$. At maturity T, we unwind our portfolio by buying back the call option and selling the underlying, resulting in another profit of

$$S(T) - C(T) = S(T) - (S(T) - K)_+ \geq 0.$$

In other words, we have found a trading strategy which requires zero initial capital but guarantees a profit without the risk of any losses. This is what we call an **arbitrage opportunity**.

Definition 1.5 (Arbitrage Opportunity)**.** An **arbitrage opportunity** is a trading strategy which guarantees a riskless profit without obliging the trader to any net cash outflows now or in the future. ◇

Note that buying a zero-coupon bond at a price of $B(t,T) < 1$ at time t and receiving 1 at time T is not an arbitrage opportunity. While we clearly make a profit of $1-B(t,T) > 0$ over the entire investment period, there is an initial cash outflow of $B(t,T)$. For a trading strategy to be an arbitrage opportunity, **all cash flows** must be positive and there must be the possibility of at least one strictly positive payoff. In essence, arbitrage opportunities are trading strategies which generate arbitrary amounts of money out of nothing.

Even if you encounter an arbitrage opportunity in a real-world financial market, you will never be able to exploit it. There is always a market participant who is faster than you and who will exploit the arbitrage opportunity with vast amounts of capital, causing it to disappear immediately. A reasonable assumption on our financial market models is therefore

Absence of Arbitrage: There are no arbitrage opportunities.

The importance of this assumption cannot be stressed enough. Absence of arbitrage is the cornerstone of option pricing theory in mathematical finance and will be the most important theme throughout this course.

To get a first impression of the power of absence of arbitrage, let us subsequently fix a financial market consisting of

- a **zero-coupon bond** B with maturity T trading at a price of $B(t,T)$ at time $t \leq T$;
- a **security** S (e.g. a stock) trading at a price of $S(t)$ at time $t \leq T$;
- a **European call** C and a **European put** P with maturity T, underlying S, and strike K, with prices $C(t)$ and $P(t)$ at time $t \leq T$, respectively;
- an **American call** C_A and an **American put** P_A with maturity T, underlying S, and strike K, with prices denoted by $C_A(t)$ and $P_A(t)$ at time $t \leq T$, respectively.

Moreover, assume that the financial market is **free of arbitrage** opportunities. In this setting, using solely the absence of arbitrage assumption, we derive universal relations between the prices of the securities. The essential argument needed for this is formulated in the following lemma.

Lemma 1.6 (Comparison of Portfolios). *Let $x \in \mathbb{R}^4$ such that*

$$x_1 B(T,T) + x_2 S(T) + x_3 C(T) + x_4 P(T) \geq 0.$$

Then it holds that

$$x_1 B(t,T) + x_2 S(t) + x_3 C(t) + x_4 P(t) \geq 0, \qquad t < T. \qquad \diamond$$

Proof. We argue by contradiction and assume that there exists $t < T$ with

$$x_1 B(t,T) + x_2 S(t) + x_3 C(t) + x_4 P(t) < 0.$$

At time t, setup a portfolio consisting of

x_1 zero-coupon bonds B,	x_2 securities S,
x_3 European calls C,	x_4 European puts P,

and consider the trading strategy which holds this portfolio on $[t,T)$ and liquidates all positions at time T. Setting up this portfolio at time t yields an immediate profit of

$$-x_1 B(t,T) - x_2 S(t) - x_3 C(t) - x_4 P(t) > 0$$

and unwinding this portfolio at time T yields another profit of

$$x_1 B(T,T) + x_2 S(T) + x_3 C(T) + x_4 P(T) \geq 0.$$

As there are no other cash in- or outflows, we hence see that this strategy is an arbitrage opportunity, which is the desired contradiction. □

The preceding lemma simply states that any portfolio which has a positive value at time T must also have a positive value at any time prior to T. It is expedient to familiarize yourself with the type of argument used in the proof of this lemma to get a better feeling of the consequences of the absence of arbitrage assumption.

Exercise 3 (Prices of Coupon Bonds). Show that absence of arbitrage implies that the time-t price of a coupon bond with face value C and coupon payments $c_1, \ldots, c_n$ at times $t_1, \ldots, t_n$ is given by

$$\sum_{i=1}^{n} c_i B(t,t_i) + C B(t,T), \qquad t \leq t_1. \qquad \diamond$$

Exercise 4 (American vs. European Options)**.** Use the absence of arbitrage assumption to show that

$$C_A(t) \geq C(t) \qquad \text{and} \qquad P_A(t) \geq P(t), \qquad t \leq T. \qquad \diamond$$

Exercise 5 (Dividend Payments)**.** Consider a dividend paying stock S and suppose that its value just before a dividend payment of $D > 0$ at time t is denoted by $S(t-)$. What is the price $S(t)$ of the stock immediately after the dividend payment? Justify your answer by arbitrage arguments. ◇

Exercise 6. Suppose that $B(t,T) \leq 1$ for all $t \leq T$ and let $T_1 < T_2 \leq T$. Denote by C_1 and C_2 European calls on the same underlying with the same strike price but differing maturities T_1 and T_2, respectively. Show that

$$C_1(t) \leq C_2(t), \qquad t \leq T_1. \qquad \diamond$$

Note that upon replacing x by $-x$ in Lemma 1.6 (Comparison of Portfolios), it follows that a portfolio with a value of zero at time T must have a value of zero at any time prior to T. This might seem rather obvious and innocent, but has tremendous consequences such as the following theorem.

Theorem 1.7 (Put-Call Parity)**.** *It holds that*

$$C(t) + KB(t,T) = P(t) + S(t), \qquad t \leq T. \qquad \diamond$$

Proof. Consider a portfolio consisting of 1 European call, K zero-coupon bonds, -1 European put, and -1 security S. The value of this portfolio at time T is

$$\begin{aligned} C(T) &+ KB(T,T) - P(T) - S(T) \\ &= (S(T) - K)_+ + K - (K - S(T))_+ - S(T) = 0. \end{aligned}$$

Hence, by Lemma 1.6 (Comparison of Portfolios),

$$C(t) + KB(t,T) - P(t) - S(t) = 0, \qquad t \leq T. \qquad \square$$

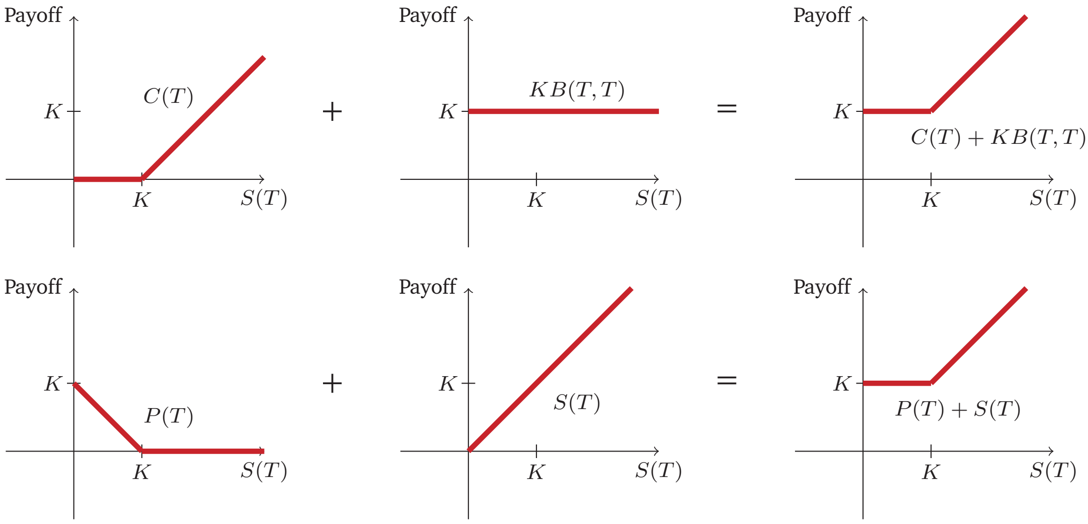

Figure 1.3. Graphical proof of the put-call parity.

The **Put-Call Parity** states that one of the securities in our market is redundant. For example, there is no need to offer European calls in this market as the same payoff could be achieved by a portfolio consisting of a European put, the underlying, and $-K$ zero-coupon bonds. We say that the European call can be **replicated** by the other securities. By absence of arbitrage, replication guarantees that the price of the European call is uniquely determined by the prices of the remaining assets.

Exercise 7 (Digital Put-Call Parity)**.** A **digital call** (resp. **digital put**) with maturity T on an underlying S pays the fixed amount of one unit of money at time T if $S(T) \geq K$ (resp. $S(T) < K$). Derive and prove a put-call parity for digital options. ◇

Put-Call Parity is so powerful because it is a **model-free result**. It does not matter how we model the price of the zero-coupon bond or the underlying. As long as there are no arbitrage opportunities in our financial market, put-call parity holds!

Another consequence of the absence of arbitrage assumption is that American calls should never be exercised before time T if interest rates are positive. In particular, this implies the surprising result that the prices of the

European and American call coincide.

Theorem 1.8 (No Early Exercise of the American Call)**.** *Assume $B(t,T) < 1$ for all $t < T$. Then it is never rational to exercise the American call before time T and hence*

$$C_A(t) = C(t), \qquad t \leq T. \qquad \diamond$$

Proof. We recall that $C_A(t) \geq C(t)$ and note that

$$C(T) = (S(T) - K)_+ \geq S(T) - KB(T,T).$$

By Lemma 1.6 (Comparison of Portfolios) and using $B(t,T) < 1$, this implies

$$C_A(t) \geq C(t) \geq S(t) - KB(t,T) > S(t) - K, \qquad t < T.$$

This is to say that, at any time $t < T$, the American call is always worth strictly more than the payoff one would receive upon exercising it. It is thus not rational to exercise the call whenever $t < T$. But then the only time at which it makes sense to exercise the American call is at maturity T, meaning that a rational investor treats the American call just like a European call, i.e. $C_A(t) = C(t)$ for all $t \leq T$. □

It is crucial in the above argument that $B(t,T) < 1$ for all $t < T$, which we think of as strictly positive interest rates. We also note that the same result is not true for the American put. Assume for example that the underlying goes bankrupt, i.e. there exists $\tau < T$ such that $S(t) = 0$ for all $\tau \leq t \leq T$. Exercising the American put at the time of default yields an immediate payoff of K, which could be invested into $K/B(\tau,T)$ zero-coupon bonds, yielding a profit of $K/B(\tau,T) > K$ at time T. On the other hand, exercising the American put at time T only yields a payoff of K at time T, so it is rational to exercise the put early.

Exercise 8 (American Put-Call Inequality and Arbitrage Bounds)**.** Suppose that $B(t,T) < 1$ for all $t \leq T$. Show that absence of arbitrage implies the price bounds

$$C_A(t) - S(t) + KB(t,T) \leq P_A(t) \leq C_A(t) - S(t) + K, \qquad t \leq T,$$

as well as

$$P(t) \leq P_A(t) \leq \big(1 - B(t,T)\big)P(t), \qquad t \leq T. \qquad \diamond$$

Exercise 9 (Arbitrage from Price Quotes)**.** Suppose you are given the following price quotes in today's newspaper:

$$C(t) = 11.42, \qquad P(t) = 3.94, \qquad P_A(t) = 5.67,$$
$$S(t) = 67.13, \qquad B(t,T) = e^{-0.02}.$$

Is there an arbitrage opportunity? If so, which trading strategy is required to exploit the arbitrage opportunity? ◇

Chapter 2

Finite Financial Market Models

While put-call parity is a powerful result due to its universality, it is not sufficient to compute option prices. Even if we specify a model for the underlying and the zero-coupon bond, we still do not have sufficient information to compute the call price as the put price is unknown as well.

To facilitate the computation of option prices, it becomes necessary to set up a financial market model, and this is exactly what we shall do in this chapter. The market model introduced here is kept as simple as possible so that we do not get distracted by mathematical technicalities. Still, the model is sufficiently rich to understand the basic concepts underlying option pricing theory. In later chapters, we will then see how to extend the results to more advanced and arguably more realistic market models.

To be precise, the market model introduced in this chapter is that of a **finite financial market** in that we assume that there are only finitely many trading dates and that at each trading date the securities in this market can only take finitely many values.

2.1 Financial Markets and Stochastic Processes

We identify financial markets with the prices of the securities traded there. For $d \in \mathbb{N}$, the financial market considered in this chapter consists of $d+1$ securities which can be traded at the trading dates $0, 1, \ldots, T$ for $T \in \mathbb{N}$. For each $i \in \{0, 1, \ldots, d\}$, we denote by

$$S^i(t) \text{ the price of the } i^{\text{th}} \text{ security at time } t = 0, 1, \ldots, T.$$

Since, in general, we cannot predict prices of securities and future prices appear random to us, it is reasonable to assume that each

$$S^i(t) : \Omega \to \mathbb{R}_+, \qquad \omega \mapsto S^i(t,\omega)$$

is a random variable defined on a given probability space $(\Omega, \mathfrak{A}, \mathbb{P})$.

In the language of stochastic analysis, for each $i = 0, 1, \dots, d$, the family $S^i = \{S^i(t)\}_{t=0,1,\dots,T}$ is referred to as an $\mathbb{R}_+$-valued **stochastic process** in the sense of the following definition.

Definition 2.1 (Stochastic Process)**.** Let $(S, \mathfrak{S})$ be a measurable space and $\mathcal{T} \subseteq \mathbb{R}$ a time index set. A family $\{X(t)\}_{t\in\mathcal{T}}$ of S-valued random variables defined on the same probability space $(\Omega, \mathfrak{A}, \mathbb{P})$ is called an S-valued **stochastic process**. ◇

Writing $S(t) = (S^0(t), S^1(t), \dots, S^d(t))$ for each $t = 0, 1, \dots, T$, we see that a financial market $S = \{S(t)\}_{t=0,1,\dots,T}$ is thus far simply assumed to be an $\mathbb{R}_+^{d+1}$-valued stochastic process.

As discussed above, to avoid unnecessary technicalities, we assume that each $S^i(t)$ takes only finitely many values, which we can guarantee by assuming that

$$\Omega = \{\omega_1, \dots, \omega_N\} \text{ is finite.}$$

The σ-field $\mathfrak{A}$ on Ω can then safely be assumed to be the power set, i.e. the set of all subsets of Ω. With this, whenever $(S, \mathfrak{S})$ is a measurable space, any function $X : \Omega \to S$ is $\mathfrak{A}$-$\mathfrak{S}$-measurable, i.e. is a random variable. Moreover, for any such X taking values in $\mathbb{R}$, the expectation

$$\mathbb{E}[X] \triangleq \sum_{\omega\in\Omega} X(\omega)\mathbb{P}[\{\omega\}]$$

is always well-defined and finite. In finite financial market models, we therefore neither have to worry about **measurability** nor **integrability**.

Without loss of generality, we furthermore rule out impossible elementary events by assuming that

$$\mathbb{P}[\{\omega\}] > 0 \qquad \text{for all } \omega \in \Omega.$$

We think of each $\omega \in \Omega$ as a possible state of the world which occurs with probability $\mathbb{P}[\{\omega\}] > 0$. In this state of the world, the prices of the securities evolve in time as

$$S^i(\cdot, \omega) : \{0, 1, \dots, T\} \to \mathbb{R}_+, \qquad t \mapsto S^i(t, \omega) \qquad \text{for } i = 0, 1, \dots, d.$$

We refer to this mapping as the **trajectory** or **path** of the i^{th} security in the state of the world ω; see Figure 2.1.

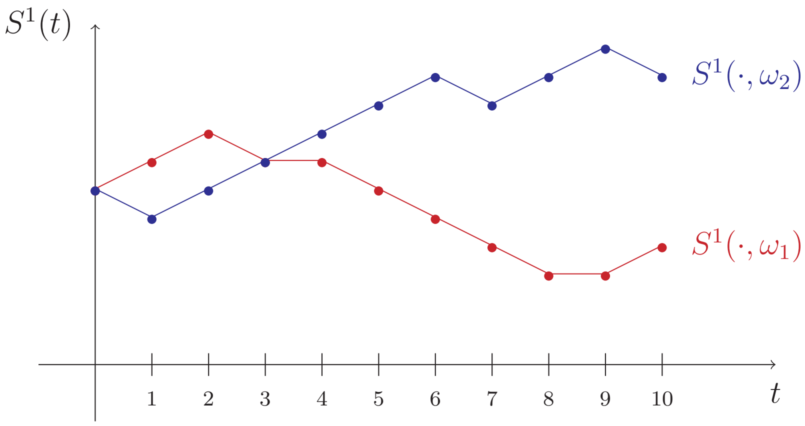

Figure 2.1. Two trajectories of the price process S^1.

Example 2.2 (CRR Model)**.** A popular model is the **Cox-Ross-Rubinstein** (CRR) model consisting of two securities S^0 and S^1. In this model, the price of S^0 evolves deterministically in time in that

$$S^0(t) \triangleq (1 + r)S^0(t - 1), \qquad t = 1, \dots, T, \qquad S^0(0) \triangleq 1,$$

for some constant $r > -1$. We may think of S^0 as the evolution of one unit of money in time and hence we may think of r as the interest rate between each trading date. Unwinding the recursive definition of S^0 yields

$$S^0(t) = (1 + r)^t, \qquad t = 0, 1, \dots, T;$$

see Figure 2.2 for a sketch of the dynamic evolution of S^0.

The security S^1 on the other hand evolves randomly in time in that between each trading date it moves by a factor of $u > -1$ with probability $p \in (0, 1)$ or by another factor of $d > -1$ for $d < u$ with probability $1 - p$. To make

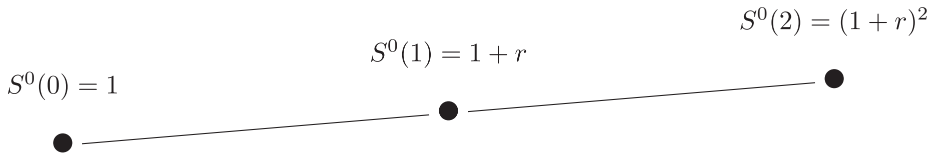

Figure 2.2. Evolution of S^0 in the CRR model.

this mathematically rigorous, we take as given a finite sequence $R_1, \dots, R_T$ of independent Bernoulli random variables taking values in $\{d, u\}$ such that

$$\mathbb{P}[R_t = u] = p \qquad \text{and} \qquad \mathbb{P}[R_t = d] = 1 - p, \qquad t = 1, \dots, T.$$

The evolution of S^1 is then defined recursively by

$$S^1(t) \triangleq (1 + R_t)S^1(t-1), \qquad t = 1, \dots, T, \qquad S^1(0) \triangleq s,$$

where $s > 0$ denotes the deterministic initial price. As in the case of S^0, we can unwind the recursive definition and express S^1 more compactly as

$$S^1(t) = s \prod_{\tau=1}^{t} (1 + R_\tau), \qquad t = 0, 1, \dots, T.$$

The evolution of S^1 is sketched in Figure 2.3 below.

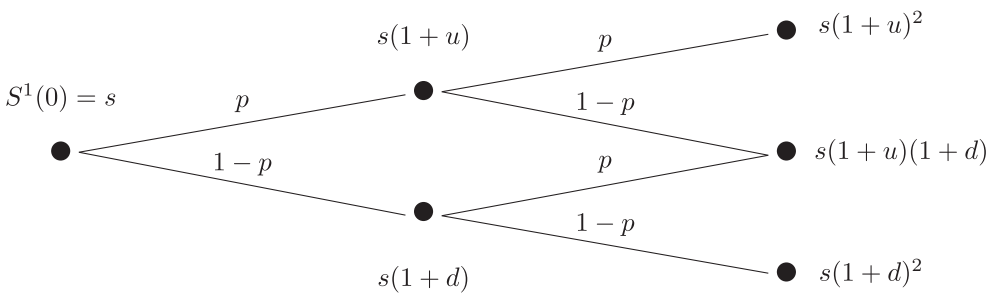

Figure 2.3. Evolution of S^1 in the CRR model.

Note that independence of $R_1, \dots, R_T$ implies that each $S^1(t)$, $t = 0, 1, \dots, T$, has a binomial distribution, which is why the CRR model is also often referred to as the **binomial model**. ◇

Example 2.3 (Trinomial Model)**.** Another popular financial market model is the **trinomial model** which is constructed just as the CRR model with the only difference being that the random variables $R_1, \ldots, R_T$ can take three different values $d, m, u \in \mathbb{R}$ with $-1 < d < m < u$ and

$$\mathbb{P}[R_t = u] = p_1, \qquad \mathbb{P}[R_t = m] = p_2, \qquad \mathbb{P}[R_t = d] = 1 - p_1 - p_2$$

for all $t = 1, \ldots, T$, where $p_1, p_2 \in (0, 1)$ with $p_1 + p_2 < 1$. ◇

In the CRR model, the security S^0 can be used for discounting as it describes the evolution of one unit of money over time. To make sure that we can also discount in our general financial market $S = (S^0, S^1, \ldots, S^d)$, we subsequently assume that S^0 is a **numéraire** in the sense of the following definition.

Definition 2.4 (Numéraire)**.** Let $\mathcal{T} \subseteq \mathbb{R}$ be a time index set with minimal element $t_0 \in \mathcal{T}$. An $\mathbb{R}$-valued process $X = \{X(t)\}_{t \in \mathcal{T}}$ is said to be a **numéraire** if $X(t_0) = 1$ and $X(t) > 0$ for all $t \in \mathcal{T}$. ◇

With $S^0(0) = 1$ and $S^0(t) > 0$ for all $t = 1, \ldots, T$, we can define the **discounted financial market** $\overline{S} = (\overline{S}^0, \overline{S}^1, \ldots, \overline{S}^d)$ by

$$\overline{S}^i(t) \triangleq \frac{S^i(t)}{S^0(t)} \qquad \text{for all } t = 0, 1, \ldots, T \text{ and } i = 0, 1, \ldots, d.$$

More generally, if $X = \{X(t)\}_{t=0,1,\ldots,T}$ is an $\mathbb{R}$-valued process, we subsequently write $\overline{X} = \{\overline{X}(t)\}_{t=0,1,\ldots,T}$ for the associated discounted process

$$\overline{X}(t) \triangleq \frac{X(t)}{S^0(t)}, \qquad t = 0, 1, \ldots, T.$$

2.2 Filtrations and the Dynamic Resolution of Information

Initially, at time $t = 0$, the prices of securities at later points in time appear random to us. On the other hand, at any time $t > 0$, we have already observed the prices $S(0), S(1), \ldots, S(t)$, i.e. we can tell exactly the outcome of these random variables. In order to formalize this mathematically, we have to model the **dynamic resolution of information** over time.

What does it mean to know the outcome of a random variable? It means that for every event involving the random variable, we can decide whether the event has occurred or not. In other words: Knowing the random variable means knowing the σ-field generated by this random variable. In the above example, observing the prices means that at any time $t = 0, 1, \dots, T$ we know at least all events contained in the σ-field $\sigma(S(0), S(1), \dots, S(t))$. As time evolves, we gather more and more information, i.e. the σ-field modeling the information available to us increases with time. Thus the flow of information can be modeled as an increasing family of σ-fields, indexed by the time index set $\{0, 1, \dots, T\}$. Such a family is called a **filtration**.

Definition 2.5 (Filtration; Filtered Probability Space)**.** Let $\mathcal{T} \subseteq \mathbb{R}$ be a time index set and $\mathfrak{F} = \{\mathfrak{F}(t)\}_{t \in \mathcal{T}}$ be a family of sub-σ-fields of $\mathfrak{A}$ satisfying

$$\mathfrak{F}(s) \subseteq \mathfrak{F}(t) \qquad \text{for all } s, t \in \mathcal{T} \text{ with } s \leq t.$$

Then $\mathfrak{F}$ is called a **filtration** of $(\Omega, \mathfrak{A})$ and the quadruple $(\Omega, \mathfrak{A}, \mathfrak{F}, \mathbb{P})$ is called a **filtered probability space**. ◇

A filtration $\mathfrak{F} = \{\mathfrak{F}(t)\}_{t=0,1,\dots,T}$ represents the **flow of information** and has (for now) nothing to do with any stochastic process. We subsequently fix such a filtration and assume that

$$\mathfrak{F}(0) = \{\emptyset, \Omega\} \qquad \text{and} \qquad \mathfrak{F}(T) = \mathfrak{A}.$$

In other words, we assume that we have no information initially at time $t = 0$ whereas all uncertainty is resolved by time T.

Exercise 10 (Random Variables on Trivial σ-Fields)**.** Let $(\Omega, \mathfrak{A}, \mathbb{P})$ be an arbitrary (possibly infinite) probability space and let $\mathfrak{G} \subseteq \mathfrak{A}$ be $\mathbb{P}$-trivial, i.e. $\mathbb{P}[G] \in \{0, 1\}$ for all $G \in \mathfrak{G}$. Show that for any real-valued $\mathfrak{G}$-measurable random variable X, there exists $\alpha \in \mathbb{R}$ such that $X = \alpha$ almost surely. ◇

Recall that knowing a random variable means knowing the σ-field it generates. Hence knowing a stochastic process $X = \{X(t)\}_{t=0,1,\dots,T}$ means knowing $X(t)$ at any time t. Since our flow of information is modeled by the filtration $\mathfrak{F} = \{\mathfrak{F}(t)\}_{t=0,1,\ldots,T}$, this is to say that $\sigma(X(t))$ is a subset of $\mathfrak{F}(t)$, i.e. $X(t)$ is $\mathfrak{F}(t)$-measurable for all $t = 0, 1, \ldots, T$.

Definition 2.6 (Adaptedness)**.** Let $\mathfrak{F} = \{\mathfrak{F}(t)\}_{t\in\mathcal{T}}$ be a filtration. We say that a stochastic process $X = \{X(t)\}_{t\in\mathcal{T}}$ is **adapted** to the filtration $\mathfrak{F}$ if

$$X(t) \text{ is } \mathfrak{F}(t)\text{-measurable for all } t \in \mathcal{T}.$$

◇

Exercise 11 (σ-Fields on Finite Probability Spaces)**.** Show that any σ-field $\mathfrak{G} \subseteq \mathfrak{A}$ on a finite probability space $(\Omega, \mathfrak{A}, \mathbb{P})$ is generated by a unique partition $\mathfrak{P}$ of Ω. ◇

Exercise 12 (Measurability on Finite Probability Spaces)**.** Let X be a random variable defined on a finite probability space $(\Omega, \mathfrak{A}, \mathbb{P})$. Let moreover $\mathfrak{G}$ be a sub-σ-field of $\mathfrak{A}$ generated by a partition $\mathfrak{P}$. Show that X is $\mathfrak{G}$-measurable if and only if X is constant on each set $P \in \mathfrak{P}$. ◇

Since we can observe prices of securities in our financial market, we subsequently assume that the $\mathbb{R}^{d+1}_+$-valued stochastic process S is adapted to the filtration $\mathfrak{F}$. With this, the modeling of the financial market is now complete. To wrap things up, we gather everything in the following definition.

Definition 2.7 (Finite Financial Market)**.** A **finite financial market** is an $\mathbb{R}^{d+1}_+$-valued and $\mathfrak{F}$-adapted stochastic process

$$S = (S^0, S^1, \ldots, S^d) = \{(S^0(t), S^1(t), \ldots, S^d(t))\}_{t=0,1,\ldots,T}$$

provided that the security S^0 is a numéraire. ◇

At this point the question arises how to choose the filtration in our model. A canonical choice is to first fix the price process $S = \{S(t)\}_{t=0,1,\ldots,T}$ and consider the minimal filtration which renders the process adapted.

Definition 2.8 (Natural Filtration)**.** Let $X = \{X(t)\}_{t\in\mathcal{T}}$ be a stochastic process. The filtration $\mathfrak{F}^X = \{\mathfrak{F}^X(t)\}_{t\in\mathcal{T}}$ given by

$$\mathfrak{F}^X(t) \triangleq \sigma(X(s) : s \in \mathcal{T},\, s \leq t), \qquad t \in \mathcal{T},$$

is called the **natural filtration** of X.

Obviously, each process is adapted to its natural filtration. If we are in the special case of $\mathfrak{F} = \mathfrak{F}^S$, it is immediately seen that $\mathfrak{F}^S(0) = \{\emptyset, \Omega\}$ if $S(0)$ is constant and $\mathfrak{F}^S(T) = \mathfrak{A}$ if $S(\cdot, \omega) \neq S(\cdot, \tilde{\omega})$ for any distinct $\omega, \tilde{\omega} \in \Omega$.

Exercise 13 (Natural Filtration in the CRR Model)**.** Consider the CRR model (S^0, S^1) introduced in Example 2.2. Show that $\mathfrak{F}^S(0) = \{\emptyset, \Omega\}$ and

$$\mathfrak{F}^S(t) = \sigma(R_1, \dots, R_t), \qquad t = 1, \dots, T.$$

What does the condition $\mathfrak{F}^S(T) = \mathfrak{A}$ tell us about the cardinality of Ω? ◇

2.3 Trading Strategies and the Wealth Process

We turn towards the modeling of **trading strategies**. To get a full description of the trading activities of an investor, we need to keep track of the number of assets S^i held in between the trading dates $0, 1, \dots, T$. This means that trading strategies can be modeled as $\mathbb{R}^{d+1}$-valued processes

$$\varphi = (\varphi_0, \varphi_1, \dots, \varphi_d) = \{(\varphi_0(t), \varphi_1(t), \dots, \varphi_d(t))\}_{t=1,\dots,T}.$$

We should be careful with the time index set here: The time indices of $\varphi = \{\varphi(t)\}_{t=1,\dots,T}$ correspond to **trading periods**, which is in contrast to the time indices in $S = \{S(t)\}_{t=0,1\dots,T}$ which correspond to **trading dates**. More precisely, for each $i = 0, 1, \dots, d$ and $t = 1, \dots, T$,

$\varphi_i(t)$ models the number of securities S^i held between $t-1$ and t.

See also Figure 2.4 for an illustration.

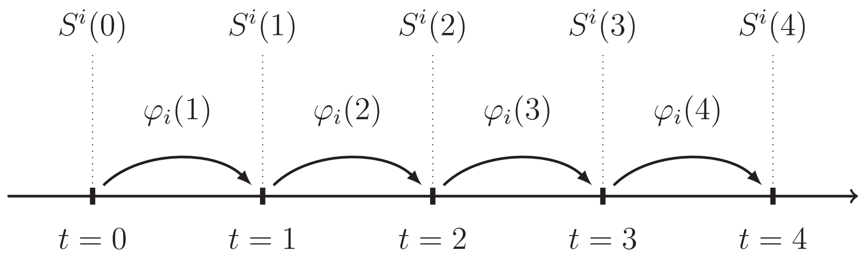

Figure 2.4. Trading periods and predictability of trading strategies.

We furthermore observe that the investor must decide at the beginning of the trading period how many shares to hold, so the choice of $\varphi_i(t)$ must be based on the information available at time $t-1$, i.e. $\varphi_i(t)$ should be $\mathfrak{F}(t-1)$-measurable. Such processes are called **predictable**.

Definition 2.9 (Predictable Process with Finite Time Index Set)**.** Given a filtration $\mathfrak{F} = \{\mathfrak{F}(t)\}_{t=0,1,\ldots,T}$, a stochastic process $X = \{X(t)\}_{t=1,\ldots,T}$ is called **predictable** if

$$X(t) \text{ is } \mathfrak{F}(t-1)\text{-measurable for all } t = 1, \ldots, T. \qquad \diamond$$

Given a predictable $\mathbb{R}^{d+1}$-valued trading strategy φ as above, the wealth before trading at any time $t = 1, \ldots, T-1$ is given by the scalar product

$$\langle \varphi(t) | S(t) \rangle = \sum_{i=0}^{d} \varphi_i(t) S^i(t),$$

whereas the wealth after trading at any time t is

$$\langle \varphi(t+1) | S(t) \rangle = \sum_{i=0}^{d} \varphi_i(t+1) S^i(t).$$

If these two quantities coincide, the trading strategy φ can be financed without any external cash in- or outflows, i.e. the strategy finances itself.

Definition 2.10 (Self-Financing Trading Strategy; Wealth Process)**.** An $\mathbb{R}^{d+1}$-valued predictable process $\varphi = \{\varphi(t)\}_{t=1,\ldots,T}$ is called **self-financing trading**

strategy with respect to the financial market S if

$$\langle \varphi(t) | S(t) \rangle = \langle \varphi(t+1) | S(t) \rangle, \qquad t = 1, \ldots, T-1. \tag{2.1}$$

With this, the $\mathbb{R}$-valued process $X^\varphi = \{X^\varphi(t)\}_{t=0,1,\ldots,T}$ given by

$$X^\varphi(t) \triangleq \langle \varphi(t) | S(t) \rangle, \qquad t = 1, \ldots, T, \qquad X^\varphi(0) \triangleq \langle \varphi(1) | S(0) \rangle,$$

is called the **wealth process** generated by φ. ◇

The advantage of restricting to self-financing trading strategies is of course that the wealth process is unambiguously defined at any trading date and we do not have to distinguish between wealth before and after trades.

We note that any $\mathbb{R}^{d+1}$-valued predictable process φ is a self-financing trading strategy with respect to S if and only if it is a self-financing trading strategy with respect to the discounted financial market $\overline{S}$. This follows immediately upon dividing Equation (2.1) by $S^0(t)$.

Given a self-financing trading strategy φ and $t = 1, \ldots, T$, the profits and losses between $t-1$ and t are

$$\sum_{i=0}^{d} \varphi_i(t) \big[S^i(t) - S^i(t-1) \big] = \langle \varphi(t) | \Delta S(t) \rangle$$

where we define

$$\Delta S(t) \triangleq S(t) - S(t-1).$$

With this, the overall profits and losses made up to (and including) time t can be obtained by summing the above over all time indices. To keep the presentation concise, we introduce a special notation for this.

Definition 2.11 (Discrete Stochastic Integral)**.** Given an $\mathbb{R}^{d+1}$-valued predictable process $\varphi = \{\varphi(t)\}_{t=1,\ldots,T}$ and an $\mathbb{R}^{d+1}$-valued adapted process $S = \{S(t)\}_{t=0,1,\ldots,T}$, we define the (discrete) **stochastic integral** of φ with respect to S as the $\mathbb{R}$-valued adapted process $\varphi \bullet S = \{\varphi \bullet S(t)\}_{t=0,1,\ldots,T}$ with[1]

$$\varphi \bullet S(t) \triangleq \sum_{\tau=1}^{t} \langle \varphi(\tau) | \Delta S(\tau) \rangle, \qquad t = 0, 1, \ldots, T.$$ ◇

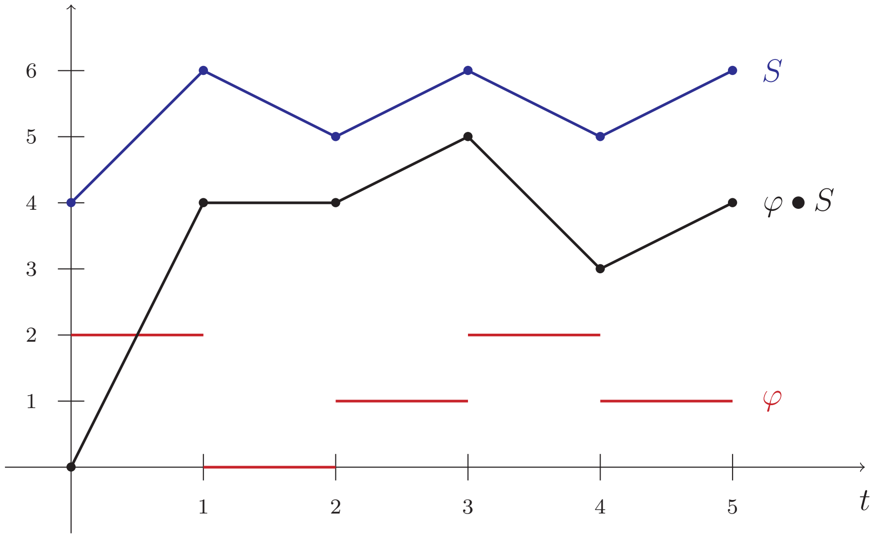

Figure 2.5. A trajectory of the stochastic integral $\varphi \bullet S$.

The stochastic integral $\varphi \bullet S(t)$ describes the cumulative profits and losses made up to and including time t. Not surprisingly, it is possible to characterize self-financing trading strategies in terms of this stochastic integral: Self-financing trading strategies are exactly the ones for which the wealth $X^{\varphi}(t)$ at time t corresponds to the initial wealth $X^{\varphi}(0)$ plus cumulative profits and losses from trading. The following result makes this precise.

Lemma 2.12 (Characterization of Self-Financing Strategies)**.** *For any $\mathbb{R}^{d+1}$-valued predictable process $\varphi = \{\varphi(t)\}_{t=1,\ldots,T}$, the following statements are equivalent:*

(i) φ *is a self-financing trading strategy with respect to* S.

(ii) It holds that $X^{\varphi}(t) = X^{\varphi}(0) + \varphi \bullet S(t)$ *for all* $t = 0, 1, \ldots, T$.

(iii) φ *is a self-financing trading strategy with respect to* $\overline{S}$.

(iv) It holds that $\overline{X}^{\varphi}(t) = X^{\varphi}(0) + \varphi \bullet \overline{S}(t)$ *for all* $t = 0, 1, \ldots, T$. ◇

Proof. We already know that (i) and (iii) are equivalent. Since S is a fi-

[1]We use the convention that $\varphi \bullet S(0) = \sum_{\tau=1}^{0} \langle \varphi(\tau) | \Delta S(\tau) \rangle = 0$.

nancial market if and only if $\overline{S}$ is a financial market and $X^{\varphi}(0) = \overline{X}^{\varphi}(0)$, it follows that (iii) and (iv) are equivalent if and only if (i) and (ii) are equivalent. Thus, we only have to prove the equivalence of (i) and (ii).

For this, let us fix $t \in \{0, 1, \ldots, T\}$ and observe that we may write

$$\begin{aligned}
X^{\varphi}(0) + \varphi \bullet S(t) &= \langle \varphi(1) | S(0) \rangle + \sum_{\tau=1}^{t} \langle \varphi(\tau) | \Delta S(\tau) \rangle \\
&= \sum_{\tau=1}^{t} \langle \varphi(\tau) | S(\tau) \rangle - \sum_{\tau=2}^{t} \langle \varphi(\tau) | S(\tau - 1) \rangle \\
&= \langle \varphi(t) | S(t) \rangle + \sum_{\tau=1}^{t-1} \langle \varphi(\tau) | S(\tau) \rangle - \sum_{\tau=1}^{t-1} \langle \varphi(\tau + 1) | S(\tau) \rangle \\
&= X^{\varphi}(t) + \sum_{\tau=1}^{t-1} \big[\langle \varphi(\tau) | S(\tau) \rangle - \langle \varphi(\tau + 1) | S(\tau) \rangle \big].
\end{aligned}$$

Thus, under (i), i.e. if φ is self-financing, the last sum equals zero and hence

$$X^{\varphi}(0) + \varphi \bullet S(t) = X^{\varphi}(t)$$

which is (ii). On the other hand, if (ii) holds, then the left hand side is equal to $X^{\varphi}(t)$ and hence we find that

$$0 = \sum_{\tau=1}^{t-1} \big[\langle \varphi(\tau) | S(\tau) \rangle - \langle \varphi(\tau + 1) | S(\tau) \rangle \big] \quad \text{for all } t = 2, \ldots, T.$$

Iteratively writing down this equation starting from $t = 2$, we find that this implies

$$\langle \varphi(t) | S(t) \rangle = \langle \varphi(t + 1) | S(t) \rangle \qquad \text{for all } t = 1, \ldots, T - 1,$$

i.e. φ is self-financing. □

Given a predictable stochastic process φ with values in $\mathbb{R}^{d+1}$, it can be quite cumbersome to check the self-financing condition

$$\langle \varphi(t) | S(t) \rangle = \langle \varphi(t + 1) | S(t) \rangle, \qquad t = 1, \ldots, T - 1,$$

since we have to do so for all $\omega \in \Omega$. Thus, it might be a good idea to come up with a simpler method of constructing self-financing trading strategies.

One possibility is to let the investor choose $\varphi_1, \dots, \varphi_d$ freely and then determine φ_0 in a way to turn $\varphi = (\varphi_0, \varphi_1, \dots, \varphi_d)$ into a self-financing trading strategy. As it turns out, φ_0 is determined uniquely if we additionally fix the initial wealth $X^\varphi(0)$.

Lemma 2.13 (Construction of Self-Financing Strategies). *Let $x \in \mathbb{R}$ and $(\varphi_1, \dots, \varphi_d) = \{(\varphi_1(t), \dots, \varphi_d(t))\}_{t=1,\dots,T}$ be an $\mathbb{R}^d$-valued predictable process. Then $\varphi_0 = \{\varphi_0(t)\}_{t=1,\dots,T}$ given by*

$$\varphi_0(t) = x + \sum_{\tau=1}^{t-1} \sum_{i=1}^{d} \varphi_i(\tau) \Delta \overline{S}^i(\tau) - \sum_{i=1}^{d} \varphi_i(t) \overline{S}^i(t-1), \qquad t = 1, \dots, T,$$

is the unique $\mathbb{R}$-valued predictable process such that $\varphi = (\varphi_0, \varphi_1, \dots, \varphi_d)$ is a self-financing trading strategy with initial wealth $X^\varphi(0) = x$. ◇

Proof. We first note that φ_0 defined as above is indeed predictable since the first sum only involves terms up to time $t-1$ and each $\varphi_i(t)$, $i = 1, \dots, d$, in the last sum is $\mathfrak{F}(t-1)$-measurable by predictability of each φ_i.

If φ_0 is an arbitrary predictable process, Lemma 2.12 (Characterization of Self-Financing Strategies) tells us that $\varphi = (\varphi_0, \varphi_1, \dots, \varphi_d)$ is a self-financing strategy with $X^\varphi(0) = x$ if and only if

$$\langle \varphi(t) | \overline{S}(t) \rangle = \overline{X}^\varphi(t) = x + \varphi \bullet \overline{S}(t) = x + \sum_{\tau=1}^{t} \sum_{i=1}^{d} \varphi_i(\tau) \Delta \overline{S}^i(\tau)$$

for all $t = 0, 1, \dots, T$. Here we have used $\overline{S}^0(\tau) = 1$ and hence $\Delta \overline{S}^0(\tau) = 0$ for all $\tau = 1, \dots, t$ to drop the index $i = 0$ in the last equality. We conclude since rearranging terms shows that

$$\begin{aligned} \varphi_0(t) = \varphi_0(t) \overline{S}^0(t) &= x + \sum_{\tau=1}^{t} \sum_{i=1}^{d} \varphi_i(\tau) \Delta \overline{S}^i(\tau) - \sum_{i=1}^{d} \varphi_i(t) \overline{S}^i(t) \\ &= x + \sum_{\tau=1}^{t-1} \sum_{i=1}^{d} \varphi_i(\tau) \Delta \overline{S}^i(\tau) - \sum_{i=1}^{d} \varphi_i(t) \overline{S}^i(t-1). \end{aligned}$$ □

The previous result allows us to give the following alternative definition of (self-financing) trading strategies.

Definition 2.14 (Trading Strategies with Initial Wealth)**.** Let (x, φ) be a pair consisting of a real number $x \in \mathbb{R}$ and an $\mathbb{R}^d$-valued predictable stochastic process $\varphi = (\varphi_1, \ldots, \varphi_d) = \{(\varphi_1(t), \ldots, \varphi_d(t))\}_{t=1,\ldots,T}$. Then we say that φ is a **trading strategy with initial wealth** x. ◇

Each trading strategy φ with initial wealth x in the sense of Definition 2.14 induces a unique self-financing trading strategy $\psi = (\varphi_0, \varphi)$ in the sense of Definition 2.10 if we define φ_0 as in Lemma 2.13. In light of this, we define the wealth process corresponding to φ by $X^\varphi \triangleq X^\psi$ and subsequently switch between these two definitions without further mention whenever we see fit, silently agreeing that we take φ_0 as in Lemma 2.13.

With this, the modeling of finite financial markets and trading strategies is now complete and we can turn towards the analysis of arbitrage opportunities and the pricing of options in such markets.

Chapter 3 Risk Neutral Pricing and Fundamental Theorems

Now that we have set up a financial market model, the first thing to do is to derive conditions under which this market is free of arbitrage opportunities. Once this is achieved, we introduce options into the financial market and show how to compute with fair prices for them.

Throughout the entire chapter, we fix a finite financial market model S in the sense of Definition 2.7 and recall that the associated discounted financial market is denoted by $\overline{S}$.

3.1 Arbitrage and the First Fundamental Theorem

Our first task is to analyze under which conditions the market S is consistent with the absence of arbitrage assumption. Since we work within a specific model, we can give a more rigorous definition of arbitrage opportunities.

Definition 3.1 (Arbitrage Opportunities in Finite Markets)**.** A self-financing trading strategy φ is called **arbitrage opportunity** if

$$X^\varphi(0) = 0, \qquad X^\varphi(T) \geq 0, \qquad \text{and} \qquad \mathbb{P}[X^\varphi(T) > 0] > 0. \qquad \diamond$$

Note that $X^\varphi(T) \geq 0$ means $X^\varphi(T,\omega) \geq 0$ for all $\omega \in \Omega$ and that we have $\mathbb{P}[X^\varphi(T) > 0] > 0$ as soon as there exists one $\omega \in \Omega$ with $X^\varphi(T,\omega) > 0$. An arbitrage opportunity is hence a trading strategy which starts from zero wealth, leads to positive wealth at time T, and there is at least one state of the world in which the wealth is even strictly positive.

In light of Lemma 2.12 (Characterization of Self-Financing Strategies), we can equivalently define arbitrage opportunities as $\mathbb{R}^d$-valued predictable processes φ satisfying[1]

$$\varphi \bullet \overline{S}(T) \geq 0 \qquad \text{and} \qquad \mathbb{P}[\varphi \bullet \overline{S}(T) > 0] > 0.$$

This representation highlights the close connection between mathematical finance and stochastic analysis, as it states that absence of arbitrage is equivalent to the statement that if a stochastic integral takes on a strictly positive value in one state of the world, there must be another state of the world in which in which the integral takes on a strictly negative value, i.e. we have a one-to-one relation between a financial concept and a statement on a property of stochastic integrals.

To get a better feeling for this and to build up some intuition, let us consider the CRR model $S = (S^0, S^1)$ as discussed in Example 2.2, where we take the underlying filtration to be the natural filtration of S, i.e. $\mathfrak{F} \triangleq \mathfrak{F}^S$. Intuitively, this model admits arbitrage if either $r \leq d < u$ since in this case S^1 outperforms S^0 in any state of the world (we can make arbitrage by buying S^1 and selling S^0 for equal amounts of money) or if $r \geq u > d$ since then S^0 outperforms S^1 (we can make arbitrage by selling S^1 and buying S^0). On the other hand, if $d < r < u$, there is always one scenario in each trading period in which S^1 outperforms S^0 and one scenario in which S^0 outperforms S^1, so it should not be possible to generate arbitrage as trading decisions must be based on the information available before the price movement. Note that the condition $d < r < u$ is equivalent to requiring the discounted price process $\overline{S}^1$ to both strictly rise and strictly fall with strictly positive probability in each trading period as showcased in Figure 3.1 below.

In particular, if we replace the probabilities p and $1-p$ for upward and downward movements on each node by suitable new values, it should be possible to ensure that, conditional on $S^1(0), S^1(1), \ldots, S^1(t-1)$, the increment $\Delta\overline{S}^1(t)$ remains constant in expectation under these new probabilities. Comparing with Figure 3.2, this is possible if and only if $d < r < u$.

[1] Observe the slight abuse of notation here since $\overline{S}$ is $(d+1)$-dimensional. This is justified since $\Delta\overline{S}^0 \equiv 0$.

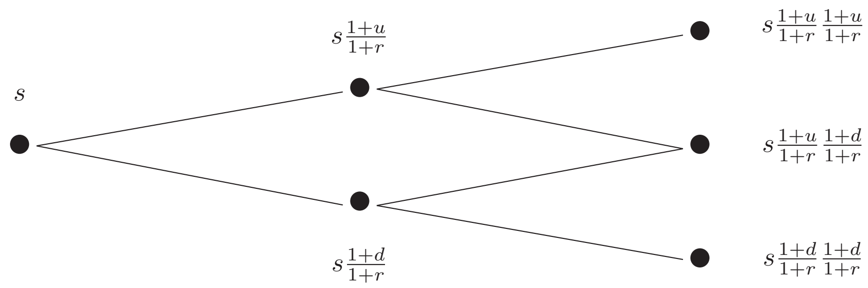

Figure 3.1. The discounted price process $\overline{S}^1$ in the CRR model if $d < r < u$.

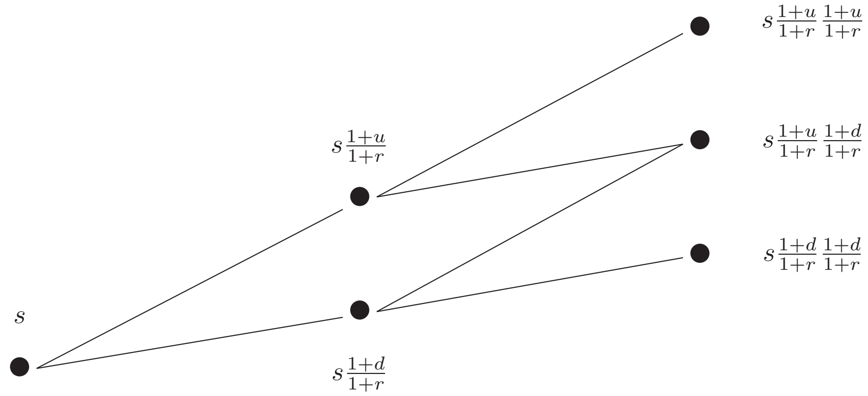

Figure 3.2. The discounted price process $\overline{S}^1$ in the CRR model if $r < d < u$.

The bottom line of this argument is that we expect the CRR model to be free of arbitrage if and only if we can change the probability measure $\mathbb{P}$ to a new probability measure $\mathbb{Q}$ such that the increments $\Delta\overline{S}(t)$ are zero in $\mathfrak{F}(t-1)$-conditional expectation under the new measure $\mathbb{Q}$. As it turns out, this statement is true even for the general market model S, and the aim of this section is to prove exactly this equivalence, which we refer to as the **First Fundamental Theorem of Asset Pricing** or **1st FTAP** for short.

Let us first translate the rather vague formulation of the 1st FTAP given above into rigorous mathematics. For this, let us recall that two probability measures $\mathbb{P}$ and $\mathbb{Q}$ defined on the same probability space are called **equivalent**, written $\mathbb{Q} \sim \mathbb{P}$ for brevity, if the nullsets of these two measures coincide, i.e. for all $A \in \mathfrak{A}$ it holds that

$$\mathbb{P}[A] = 0 \qquad \text{if and only if} \qquad \mathbb{Q}[A] = 0.$$

Under our finiteness assumption on the probability space $(\Omega, \mathfrak{A}, \mathbb{P})$, the above statement simply means that a probability measure $\mathbb{Q}$ is equivalent to $\mathbb{P}$ if and only if

$$\mathbb{Q}[\{\omega\}] > 0 \qquad \text{for all } \omega \in \Omega.$$

This is to say that under the new measure $\mathbb{Q}$, it remains true that there are no impossible elementary events.

Under the measure $\mathbb{P}$, for any $t = 0, 1, \ldots, T$, we write $\mathbb{E}_t \triangleq \mathbb{E}[\,\cdot\,|\mathfrak{F}(t)]$ for the conditional expectation given $\mathfrak{F}(t)$. Note that $\mathbb{E} = \mathbb{E}_0$ since $\mathfrak{F}(0)$ is trivial. Moreover, we write $\mathbb{E}^{\mathbb{Q}}$ for the unconditional expectation under $\mathbb{Q}$ and we set $\mathbb{E}_t^{\mathbb{Q}} \triangleq \mathbb{E}^{\mathbb{Q}}[\,\cdot\,|\mathfrak{F}(t)]$ for the $\mathfrak{F}(t)$-conditional expectation under $\mathbb{Q}$. We can now begin formalizing the statement of the 1^{st} FTAP.

Definition 3.2 ((Sub-/Super-) Martingales on Finite Probability Spaces)**.** Let $Y = \{Y(t)\}_{t=0,1,\ldots,T}$ be an adapted real-valued process defined on a finite filtered probability space $(\Omega, \mathfrak{A}, \mathfrak{F}, \mathbb{P})$. We refer to Y

(i)	as a **supermartingale**	if	$\mathbb{E}_{t-1}[\Delta Y(t)] \leq 0,$	$t = 1, \ldots, T;$
(ii)	as a **submartingale**	if	$\mathbb{E}_{t-1}[\Delta Y(t)] \geq 0,$	$t = 1, \ldots, T;$
(iii)	as a **martingale**	if	$\mathbb{E}_{t-1}[\Delta Y(t)] = 0,$	$t = 1, \ldots, T.$ ◇

Adaptedness of Y means that $Y(t-1)$ is $\mathfrak{F}(t-1)$-measurable and hence the martingale condition rewrites equivalently as

$$\mathbb{E}_{t-1}[Y(t)] = Y(t-1), \qquad t = 1, \ldots, T.$$

Moreover, whenever $s, t \in \{0, 1, \ldots, T\}$ with $s < t$, we note that the tower property of conditional expectation and the martingale property imply that

$$\mathbb{E}_s[Y(t)] = \mathbb{E}_s\big[\mathbb{E}_{t-1}[Y(t)]\big] = \mathbb{E}_s[Y(t-1)] = \ldots = \mathbb{E}_s[Y(s+1)] = Y(s).$$

Taking expectations on both sides, it furthermore follows that

$$\mathbb{E}[Y(t)] = \mathbb{E}[Y(s)] \quad \text{and hence also} \quad \mathbb{E}[Y(t)] = \mathbb{E}[Y(0)] = Y(0) \qquad (3.1)$$

by choosing $s = 0$ and since $Y(0)$ is $\mathfrak{F}(0)$-measurable and hence constant. Analogous statements are valid if Y is a supermartingale or a submartingale.

In the statement of the 1^{st} FTAP, the discounted market $\overline{S}$ does not satisfy the martingale property under the original measure $\mathbb{P}$, but under the equivalent measure $\mathbb{Q} \sim \mathbb{P}$. We reserve a special name for this kind of measure.

Definition 3.3 (Equivalent Martingale Measure). A probability measure $\mathbb{Q}$ on $(\Omega, \mathfrak{A})$ is called **equivalent martingale measure** or **EMM** for S if $\mathbb{Q} \sim \mathbb{P}$ and the discounted market $\overline{S}$ is a martingale under $\mathbb{Q}$, i.e.

$$\mathbb{E}_{t-1}^{\mathbb{Q}}[\overline{S}(t)] = \overline{S}(t-1), \qquad t = 1, \ldots, T. \qquad \diamond$$

With this definition in place, the rather vague statement of the 1^{st} FTAP can now be made precise: We expect that the finite financial market S is free of arbitrage opportunities if and only if there exists an EMM $\mathbb{Q}$. Before we can prove this statement, however, we first need an alternative characterization of EMMs.

Lemma 3.4 (Characterization of EMMs). *A probability measure $\mathbb{Q} \sim \mathbb{P}$ is an EMM for S if and only if*

$$\mathbb{E}^{\mathbb{Q}}[\overline{X}^{\varphi}(T)] = 0 \text{ for any trading strategy } \varphi \text{ with initial wealth zero.} \tag{3.2}$$

In this case, $\overline{X}^{\varphi}$ is a $\mathbb{Q}$-martingale for any self-financing trading strategy φ (i.e. with arbitrary initial wealth). $\diamond$

Proof. Suppose that $\mathbb{Q}$ is an EMM and let φ be an arbitrary self-financing trading strategy. For any $t = 1, \ldots, T$, using the definition and adaptedness of the discrete stochastic integral followed by the predictability of φ, we find

$$\begin{aligned}
\mathbb{E}_{t-1}^{\mathbb{Q}}[\varphi \bullet \overline{S}(t)] &= \mathbb{E}_{t-1}^{\mathbb{Q}}[\varphi \bullet \overline{S}(t-1) + \langle \varphi(t) | \Delta \overline{S}(t) \rangle] \\
&= \varphi \bullet \overline{S}(t-1) + \mathbb{E}_{t-1}^{\mathbb{Q}}\Big[\sum_{i=0}^{d} \varphi_i(t) \Delta \overline{S}^i(t)\Big] \\
&= \varphi \bullet \overline{S}(t-1) + \sum_{i=0}^{d} \varphi_i(t) \mathbb{E}_{t-1}^{\mathbb{Q}}[\Delta \overline{S}^i(t)] = \varphi \bullet \overline{S}(t-1),
\end{aligned}$$

where we have used the $\mathbb{Q}$-martingale property of $\overline{S}$ for the last equality. We have therefore argued that $\varphi \bullet \overline{S}$ is a $\mathbb{Q}$-martingale whenever $\overline{S}$ is a

$\mathbb{Q}$-martingale. Since adding the constant value $X^{\varphi}(0)$ does not affect the martingale property, it follows that in this case also $\overline{X}^{\varphi} = X^{\varphi}(0) + \varphi \bullet \overline{S}$ is a $\mathbb{Q}$-martingale. From Equation (3.1), we conclude that

$$\mathbb{E}[\overline{X}^{\varphi}(T)] = 0 \qquad \text{if } \varphi \text{ is a trading strategy with initial wealth zero.}$$

For the converse direction, we fix $t \in \{1, \dots, T\}$, $i \in \{1, \dots, d\}$, and set

$$F \triangleq \left\{\mathbb{E}^{\mathbb{Q}}_{t-1}[\Delta\overline{S}^{i}(t)] > 0\right\} \in \mathfrak{F}(t-1).$$

With this, we define a trading strategy φ with initial wealth zero by setting $\varphi_i(t) = \mathbb{1}_F$ and

$$\varphi_j(s) = 0 \qquad \text{if either } j \in \{1, \dots, d\} \setminus \{i\} \text{ or } s \in \{1, \dots, T\} \setminus \{t\},$$

i.e. φ is the strategy which holds 1 unit of S^i between $t-1$ and t on the event F. Thus, it follows that

$$\overline{X}^{\varphi}(T) = \varphi_i(t)\Delta\overline{S}^{i}(t) = \mathbb{1}_F\Delta\overline{S}^{i}(t),$$

and upon assuming that Equation (3.2) is valid, it follows that

$$\mathbb{E}^{\mathbb{Q}}\left[\mathbb{1}_{\{\mathbb{E}^{\mathbb{Q}}_{t-1}[\Delta\overline{S}^{i}(t)]>0\}}\mathbb{E}^{\mathbb{Q}}_{t-1}[\Delta\overline{S}^{i}(t)]\right] = \mathbb{E}^{\mathbb{Q}}[\mathbb{1}_F\Delta\overline{S}^{i}(t)] = \mathbb{E}^{\mathbb{Q}}[\overline{X}^{\varphi}(T)] = 0.$$

But this is only possible if F has probability zero, i.e. $\mathbb{E}^{\mathbb{Q}}_{t-1}[\Delta\overline{S}^{i}(t)] \leq 0$. Repeating the same argument with $F \triangleq \{\mathbb{E}^{\mathbb{Q}}_{t-1}[\Delta\overline{S}^{i}(t)] < 0\}$ entails that

$$\mathbb{E}^{\mathbb{Q}}_{t-1}[\Delta\overline{S}^{i}(t)] = 0, \qquad t = 1, \dots, T,\ i = 1, \dots, d,$$

i.e. $\overline{S}$ is a $\mathbb{Q}$-martingale and hence $\mathbb{Q}$ an EMM. □

Let us highlight once more that in the previous proof we have in particular argued that if $\overline{S}$ is a $\mathbb{Q}$-martingale, so is any discrete stochastic integral $\varphi \bullet \overline{S}$, i.e. the discrete stochastic integral preserves the martingale property of the integrator, a property which will become quite significant later.

Let us now come to the first fundamental theorem of asset pricing.

Theorem 3.5 (First Fundamental Theorem of Asset Pricing)**.** *In a finite financial market S, the following statements are equivalent:*

(i) There are no arbitrage opportunities in S.

(ii) There exists at least one EMM $\mathbb{Q}$. ◇

Proof. Step 1. We show that (ii) implies (i). For this, let us assume by contradiction that there exists an arbitrage opportunity φ. As $\mathbb{Q} \sim \mathbb{P}$ and S^0 is a numéraire, it follows that

$$\overline{X}^{\varphi}(T) \geq 0 \qquad \text{and} \qquad \mathbb{Q}[\overline{X}^{\varphi}(T) > 0] > 0, \qquad \text{i.e.} \qquad \mathbb{E}^{\mathbb{Q}}[\overline{X}^{\varphi}(T)] > 0.$$

But this is a contradiction to Lemma 3.4 (Characterization of EMMs).

Step 2. We are left with showing that (i) implies (ii). Since $\Omega = \{\omega_1, \ldots, \omega_N\}$ is finite, we can identify the space L of random variables on Ω with the Euclidean space $\mathbb{R}^N$ by identifying each random variable X with the vector $(X(\omega_1), \ldots, X(\omega_N))$. The scalar product on L induced by the scalar product on $\mathbb{R}^N$ is given by

$$\langle X|Y\rangle = \sum_{\omega\in\Omega} X(\omega)Y(\omega), \qquad X, Y \in \mathrm{L}.$$

We now separate the set of attainable payoffs from the set of arbitrage payoffs by a hyperplane as sketched in Figure 3.3.

For this, we define a compact and convex subset of L by

$$\mathrm{K} \triangleq \Big\{\textstyle\sum_{\omega\in\Omega}\lambda_\omega \mathbb{1}_{\{\omega\}} \in \mathrm{L} : \lambda_\omega \geq 0,\ \omega \in \Omega \text{ and } \sum_{\omega\in\Omega}\lambda_\omega = 1\Big\},$$

and denote by P_0 the linear space of discounted payoffs attainable from zero initial wealth given by

$$\mathrm{P}_0 \triangleq \{\overline{X}^{\varphi}(T) \in \mathrm{L} : \varphi \text{ trading strategy with initial wealth zero}\}.$$

By assumption, there are no arbitrage opportunities in S, and hence it follows that

$$\mathrm{P}_0 \cap \mathrm{K} = \emptyset.$$

Since P_0 is linear and closed and K is convex and compact, the hyperplane separation theorem yields the existence of a random variable $0 \neq Z \in \mathrm{L}$ and a constant $\beta > 0$ such that

$$\langle Z|X\rangle = 0 < \beta \leq \langle Z|Y\rangle \qquad \text{for all } X \in \mathrm{P}_0 \text{ and } Y \in \mathrm{K}. \tag{3.3}$$

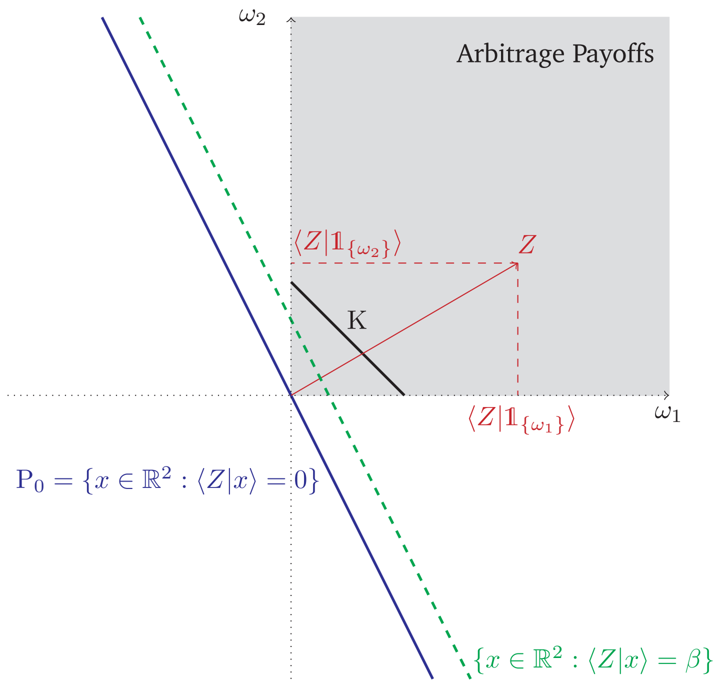

Figure 3.3. Sketch of the separation argument in the proof of the 1^{st} FTAP.

Since $\mathbb{1}_{\{\omega\}} \in \mathrm{K}$, we see in particular that

$$Z(\omega) = \langle Z|\mathbb{1}_{\{\omega\}}\rangle \geq \beta > 0, \qquad \omega \in \Omega.$$

But then we can define a probability measure $\mathbb{Q} \sim \mathbb{P}$ by setting

$$\mathbb{Q}[\{\omega\}] \triangleq \frac{1}{\gamma} Z(\omega) > 0, \qquad \omega \in \Omega, \qquad \text{where } \gamma \triangleq \sum_{\omega \in \Omega} Z(\omega) > 0.$$

Now let φ be an arbitrary trading strategy with initial wealth zero. Then $\overline{X}^{\varphi}(T) \in \mathrm{P}_0$ and hence by Equation (3.3)

$$\mathbb{E}^{\mathbb{Q}}[\overline{X}^{\varphi}(T)] = \sum_{\omega \in \Omega} \overline{X}^{\varphi}(T, \omega)\mathbb{Q}[\{\omega\}] = \frac{1}{\gamma}\langle \overline{X}^{\varphi}(T)|Z\rangle = 0.$$

But then Lemma 3.4 (Characterization of EMMs) implies that $\mathbb{Q}$ is an EMM and the proof is complete. □

To convince ourselves of the usefulness of the 1^{st} FTAP, let us apply it to the CRR model.

Example 3.6 (Arbitrage in the CRR Model)**.** Let $S = (S^0, S^1)$ be the CRR model with $\mathfrak{F} = \mathfrak{F}^S$ the natural filtration of S. Recall that $\mathfrak{F}^S(T) = \mathfrak{A}$ if and only if $S^1(\cdot, \omega) \neq S^1(\cdot, \tilde{\omega})$ for any distinct $\omega, \tilde{\omega} \in \Omega$. This implies that Ω consists of exactly 2^T elements, i.e. each $\omega \in \Omega$ corresponds to exactly one trajectory of S^1. The i.i.d. property of $R_1, \dots, R_T$ then implies that

$$\mathbb{P}[\{\omega\}] = p^{U(\omega)}(1-p)^{D(\omega)}, \qquad \omega \in \Omega,$$

where $U : \Omega \to \{0, 1, \dots, T\}$ and $D : \Omega \to \{0, 1, \dots, T\}$ count the number of up and down movements in each state of the world, i.e.

$$U(\omega) \triangleq \sum_{t=1}^{T} \mathbb{1}_{\{R_t(\omega)=u\}} \qquad \text{and} \qquad D(\omega) \triangleq \sum_{t=1}^{T} \mathbb{1}_{\{R_t(\omega)=d\}}, \qquad \omega \in \Omega.$$

For any $q \in [0, 1]$, we let $\mathbb{Q}$ denote the probability measure on $(\Omega, \mathfrak{A})$ characterized by

$$\mathbb{Q}[R_t = u] = q \qquad \text{and} \qquad \mathbb{Q}[R_t = d] = 1 - q, \qquad t = 1, \dots, T,$$

which is to say that

$$\mathbb{Q}[\{\omega\}] = q^{U(\omega)}(1-q)^{D(\omega)}, \qquad \omega \in \Omega.$$

Observe that, for any sequence $x_1, \dots, x_T$ taking values in $\{u, d\}$, there exists exactly one $\omega \in \Omega$ such that

$$\big(R_1, \dots, R_T\big)(\omega) = (x_1, \dots, x_T)$$

and hence

$$\mathbb{Q}[R_1 = x_1, \dots, R_T = x_T] = \mathbb{Q}[\{\omega\}] = q^{U(\omega)}(1-q)^{D(\omega)} = \prod_{t=1}^{T} \mathbb{Q}[R_t = x_t],$$

i.e. $R_1, \dots, R_T$ are also independent under $\mathbb{Q}$. Now $\mathbb{Q}$ is an EMM if and only if $q \in (0, 1)$ and $\overline{S}^1$ is a $\mathbb{Q}$-martingale, i.e.

$$\begin{aligned} 0 = \mathbb{E}^{\mathbb{Q}}_{t-1}[\Delta \overline{S}^1(t)] &= \mathbb{E}^{\mathbb{Q}}_{t-1}\Big[\frac{1+R_t}{1+r}\,\overline{S}(t-1) - \overline{S}(t-1)\Big] \\ &= \frac{\overline{S}^1(t-1)}{1+r}\mathbb{E}^{\mathbb{Q}}_{t-1}\big[1 + R_t - (1+r)\big], \quad t = 1, \dots, T. \qquad (3.4) \end{aligned}$$

Since R_t is independent of $R_1, \dots, R_{t-1}$, it follows that it is also independent of $\mathfrak{F}(t-1) = \sigma(S(0), S(1), \dots, S(t-1))$ and hence

$$\mathbb{E}^{\mathbb{Q}}_{t-1}[R_t] = \mathbb{E}^{\mathbb{Q}}[R_t] = qu + (1-q)d, \qquad t = 1, \dots, T.$$

But then, using this together with Equation (3.4), $\mathbb{Q}$ is an EMM if and only if $q \in (0,1)$ and

$$0 = \mathbb{E}^{\mathbb{Q}}_{t-1}\big[1 + R_t - (1+r)\big] = qu + (1-q)d - r,$$

which holds if and only if

$$q = \frac{r-d}{u-d} \in (0,1), \qquad \text{i.e.} \qquad d < r < u.$$

In other words, it follows from the first fundamental theorem of asset pricing that the CRR model is free of arbitrage if and only if $d < r < u$. ◇

Exercise 14 (Arbitrage in the Trinomial Model)**.** Under which assumptions on the parameters d, m, u, r is the trinomial model free of arbitrage? ◇

Another consequence of the characterization of EMMs is the following **law of one price**, the importance of which will become clear very soon.

Theorem 3.7 (Law of One Price)**.** *Assume that S is free of arbitrage opportunities and let φ be a self-financing trading strategy. Then*

$$X^{\varphi}(t) = S^0(t)\mathbb{E}^{\mathbb{Q}}_t\Big[\frac{X^{\varphi}(T)}{S^0(T)}\Big], \qquad t = 0, 1, \dots, T, \tag{3.5}$$

for all equivalent martingale measures $\mathbb{Q}$. ◇

Proof. Lemma 3.4 (Characterization of EMMs) implies that $\overline{X}^{\varphi}$ is a $\mathbb{Q}$-martingale for each EMM $\mathbb{Q}$ and hence

$$\overline{X}^{\varphi}(t) = \mathbb{E}^{\mathbb{Q}}_t[\overline{X}^{\varphi}(T)], \qquad t = 0, 1, \dots, T.$$

Multiplying both sides by $S^0(t)$ yields Equation (3.5). □

3.2 Risk Neutral Pricing in Finite Market Models

Now that we have a satisfactory answer to the question of existence or absence of arbitrage in finite financial markets, we can start to price options. As it turns out, the first fundamental theorem of asset pricing is quite helpful here as well as the following line of arguments shows.

Consider a market S free of arbitrage opportunities and suppose that you are given the task to introduce a European option to this market. Of course you do not want to create arbitrage opportunities by specifying the wrong price process $P = \{P(t)\}_{t=0,1,\dots,T}$. By the 1^{st} FTAP, you know what to do: You have to ensure that the extended market (S, P) admits an equivalent martingale measure. One way to achieve this is to choose any EMM $\mathbb{Q}$ in the original market S and construct P so that $\overline{P}$ is a $\mathbb{Q}$-martingale. In this section, we show how this is done.

Recall that options are derivative instruments with solely positive payoffs. If the option is European, it has exactly one payoff which the owner of the option receives at time T. In other words, European options can be identified with $\mathfrak{F}(T) = \mathfrak{A}$-measurable random variables.

Definition 3.8 (Option; Attainable; Replication Strategy)**.** Any $\mathbb{R}_+$-valued random variable ξ is referred to as a (European) **option**. We say that ξ is **attainable** if there is a self-financing strategy φ with $X^\varphi(T) = \xi$, in which case we refer to φ as a **replication strategy** for ξ. ◇

For example, a European call option with maturity T, strike K, and underlying S^i is an option in the sense of the previous definition if we set

$$\xi \triangleq (S^i(T) - K)_+.$$

Exercise 15 (Attainability in the CRR Model)**.** Consider the CRR model and an option ξ of the form $\xi \triangleq g(S^1(T))$ for some function $g : \mathbb{R}_+ \to \mathbb{R}_+$.

(i) Construct a replication strategy for ξ in case of $T = 1$.

(ii) Construct a replication strategy for ξ for $T = 2$ in the special case of a European call, i.e. $g(x) = (x - K)_+$ for $K \triangleq s(1+d)(1+u)$. ◇

Exercise 16 (Attainability in the Trinomial Model). Consider the trinomial model with $T = 1$ and fix $K > 0$ with $s(1+d) < K < s(1+u)$. Consider the payoffs

$$\xi_P \triangleq \left(K - S^1(1)\right)_+ \qquad \text{and} \qquad \xi_F \triangleq S^1(1) - K.$$

Are these payoffs attainable? ◇

It should be clear by now that the price of an option should not generate any arbitrage opportunities. It is, however, unclear at this point (and in general simply not true) that such arbitrage-free prices are unique. It is therefore reasonable to introduce another pricing concept which is both rational and unique by construction.

Definition 3.9 ((European) Superhedging Strategy; Superhedging Price). Let ξ be a European option. Any self-financing trading strategy φ satisfying $X^\varphi(T) \geq \xi$ is called **superhedging strategy** for ξ. With this, we refer to the stochastic process $\hat{P} = \{\hat{P}(t)\}_{t=0,1,\ldots,T}$ defined by

$$\hat{P}(t) \triangleq \inf\left\{X^\varphi(t) : \varphi \text{ is a superhedging strategy for } \xi\right\}$$

for $t = 0, 1, \ldots, T$ as the **superhedging price** of ξ. ◇

The superhedging price is thus the minimal capital required to guarantee that the seller of the option can generate enough wealth by clever trading to deliver the payoff ξ at time T. While this is clearly a reasonable concept, we note that the price might still appear way to high from the buyer's point of view unless the option is attainable.

Exercise 17 (Overexpensive Superhedging Prices). Let $S = (S^0, S^1)$ be a finite financial market with $S^0(t) = S^1(t) = 1$ for all $t = 0, 1, \ldots, T$. Compute the superhedging price of an arbitrary option ξ. ◇

We now come to the main result of this section which shows how to produce all arbitrage free prices of an option ξ. Note that the time-T price $P(T)$ of ξ

must be given by $P(T) = \xi$ by absence of arbitrage.

Theorem 3.10 (Risk Neutral Pricing). *Let S be a financial market free of arbitrage opportunities, let ξ be a European option, and fix an adapted process $P = \{P(t)\}_{t=0,1,\dots,T}$ with $P(T) = \xi$. Then there are no arbitrage opportunities in the extended financial market (S, P) if and only if*

$$P(t) = S^0(t)\mathbb{E}_t^{\mathbb{Q}}\Big[\frac{\xi}{S^0(T)}\Big], \qquad t = 0, 1, \dots, T, \tag{3.6}$$

for an equivalent martingale measure $\mathbb{Q}$ in the original market S. Moreover, if there exists a replication strategy φ for ξ, we have

$$P(t) = X^{\varphi}(t), \qquad t = 0, 1, \dots, T,$$

and the superhedging price of ξ is given by

$$\hat{P}(t) = S^0(t)\mathbb{E}_t^{\mathbb{Q}}\Big[\frac{\xi}{S^0(T)}\Big], \qquad t = 0, 1 \dots, T$$

for any equivalent martingale measure $\mathbb{Q}$ in the market S. ◇

Proof. Let $\mathbb{Q}$ be an EMM in the extended market (S, P). Then $\mathbb{Q}$ is clearly also an EMM in the original market S and $\overline{P}$ is a $\mathbb{Q}$-martingale, implying

$$P(t) = S^0(t)\overline{P}(t) = S^0(t)\mathbb{E}_t^{\mathbb{Q}}[\overline{P}(T)] = S^0(t)\mathbb{E}_t^{\mathbb{Q}}\Big[\frac{\xi}{S^0(T)}\Big], \quad t = 0, 1, \dots, T.$$

On the other hand, if $\mathbb{Q}$ is an EMM in S and P is defined as in Equation (3.6), it follows that

$$\overline{P}(t-1) = \mathbb{E}_{t-1}^{\mathbb{Q}}\Big[\frac{\xi}{S^0(T)}\Big] = \mathbb{E}_{t-1}^{\mathbb{Q}}\Big[\mathbb{E}_t^{\mathbb{Q}}\Big[\frac{\xi}{S^0(T)}\Big]\Big] = \mathbb{E}_{t-1}^{\mathbb{Q}}[\overline{P}(t)], \quad t = 1, \dots, T,$$

i.e. $\overline{P}$ is a $\mathbb{Q}$-martingale. But then $\mathbb{Q}$ is an EMM for (S, P), and this market is hence free of arbitrage by the 1^{st} FTAP.

Now let φ be a replication strategy for ξ and let ψ be a superhedging strategy for ξ, i.e. $X^{\psi}(T) \geq \xi = X^{\varphi}(T)$. Then Theorem 3.7 (Law of One Price) implies

$$X^{\varphi}(t) = S^0(t)\mathbb{E}_t^{\mathbb{Q}}\Big[\frac{X^{\varphi}(T)}{S^0(T)}\Big] \leq S^0(t)\mathbb{E}_t^{\mathbb{Q}}\Big[\frac{X^{\psi}(T)}{S^0(T)}\Big] = X^{\psi}(t), \quad t = 0, 1, \dots, T.$$

Since ψ was chosen arbitrarily and since the replication strategy φ is clearly also a superhedging strategy, we conclude that

$$\hat{P}(t) = X^{\varphi}(t) = S^0(t)\mathbb{E}_t^{\mathbb{Q}}\Big[\frac{X^{\varphi}(T)}{S^0(T)}\Big] = P(t), \quad t = 0, 1, \ldots, T. \qquad \square$$

Let us apply the risk neutral pricing theorem to price European call options in the CRR model.

Example 3.11 (Risk Neutral Pricing in the CRR Model)**.** Consider the CRR model $S = (S^0, S^1)$ with $d < r < u$ and recall that an EMM $\mathbb{Q}$ is given by

$$\mathbb{Q}[\{\omega\}] = q^{U(\omega)}(1-q)^{D(\omega)}, \qquad \omega \in \Omega,$$

where $q \triangleq (r-d)/(u-d)$. Consider a European call on S^1 with strike K, i.e.

$$\xi \triangleq (S^1(T) - K)_+.$$

By risk neutral pricing, an arbitrage free price of ξ at time zero is

$$\begin{aligned} P(0) &= \mathbb{E}^{\mathbb{Q}}\Big[\frac{\xi}{S^0(T)}\Big] \\ &= \frac{1}{(1+r)^T}\mathbb{E}^{\mathbb{Q}}\big[(S^1(T) - K)_+\big] \\ &= \frac{1}{(1+r)^T}\sum_{\omega\in\Omega} q^{U(\omega)}(1-q)^{D(\omega)}\Big(s(1+u)^{U(\omega)}(1+d)^{D(\omega)} - K\Big)_+. \end{aligned}$$

Since for each $n = 0, 1, \ldots, T$ it holds that

$$\text{there are } \binom{T}{n} = \frac{T!}{n!(T-n)!} \text{ many } \omega \in \Omega \text{ with } U(\omega) = n,$$

it follows that

$$P(0) = \frac{1}{(1+r)^T}\sum_{n=0}^{T}\binom{T}{n} q^n(1-q)^{T-n}\Big(s(1+u)^n(1+d)^{T-n} - K\Big)_+. \qquad \diamond$$

It is important to note that the measure in the risk neutral pricing formula must be an EMM $\mathbb{Q}$ and not the original measure $\mathbb{P}$. Using $\mathbb{P}$ as the pricing measure typically results in arbitrage as the following exercise illuminates.

Exercise 18 (Arbitrage under the Physical Measure)**.** Give an example of an arbitrage free finite financial market S and an option ξ such that the price

$$P(t) \triangleq S^0(t)\mathbb{E}_t^{\mathbb{P}}\Big[\frac{\xi}{S^0(T)}\Big], \qquad t = 0, 1, \dots, T,$$

leads to arbitrage opportunities in the extended financial market (S, P). ◇

Exercise 19 (Attainability in a Market with Transaction Costs)**.** Consider a market $S = (S^0, S^1)$ with $\Omega = \{\omega_1, \omega_2\}$, $T = 1$ and $S^0(0) = S^0(1) = 1$. We assume that trading in the security S^1 incurs transaction costs, which is to say that S^1 is bought at the ask price S^{ask} and sold at the bid price S^{bid} with

$$\begin{aligned} &S^{\mathsf{ask}}(0) = 5, && S^{\mathsf{ask}}(1, \omega_1) = 3, && S^{\mathsf{ask}}(1, \omega_2) = 6,\\ &S^{\mathsf{bid}}(0) = 3, && S^{\mathsf{bid}}(1, \omega_1) = 2, && S^{\mathsf{bid}}(1, \omega_2) = 4. \end{aligned}$$

We assume moreover that any position in S^1 has to be liquidated at time T and fix an option ξ with $\xi(\omega_1) = 0$ and $\xi(\omega_2) = 2$.

(i) Construct a replication strategy for ξ.

(ii) Construct a superhedging strategy for ξ requiring strictly less initial capital than the replication strategy constructed in (i). ◇

3.3 Completeness and the Second Fundamental Theorem

The significance of the risk neutral pricing theorem is that the arbitrage free prices are exactly those given by Equation (3.6) for some EMM $\mathbb{Q}$. In particular, if the EMM is unique, there exists exactly one arbitrage free price. Moreover, again by risk neutral pricing, the arbitrage free price of an option is unique if the option can be replicated. This suggests that there is a link between uniqueness of EMMs and attainability of options.

Definition 3.12 (Complete Market)**.** We say that S is **complete** if it does not admit any arbitrage opportunities and every option is attainable. ◇

Complete markets are particularly tractable models as it follows from risk neutral pricing that all arbitrage free prices are uniquely determined by Equation (3.6). Since the set of all options corresponds to the set of all positive random variables, it appears likely that this can only be true if there is at most one EMM $\mathbb{Q}$, i.e. the market S is complete if and only if there exists a unique EMM. This is the **Second Fundamental Theorem of Asset Pricing** or **2nd FTAP** for short.

Theorem 3.13 (Second Fundamental Theorem of Asset Pricing)**.** *In a finite financial market S, the following statements are equivalent:*

(i) S is complete.

(ii) There exists a unique EMM $\mathbb{Q}$. ◇

Proof. Step 1. We show that (i) implies (ii). For this, suppose that S is complete so that there exists at least one EMM $\mathbb{Q}$ by the 1st FTAP. Now assume by contradiction that there exists another EMM $\mathbb{Q}'$. For $F \in \mathfrak{A}$, consider the option $\xi \triangleq \mathbb{1}_F S^0(T)$. By completeness, this option is attainable and hence there exists a replication strategy φ. But then, by Theorem 3.7 (Law of One Price), we find that

$$\mathbb{Q}[F] = \mathbb{E}^{\mathbb{Q}}[\mathbb{1}_F] = \mathbb{E}^{\mathbb{Q}}\Big[\frac{\xi}{S^0(T)}\Big] = \mathbb{E}^{\mathbb{Q}}\Big[\frac{X^\varphi(T)}{S^0(T)}\Big] = X^\varphi(0)$$
$$= \mathbb{E}^{\mathbb{Q}'}\Big[\frac{X^\varphi(T)}{S^0(T)}\Big] = \ldots = \mathbb{Q}'[F].$$

Since $F \in \mathfrak{A}$ was chosen arbitrarily, it follows that $\mathbb{Q} = \mathbb{Q}'$, i.e. the EMM in this market is unique.

Step 2. Recall that L denotes the space of all random variables and define

$$\mathrm{P} \triangleq \{\overline{X}^\varphi(T) \in \mathrm{L} : \varphi \text{ is a self-financing trading strategy}\}.$$

We proceed to show that $\mathrm{P} = \mathrm{L}$ if the EMM $\mathbb{Q}$ is unique. It is clear that P is a linear subspace of the linear space L and that

$$(X, Y) \mapsto \mathbb{E}^{\mathbb{Q}}[XY], \qquad X, Y \in \mathrm{L},$$

defines a scalar product on L. Suppose by contradiction that $\mathrm{P} \subsetneq \mathrm{L}$. Then there exists $Y \neq 0$ with $Y \in \mathrm{L} \setminus \mathrm{P}$ such that

$$\mathbb{E}^{\mathbb{Q}}[XY] = 0 \qquad \text{for all } X \in \mathrm{P}. \tag{3.7}$$

Since $\max_{\omega\in\Omega} |Y(\omega)| < \infty$, we may assume that $|Y(\omega)| \leq 1/2$ for all $\omega \in \Omega$ after possibly rescaling Y. But then we can define a measure $\mathbb{Q}' \sim \mathbb{P}$ by

$$\mathbb{Q}'[\{\omega\}] \triangleq (Y(\omega) + 1)\mathbb{Q}[\{\omega\}], \qquad \omega \in \Omega.$$

For any $X \in \mathrm{L}$, we observe that

$$\mathbb{E}^{\mathbb{Q}'}[X] = \sum_{\omega\in\Omega} X(\omega)(Y(\omega) + 1)\mathbb{Q}[\{\omega\}] = \mathbb{E}^{\mathbb{Q}}[XY] + \mathbb{E}^{\mathbb{Q}}[X].$$

Since the constant random variable $\mathbb{1}_\Omega$ is clearly in P (choose an initial wealth of 1 and consider the no-trading strategy), it follows from Equation (3.7) that $\mathbb{E}^{\mathbb{Q}}[Y] = \mathbb{E}^{\mathbb{Q}}[Y\mathbb{1}_\Omega] = 0$ and thus

$$\mathbb{Q}'[\Omega] = \mathbb{E}^{\mathbb{Q}'}[\mathbb{1}_\Omega] = \mathbb{E}^{\mathbb{Q}}[Y] + \mathbb{E}^{\mathbb{Q}}[\mathbb{1}_\Omega] = \mathbb{E}^{\mathbb{Q}}[\mathbb{1}_\Omega] = \mathbb{Q}[\Omega] = 1,$$

i.e. $\mathbb{Q}'$ is a probability measure. Let now φ be an arbitrary trading strategy with initial wealth zero. Then $\overline{X}^\varphi(T) \in \mathrm{P}$ by definition and thus by Equation (3.7) and Lemma 3.4 (Characterization of EMMs)

$$\mathbb{E}^{\mathbb{Q}'}[\overline{X}^\varphi(T)] = \mathbb{E}^{\mathbb{Q}}[\overline{X}^\varphi(T)Y] + \mathbb{E}^{\mathbb{Q}}[\overline{X}^\varphi(T)] = \mathbb{E}^{\mathbb{Q}}[\overline{X}^\varphi(T)] = 0.$$

But then, by Lemma 3.4 again, $\mathbb{Q}'$ must also be an EMM. Since $\mathbb{Q}' \neq \mathbb{Q}$ we find the desired contradiction to the uniqueness of $\mathbb{Q}$ and hence we must have $\mathrm{P} = \mathrm{L}$.

Step 3. We show that (ii) implies (i). Let $\mathbb{Q}$ be the unique EMM and let ξ be an arbitrary option. Then $\xi/S^0(T) \in \mathrm{L} = \mathrm{P}$ by the previous step, and hence by the definition of P there exists a self-financing strategy φ with $\overline{X}^\varphi(T) = \xi/S^0(T)$, i.e.

$$\xi = S^0(T)\overline{X}^\varphi(T) = X^\varphi(T).$$

In other words, ξ is attainable and thus S and the proof are complete. □

Let us apply the 2^{nd} FTAP to the CRR model. In Example 3.6 we have argued that the CRR model is free of arbitrage opportunities if and only if $d < r < u$. To see that the model is also complete in this case, we have to argue that the EMM is unique.

Example 3.14 (Completeness of the CRR Model)**.** Consider the CRR model with $d < r < u$ and denote by $\mathbb{Q}$ an arbitrary EMM. Then it holds that

$$\overline{S}^1(t-1) = \mathbb{E}^{\mathbb{Q}}_{t-1}[\overline{S}^1(t)] = \frac{\overline{S}^1(t-1)}{1+r}\mathbb{E}^{\mathbb{Q}}_{t-1}[1+R_t], \qquad t = 1, \dots, T,$$

which is only possible if

$$\mathbb{E}^{\mathbb{Q}}_{t-1}[R_t] = r, \qquad t = 1, \dots, T.$$

In particular, each $\mathbb{E}^{\mathbb{Q}}_{t-1}[R_t]$ is equal to a constant and thus $\mathfrak{F}(0)$-measurable[2]. But then $\mathbb{E}^{\mathbb{Q}}[R_t] = \mathbb{E}^{\mathbb{Q}}_0[\mathbb{E}^{\mathbb{Q}}_{t-1}[R_t]] = \mathbb{E}^{\mathbb{Q}}_{t-1}[R_t]$ and thus

$$r = \mathbb{E}^{\mathbb{Q}}_{t-1}[R_t] = \mathbb{E}^{\mathbb{Q}}[R_t] = u\,\mathbb{Q}[R_t = u] + d\,\mathbb{Q}[R_t = d], \qquad t = 1, \dots, T.$$

Using $\mathbb{Q}[R_t = u] + \mathbb{Q}[R_t = d] = 1$, this is only possible if

$$\mathbb{Q}[R_t = u] = q \triangleq \frac{r-d}{u-d} \qquad \text{and} \qquad \mathbb{Q}[R_t = d] = 1 - q, \qquad t = 1, \dots, T.$$

As in Example 3.6, this implies that

$$\mathbb{Q}[\{\omega\}] = q^{U(\omega)}(1-q)^{D(\omega)}, \qquad \omega \in \Omega,$$

so that $\mathbb{Q}$ is uniquely determined. ◇

Exercise 20 (Incompleteness of the Trinomial Model)**.** Compute all EMMs in the trinomial model and conclude that this model is incomplete. ◇

3.4 Incomplete Markets and Superhedging Duality

Let us now take a closer look at the case of arbitrage free markets which are **incomplete**, i.e. for which there exist options which are **not attainable**.

[2]This even yields independence of $R_1, \dots, R_T$. Indeed, as $\mathfrak{F}(t-1) = \sigma(R_1, \dots, R_{t-1})$,

$$\begin{aligned} r = \mathbb{E}^{\mathbb{Q}}_{t-1}[R_t] &= u\,\mathbb{Q}[R_t = u|\mathfrak{F}^S(t-1)] + d\big(1 - \mathbb{Q}[R_t = u|\mathfrak{F}^S(t-1)]\big) \\ &= u\,\mathbb{Q}[R_t = u|R_1, \dots, R_{t-1}] + d\big(1 - \mathbb{Q}[R_t = u|R_1, \dots, R_{t-1}]\big), \end{aligned}$$

i.e. $\mathbb{Q}[R_t = u|R_1, \dots, R_{t-1}]$ is constant (and necessarily equal to $\mathbb{Q}[R_t = u]$), which is only possible if R_t is independent of $R_1, \dots, R_{t-1}$.

By the 1^{st} FTAP and risk neutral pricing, there exists an entire **range of arbitrage free prices** for these options. Moreover, in this situation, the role of the superhedging price $\hat{P}$ is thus far unclear in the sense that we do not know (yet) if it is arbitrage free, or, more generally, how it relates to the set of arbitrage free prices.

Definition 3.15 (Lower Price; Upper Price)**.** Let S be an arbitrage free financial market and let ξ be an option. Then we refer to

$$\underline{\pi}_\xi \triangleq \inf_{\mathbb{Q}\in\mathcal{M}_e(S)} \mathbb{E}^{\mathbb{Q}}\Big[\frac{\xi}{S^0(T)}\Big] \qquad \text{and} \qquad \overline{\pi}_\xi \triangleq \sup_{\mathbb{Q}\in\mathcal{M}_e(S)} \mathbb{E}^{\mathbb{Q}}\Big[\frac{\xi}{S^0(T)}\Big]$$

as the **lower price** and the **upper price** of ξ, respectively, where

$$\mathcal{M}_e(S) \triangleq \{\mathbb{Q} : \mathfrak{A} \to [0,1] : \mathbb{Q} \text{ is an EMM for } S\}$$

denotes the set of equivalent martingale measures for S. ◇

Intuition dictates that the lower and upper price characterize the entire set of arbitrage free prices: Any arbitrage free price (at time $t=0$) should be contained in the interval $[\underline{\pi}_\xi, \overline{\pi}_\xi]$. At first sight, it is however not clear if the lower and upper prices themselves are arbitrage free. Our first task is to clarify this situation.

Theorem 3.16 (Set of Arbitrage Free Prices)**.** *In an arbitrage free financial market S, let ξ be an option with superhedging price $\hat{P}$.*

(i) If ξ is attainable, its unique arbitrage free price at time $t=0$ is given by

$$\underline{\pi}_\xi = \overline{\pi}_\xi = \hat{P}(0) = \mathbb{E}^{\mathbb{Q}}\Big[\frac{\xi}{S^0(T)}\Big], \qquad \mathbb{Q}\in\mathcal{M}_e(S).$$

(ii) If ξ is not attainable, then $\underline{\pi}_\xi < \overline{\pi}_\xi$ and

$$\big(\underline{\pi}_\xi, \overline{\pi}_\xi\big) = \Big\{\mathbb{E}^{\mathbb{Q}}\Big[\frac{\xi}{S^0(T)}\Big] : \mathbb{Q}\in\mathcal{M}_e(S)\Big\}. \qquad ◇$$

Proof. Note that (i) is an immediate consequence of Theorem 3.10 (Risk Neutral Pricing), so we only have to prove (ii).

Step 1. We write

$$I \triangleq \Big\{\mathbb{E}^{\mathbb{Q}}\Big[\frac{\xi}{S^0(T)}\Big] : \mathbb{Q} \in \mathcal{M}_e(S)\Big\}$$

and observe that I is a bounded interval. Indeed, boundedness follows from

$$0 \leq \mathbb{E}^{\mathbb{Q}}\Big[\frac{\xi}{S^0(T)}\Big] \leq \sup_{\omega\in\Omega} \frac{\xi(\omega)}{S^0(T,\omega)} < +\infty, \qquad \mathbb{Q} \in \mathcal{M}_e(S).$$

Regarding the interval structure, let $p_1, p_2 \in I$ and fix $\lambda \in [0,1]$. Then, by definition of I, we can find $\mathbb{Q}_1, \mathbb{Q}_2 \in \mathcal{M}_e(S)$ such that

$$p_1 = \mathbb{E}^{\mathbb{Q}_1}\Big[\frac{\xi}{S^0(T)}\Big] \qquad \text{and} \qquad p_2 = \mathbb{E}^{\mathbb{Q}_2}\Big[\frac{\xi}{S^0(T)}\Big].$$

Observe that $\mathbb{Q} \triangleq \lambda\mathbb{Q}_1 + (1-\lambda)\mathbb{Q}_2 \in \mathcal{M}_e(S)$ and thus

$$\lambda p_1 + (1-\lambda)p_2 = \lambda\mathbb{E}^{\mathbb{Q}_1}\Big[\frac{\xi}{S^0(T)}\Big] + (1-\lambda)\mathbb{E}^{\mathbb{Q}_2}\Big[\frac{\xi}{S^0(T)}\Big] = \mathbb{E}^{\mathbb{Q}}\Big[\frac{\xi}{S^0(T)}\Big] \in I,$$

showing that I is a convex subset of the real line and therefore an interval.

Step 2. To conclude, it suffices to argue that for all $p \in I$ we can find $p_-, p_+ \in I$ with $p_- < p < p_+$. For this, fix $p \in I$ and note that, by definition, there exists $\mathbb{Q} \in \mathcal{M}_e(S)$ with $p = \mathbb{E}^{\mathbb{Q}}[\xi/S^0(T)]$. We proceed to construct $\mathbb{Q}_-, \mathbb{Q}_+ \in \mathcal{M}_e(S)$ such that

$$\mathbb{E}^{\mathbb{Q}_-}\Big[\frac{\xi}{S^0(T)}\Big] < p = \mathbb{E}^{\mathbb{Q}}\Big[\frac{\xi}{S^0(T)}\Big] < \mathbb{E}^{\mathbb{Q}_+}\Big[\frac{\xi}{S^0(T)}\Big].$$

To begin with, we introduce an adapted process $M = \{M(t)\}_{t=0,1,\dots,T}$ through

$$M(t) \triangleq \mathbb{E}^{\mathbb{Q}}_t\Big[\frac{\xi}{S^0(T)}\Big], \qquad t = 0,1,\dots,T,$$

and note that M is a martingale since, for all $t = 1,\dots,T$,

$$M(t-1) = \mathbb{E}^{\mathbb{Q}}_{t-1}\Big[\frac{\xi}{S^0(T)}\Big] = \mathbb{E}^{\mathbb{Q}}_{t-1}\Big[\mathbb{E}^{\mathbb{Q}}_t\Big[\frac{\xi}{S^0(T)}\Big]\Big] = \mathbb{E}^{\mathbb{Q}}_{t-1}\big[M(t)\big].$$

We observe that we may write

$$M(0) + \sum_{t=1}^{T} \Delta M(t) = M(T) = \frac{\xi}{S^0(T)},$$

and thus, since ξ is not attainable, there exists $t \in \{1, \dots, T\}$ such that

$$\Delta M(t) \notin \mathrm{P}_t \triangleq \big\{ \langle \varphi(t) | \Delta \overline{S}(t) \rangle : \varphi(t) \text{ is } \mathbb{R}^d\text{-valued and } \mathfrak{F}(t-1)\text{-measurable} \big\}. \quad (3.8)$$

Writing $\Omega = \{\omega_1, \dots, \omega_N\}$, we identify $\mathfrak{F}(t)$-measurable random variables X with vectors in $\mathbb{R}^N$ of the form

$$\Big(X(\omega_1) \mathbb{Q}[\{\omega_1\}]^{1/2}, \dots, X(\omega_N) \mathbb{Q}[\{\omega_N\}]^{1/2} \Big).$$

The scalar product on $\mathbb{R}^N$ then induces a scalar product $\langle \cdot | \cdot \rangle_{\mathbb{Q}}$ on the space L_t of $\mathfrak{F}(t)$-measurable random variables as follows:

$$\langle X | Y \rangle_{\mathbb{Q}} = \sum_{\omega \in \Omega} X(\omega) Y(\omega) \mathbb{Q}[\{\omega\}] = \mathbb{E}^{\mathbb{Q}}[XY], \qquad X, Y \in \mathrm{L}_t.$$

Applying the hyperplane separation theorem[3] to the singleton set $\{\Delta M(t)\}$ and the closed linear space P_t allows us to find $Z \in \mathrm{L}_t$ such that

$$\mathbb{E}^{\mathbb{Q}}[YZ] = \langle Y | Z \rangle_{\mathbb{Q}} = 0 < \langle \Delta M(t) | Z \rangle_{\mathbb{Q}} = \mathbb{E}^{\mathbb{Q}}[\Delta M(t) Z], \quad Y \in \mathrm{P}_t. \quad (3.9)$$

After rescaling, we may assume that $|Z(\omega)| < 1/3$ for all $\omega \in \Omega$ so that

$$Z_+ \triangleq 1 + Z - \mathbb{E}^{\mathbb{Q}}_{t-1}[Z] > 0 \qquad \text{and} \qquad Z_- \triangleq 1 - Z + \mathbb{E}^{\mathbb{Q}}_{t-1}[Z] > 0.$$

With this, we can now define two measures $\mathbb{Q}_+, \mathbb{Q}_-$ on $(\Omega, \mathfrak{A})$ by

$$\mathbb{Q}_+[\{\omega\}] \triangleq Z_+(\omega) \mathbb{Q}[\{\omega\}] \qquad \text{and} \qquad \mathbb{Q}_-[\{\omega\}] \triangleq Z_-(\omega) \mathbb{Q}[\{\omega\}], \qquad \omega \in \Omega.$$

But then, by the tower property of conditional expectation,

$$\begin{aligned}
\mathbb{E}^{\mathbb{Q}_\pm}\Big[\frac{\xi}{S^0(T)}\Big] &= \mathbb{E}^{\mathbb{Q}}\Big[\frac{\xi}{S^0(T)} Z_\pm\Big] \\
&= \mathbb{E}^{\mathbb{Q}}\Big[\frac{\xi}{S^0(T)}\Big] \pm \mathbb{E}^{\mathbb{Q}}\Big[\frac{\xi}{S^0(T)} Z\Big] \mp \mathbb{E}^{\mathbb{Q}}\Big[\frac{\xi}{S^0(T)} \mathbb{E}_{t-1}[Z]\Big] \\
&= p \pm \mathbb{E}^{\mathbb{Q}}\Big[\mathbb{E}^{\mathbb{Q}}_t\Big[\frac{\xi}{S^0(T)}\Big] Z\Big] \mp \mathbb{E}^{\mathbb{Q}}\Big[\mathbb{E}^{\mathbb{Q}}_{t-1}\Big[\frac{\xi}{S^0(T)}\Big] \mathbb{E}_{t-1}[Z]\Big]
\end{aligned}$$

[3]Note that $(x_1, \dots, x_N) \mapsto Q(x_1, \dots, x_N) \triangleq (x_1 \mathbb{Q}[\{\omega_1\}]^{1/2}, \dots, x_N \mathbb{Q}[\{\omega_N\}]^{1/2})$ is a linear bijection from $\mathbb{R}^N$ into itself. Thus the image sets $Q(\{\Delta M(t)\})$ and $Q(\mathrm{P}_t)$ are disjoint, $Q(\{\Delta M(t)\})$ is compact and convex, and $Q(\mathrm{P}_t)$ is linear and closed. Thus we can indeed apply the hyperplane separation theorem in $\mathbb{R}^N$.

$$= p \pm \mathbb{E}^{\mathbb{Q}}\big[M(t)Z\big] \mp \mathbb{E}^{\mathbb{Q}}\big[M(t-1)\mathbb{E}_{t-1}[Z]\big] = p \pm \mathbb{E}^{\mathbb{Q}}\big[\Delta M(t)Z\big].$$

Using Equation (3.9), this implies

$$\mathbb{E}^{\mathbb{Q}_-}\Big[\frac{\xi}{S^0(T)}\Big] = p - \mathbb{E}^{\mathbb{Q}}\big[\Delta M(t)Z\big] < p < p + \mathbb{E}^{\mathbb{Q}}\big[\Delta M(t)Z\big] = \mathbb{E}^{\mathbb{Q}_+}\Big[\frac{\xi}{S^0(T)}\Big].$$

Therefore, in order to conclude, we are only left with showing that $\mathbb{Q}_-$ and $\mathbb{Q}_+$ are EMMs.

Step 3. We show that $\mathbb{Q}_+ \in \mathcal{M}_e(S)$, the proof for $\mathbb{Q}_- \in \mathcal{M}_e(S)$ works analogously. Let us first argue that $\mathbb{Q}_+$ is indeed a probability measure. For this, it suffices to observe that

$$\begin{aligned}\mathbb{Q}_+[\Omega] = \mathbb{E}^{\mathbb{Q}_+}[\mathbb{1}_\Omega] = \mathbb{E}^{\mathbb{Q}}[\mathbb{1}_\Omega] + \mathbb{E}^{\mathbb{Q}}[\mathbb{1}_\Omega Z] - \mathbb{E}^{\mathbb{Q}}\big[\mathbb{1}_\Omega \mathbb{E}^{\mathbb{Q}}_{t-1}[Z]\big]\\ = \mathbb{Q}[\Omega] + \mathbb{E}^{\mathbb{Q}}[Z] - \mathbb{E}^{\mathbb{Q}}\big[\mathbb{E}^{\mathbb{Q}}_{t-1}[Z]\big] = 1.\end{aligned}$$

Let now φ be a trading strategy with initial wealth zero. Then Lemma 3.4 (Characterization of EMMs) shows that

$$\begin{aligned}\mathbb{E}^{\mathbb{Q}_+}\big[\overline{X}^\varphi(T)\big] &= \mathbb{E}^{\mathbb{Q}}\big[\overline{X}^\varphi(T)\big] + \mathbb{E}^{\mathbb{Q}}\big[\overline{X}^\varphi(T)\big(Z - \mathbb{E}^{\mathbb{Q}}_{t-1}[Z]\big)\big]\\ &= 0 + \mathbb{E}^{\mathbb{Q}}\big[\overline{X}^\varphi(T)\big(Z - \mathbb{E}^{\mathbb{Q}}_{t-1}[Z]\big)\big]\end{aligned}$$

since $\mathbb{Q}$ is an EMM. Now Z is $\mathfrak{F}(t)$-measurable, hence we can condition on $\mathfrak{F}(t)$ and use the $\mathbb{Q}$-martingale property of $\overline{X}^\varphi$ to obtain

$$\begin{aligned}&\mathbb{E}^{\mathbb{Q}_+}\big[\overline{X}^\varphi(T)\big]\\ &= \mathbb{E}^{\mathbb{Q}}\Big[\mathbb{E}^{\mathbb{Q}}_t\big[\overline{X}^\varphi(T)\big]\big(Z - \mathbb{E}^{\mathbb{Q}}_{t-1}[Z]\big)\Big]\\ &= \mathbb{E}^{\mathbb{Q}}\Big[\overline{X}^\varphi(t)\big(Z - \mathbb{E}^{\mathbb{Q}}_{t-1}[Z]\big)\Big]\\ &= \mathbb{E}^{\mathbb{Q}}\Big[\overline{X}^\varphi(t-1)\big(Z - \mathbb{E}^{\mathbb{Q}}_{t-1}[Z]\big)\Big] + \mathbb{E}^{\mathbb{Q}}\Big[\langle\varphi(t)|\Delta\overline{S}(t)\rangle\big(Z - \mathbb{E}^{\mathbb{Q}}_{t-1}[Z]\big)\Big].\end{aligned}$$

Regarding the first expectation, conditioning on $\mathfrak{F}(t-1)$ yields

$$\mathbb{E}^{\mathbb{Q}}\Big[\overline{X}^\varphi(t-1)\big(Z - \mathbb{E}^{\mathbb{Q}}_{t-1}[Z]\big)\Big] = \mathbb{E}^{\mathbb{Q}}\Big[\overline{X}^\varphi(t-1)\big(\mathbb{E}^{\mathbb{Q}}_{t-1}[Z] - \mathbb{E}^{\mathbb{Q}}_{t-1}[Z]\big)\Big] = 0.$$

For the second expectation, we condition on $\mathfrak{F}(t-1)$ and use the predictability of φ (for the first equality) followed by $\langle\varphi(t)|\Delta\overline{S}(t)\rangle \in \mathrm{P}_t$, Equation (3.9), and the $\mathbb{Q}$-martingale property of $\overline{S}$ (for the second equality) to obtain

$$\mathbb{E}^{\mathbb{Q}}\Big[\langle\varphi(t)|\Delta\overline{S}(t)\rangle\big(Z-\mathbb{E}^{\mathbb{Q}}_{t-1}[Z]\big)\Big]$$
$$=\mathbb{E}^{\mathbb{Q}}\Big[\langle\varphi(t)|\Delta\overline{S}(t)\rangle Z\Big]-\mathbb{E}^{\mathbb{Q}}\Big[\langle\varphi(t)|\mathbb{E}^{\mathbb{Q}}_{t-1}[\Delta\overline{S}(t)]\rangle\mathbb{E}^{\mathbb{Q}}_{t-1}[Z]\Big]=0.$$

In combination, this implies

$$\mathbb{E}^{\mathbb{Q}_+}\big[\overline{X}^{\varphi}(T)\big]=0$$

and thus $\mathbb{Q}_+$ is an EMM by Lemma 3.4 (Characterization of EMMs). □

Our next aim is to find a connection between the set of arbitrage free prices and the **superhedging price** of unattainable options. We shall see that the superhedging price $\hat{P}(0)$ coincides with the upper price $\overline{\pi}_{\xi}$. Before we can establish this result, we first need to take a closer look at the structure of certain supermartingales.

Theorem 3.17 (Optional Decomposition)**.** *Let S be an arbitrage free financial market and let $X=\{X(t)\}_{t=0,1,\dots,T}$ be an adapted process. Then the following statements are equivalent:*

(i) X is a $\mathbb{Q}$-supermartingale for every $\mathbb{Q}\in\mathcal{M}_{\mathrm{e}}(S)$.

(ii) There exists a trading strategy φ with initial wealth zero and an increasing adapted process $C=\{C(t)\}_{t=0,1,\dots,T}$ with $C(0)=0$ such that X admits the optional decomposition

$$X(t)=X(0)+\varphi\bullet\overline{S}(t)-C(t),\qquad t=0,1,\dots,T.$$ ◇

Proof. Step 1. (ii) implies (i). Let $\mathbb{Q}\in\mathcal{M}_{\mathrm{e}}(S)$ be an EMM and suppose that (ii) holds. Then Lemma 3.4 (Characterization of EMMs) implies that $\varphi\bullet\overline{S}$ is a $\mathbb{Q}$-martingale and hence

$$\mathbb{E}^{\mathbb{Q}}_{t-1}\big[\Delta X(t)\big]=\mathbb{E}^{\mathbb{Q}}_{t-1}\big[\Delta\big(\varphi\bullet\overline{S}\big)(t)\big]-\mathbb{E}^{\mathbb{Q}}_{t-1}\big[\Delta C(t)\big]=-\mathbb{E}^{\mathbb{Q}}_{t-1}\big[\Delta C(t)\big]\le 0$$

for all $t=1,\dots,T$ since C is increasing. But then X is a $\mathbb{Q}$-supermartingale as claimed.

Step 2. (i) implies (ii). It is sufficient to show that for each $t = 1, \dots, T$, we can find an $\mathfrak{F}(t-1)$-measurable and $\mathbb{R}^d$-valued random variable φ_t as well as an $\mathfrak{F}(t)$-measurable and $\mathbb{R}_+$-valued random variable C_t such that

$$\Delta X(t) = \langle \varphi_t | \Delta \overline{S}(t) \rangle - C_t. \tag{3.10}$$

If this is the case, we can define a trading strategy φ with initial wealth zero and an increasing adapted process C with $C(0) \triangleq 0$ by setting

$$\varphi(t) = \varphi_t \qquad \text{and} \qquad C(t) \triangleq \sum_{s=1}^{t} C_s, \qquad t = 1, \dots, T,$$

and it follows that

$$X(t) - X(0) = \sum_{\tau=1}^{t} \Delta X(\tau) = \sum_{\tau=1}^{t} \langle \varphi_\tau | \Delta \overline{S}(\tau) \rangle - \sum_{\tau=1}^{t} C_\tau = \varphi \bullet \overline{S}(t) - C(t)$$

for all $t = 0, 1, \dots, T$, which concludes the proof. To establish Equation (3.10), we use a hyperplane separation argument yet again.

Step 3. For $t = 1, \dots, T$, recall the set P_t defined in Equation (3.8) and denote by L_t^+ the set of $\mathbb{R}_+$-valued, $\mathfrak{F}(t)$–measurable random variables. Then Equation (3.10) rewrites as

$$\Delta X(t) \in \mathrm{K}_t \triangleq \mathrm{P}_t - \mathrm{L}_t^+, \qquad t = 1, \dots, T.$$

We assume by contradiction that there exists some t for which this is not the case, i.e. $\Delta X(t) \notin \mathrm{K}_t$. We observe that K_t is a closed convex cone: If $Y_1, Y_2 \in \mathrm{K}_t$ and $\lambda_1, \lambda_2 \geq 0$, then $\lambda_1 Y_1 + \lambda_2 Y_2 \in \mathrm{K}_t$. Let us now apply the hyperplane separation theorem to separate K_t from the singleton $\{\Delta X(t)\}$. Then, for an arbitrary but fixed EMM $\mathbb{Q} \in \mathcal{M}_{\mathrm{e}}(S)$, we can find an $\mathfrak{F}(t)$-measurable random variable Z such that

$$\alpha \triangleq \sup_{Y \in \mathrm{K}_t} \mathbb{E}^{\mathbb{Q}}[YZ] = \langle Y | Z \rangle_{\mathbb{Q}} < \langle \Delta X(t) | Z \rangle_{\mathbb{Q}} = \mathbb{E}^{\mathbb{Q}}\big[\Delta X(t) Z\big] < +\infty. \tag{3.11}$$

As $0 \in \mathrm{K}_t$, it follows that $\alpha \geq 0$. But as K_t is a cone, we must have $\alpha = 0$.[4] Moreover, since $\mathrm{P}_t \subset \mathrm{K}_t$ is linear, we see that

$$\mathbb{E}^{\mathbb{Q}}[YZ] = 0, \qquad Y \in \mathrm{P}_t. \tag{3.12}$$

[4]Indeed, otherwise there exists $Y \in \mathrm{K}_t$ with $\langle Y | Z \rangle_{\mathbb{P}} > 0$. But as $\lambda Y \in \mathrm{K}_t$ for any $\lambda \geq 0$ and $\langle \lambda Y | Z \rangle_{\mathbb{P}} = \lambda \langle Y | Z \rangle_{\mathbb{P}}$, thus would yield the contradiction $\alpha = \infty$.

Now denote by $\{F_1, \dots, F_n\}$ the unique partition of Ω generating $\mathfrak{F}(t)$. Fix $\omega \in \Omega$ and choose $i = 1, \dots, n$ such that $\omega \in F_i$. Observe that Z is $\mathfrak{F}(t)$-measurable and hence in particular constant on F_i in that $Z(\omega) = Z(\bar{\omega})$ for each $\bar{\omega} \in F_i$. Moreover, $\mathbb{1}_{F_i}$ is positive and $\mathfrak{F}(t)$-measurable, from which we conclude that $-\mathbb{1}_{F_i} \in K_t$. But then, by Equation (3.11),

$$0 \geq \langle -\mathbb{1}_{F_i} | Z \rangle_{\mathbb{Q}} = -\mathbb{E}^{\mathbb{Q}}[\mathbb{1}_{F_i} Z] = -Z(\omega)\mathbb{Q}[F_i],$$

implying that $Z(\omega) \geq 0$ for all $\omega \in \Omega$. For $\varepsilon \in (0,1)$, let us now define a measure $\mathbb{Q}^\varepsilon \sim \mathbb{P}$ by

$$\mathbb{Q}^\varepsilon[\{\omega\}] \triangleq Z_\varepsilon(\omega)\mathbb{Q}[\{\omega\}], \quad \omega \in \Omega, \quad \text{where } Z_\varepsilon \triangleq \frac{\varepsilon + (1-\varepsilon)Z}{\mathbb{E}^{\mathbb{Q}}_{t-1}[\varepsilon + (1-\varepsilon)Z]} > 0.$$

Note that Z_ε is $\mathfrak{F}(t)$-measurable since Z is $\mathfrak{F}(t)$-measurable and that

$$\mathbb{E}^{\mathbb{Q}}_{t-1}[Z_\varepsilon] = \frac{\mathbb{E}^{\mathbb{Q}}_{t-1}[\varepsilon + (1-\varepsilon)Z]}{\mathbb{E}^{\mathbb{Q}}_{t-1}[\varepsilon + (1-\varepsilon)Z]} = 1.$$

Hence $\mathbb{Q}^\varepsilon$ is a probability measure since

$$\mathbb{Q}^\varepsilon[\Omega] = \mathbb{E}^{\mathbb{Q}}[\mathbb{1}_\Omega Z_\varepsilon] = \mathbb{E}^{\mathbb{Q}}\big[\mathbb{E}^{\mathbb{Q}}_{t-1}[Z_\varepsilon]\big] = 1.$$

Moreover, if φ is a trading strategy with initial wealth zero, it follows from the $\mathfrak{F}(t)$-measurability of Z_ε (for the first equality) and the martingale property of $\overline{X}^\varphi$ under $\mathbb{Q}$ (for the second equality), that

$$\begin{aligned}\mathbb{E}^{\mathbb{Q}^\varepsilon}\big[\overline{X}^\varphi(T)\big] &= \mathbb{E}^{\mathbb{Q}}\Big[\mathbb{E}^{\mathbb{Q}}_t\big[\overline{X}^\varphi(T)\big] Z_\varepsilon\Big] \\ &= \mathbb{E}^{\mathbb{Q}}\big[\overline{X}^\varphi(t) Z_\varepsilon\big] = \mathbb{E}^{\mathbb{Q}}\big[\overline{X}^\varphi(t-1) Z_\varepsilon\big] + \mathbb{E}^{\mathbb{Q}}\big[\langle \varphi(t) | \Delta \overline{S}(t) \rangle Z_\varepsilon\big].\end{aligned}$$

The first expectation here is zero since conditioning on $\mathfrak{F}(t-1)$, the fact that $\mathbb{E}^{\mathbb{Q}}_{t-1}[Z_\varepsilon] = 1$, and the $\mathbb{Q}$-martingale property of $\overline{X}^\varphi$ yield

$$\mathbb{E}^{\mathbb{Q}}\big[\overline{X}^\varphi(t-1) Z_\varepsilon\big] = \mathbb{E}^{\mathbb{Q}}\big[\overline{X}^\varphi(t-1)\mathbb{E}^{\mathbb{Q}}_{t-1}[Z_\varepsilon]\big] = \mathbb{E}^{\mathbb{Q}}\big[\overline{X}^\varphi(t-1)\big] = \overline{X}^\varphi(0) = 0.$$

Regarding the second expectation, we introduce the $\mathfrak{F}(t-1)$-measurable random variable

$$\varphi_\varepsilon(t) \triangleq \frac{\varphi(t)}{\mathbb{E}^{\mathbb{Q}}_{t-1}[\varepsilon + (1-\varepsilon)Z]}$$

and use this to compute

$$\mathbb{E}^{\mathbb{Q}^\varepsilon}\big[\overline{X}^\varphi(T)\big] = \mathbb{E}^{\mathbb{Q}}\big[\langle \varphi(t) | \Delta \overline{S}(t) \rangle Z_\varepsilon\big]$$

$$
\begin{aligned}
&= \varepsilon \mathbb{E}^{\mathbb{Q}}\Big[\frac{\langle \varphi(t) | \Delta \overline{S}(t)\rangle}{\mathbb{E}^{\mathbb{Q}}_{t-1}[\varepsilon + (1-\varepsilon)Z]}\Big] + (1-\varepsilon)\mathbb{E}^{\mathbb{Q}}\Big[\frac{\langle \varphi(t) | \Delta \overline{S}(t)\rangle}{\mathbb{E}^{\mathbb{Q}}_{t-1}[\varepsilon + (1-\varepsilon)Z]} Z\Big] \\
&= \varepsilon \mathbb{E}^{\mathbb{Q}}\Big[\frac{\langle \varphi(t) | \mathbb{E}^{\mathbb{Q}}_{t-1}[\Delta \overline{S}(t)]\rangle}{\mathbb{E}^{\mathbb{Q}}_{t-1}[\varepsilon + (1-\varepsilon)Z]}\Big] + (1-\varepsilon)\mathbb{E}^{\mathbb{Q}}\big[\langle \varphi_\varepsilon(t) | \overline{S}(t)\rangle Z\big] = 0,
\end{aligned}
$$

where we have used that $\overline{S}$ is a $\mathbb{Q}$-martingale for the first term to be zero and $\langle \varphi_\varepsilon(t) | \overline{S}(t)\rangle \in \mathrm{P}_t$ and Equation (3.12) for the second term to be zero. Altogether, we have thus argued that $\mathbb{E}^{\mathbb{Q}_\varepsilon}[\overline{X}^\varphi(T)] = 0$ and hence $\mathbb{Q}_\varepsilon$ is an EMM by Lemma 3.4 (Characterization of EMMs).

Step 4. By assumption, the process X is a supermartingale under any measure $\mathbb{Q}_\varepsilon$ with $\varepsilon \in (0,1)$. Therefore $\mathbb{E}^{\mathbb{Q}_\varepsilon}_{t-1}[\Delta X(t)] \leq 0$ and thus, since $Z \geq 0$,

$$
\mathbb{E}^{\mathbb{Q}_\varepsilon}_{t-1}\big[\Delta X(t) \mathbb{E}^{\mathbb{Q}}_{t-1}[\varepsilon + (1-\varepsilon)Z]\big] = \mathbb{E}^{\mathbb{Q}_\varepsilon}_{t-1}[\Delta X(t)]\, \mathbb{E}^{\mathbb{Q}}_{t-1}[\varepsilon + (1-\varepsilon)Z] \leq 0.
$$

Taking expectations and using the definition of Z_ε thus implies

$$
\begin{aligned}
0 &\geq \mathbb{E}^{\mathbb{Q}_\varepsilon}\Big[\mathbb{E}^{\mathbb{Q}_\varepsilon}_{t-1}\big[\Delta X(t) \mathbb{E}^{\mathbb{Q}}_{t-1}[\varepsilon + (1-\varepsilon)Z]\big]\Big] \\
&= \mathbb{E}^{\mathbb{Q}_\varepsilon}\Big[\Delta X(t) \mathbb{E}^{\mathbb{Q}}_{t-1}[\varepsilon + (1-\varepsilon)Z]\Big] \\
&= \mathbb{E}^{\mathbb{Q}_\varepsilon}\Big[(\varepsilon + (1-\varepsilon)Z)\Delta X(t) \frac{1}{Z_\varepsilon}\Big] = \mathbb{E}^{\mathbb{Q}}\big[(\varepsilon + (1-\varepsilon)Z)\Delta X(t)\big].
\end{aligned}
$$

As $\varepsilon \downarrow 0$, we arrive at

$$
\mathbb{E}^{\mathbb{Q}}[Z \Delta X(t)] = \lim_{\varepsilon \downarrow 0} \mathbb{E}^{\mathbb{Q}}\big[(\varepsilon + (1-\varepsilon)Z)\Delta X(t)\big] \leq 0,
$$

which, together with $\alpha = 0$, contradicts Equation (3.11) and hence the proof is complete. □

By Lemma 3.4 (Characterization of EMMs), any discounted wealth process $\overline{X}^\varphi$ is a martingale under any EMM $\mathbb{Q} \in \mathcal{M}_{\mathrm{e}}(S)$ and can be written as

$$
\overline{X}^\varphi = \overline{X}^\varphi(0) + \varphi \bullet \overline{S}.
$$

On the other hand, if X is a supermartingale under any EMM $\mathbb{Q}$, the optional decomposition theorem states that it is of the form

$$
X = X(0) + \varphi \bullet \overline{S} - C
$$

for an increasing process C. This allows for an **economic interpretation** of X as follows: $X(0)$ may be thought of as the initial wealth, the martingale component $\varphi \bullet \overline{S}$ are the profits and losses from trading, and the increasing component C may be thought of as **cumulative consumption**. In other words, X may be thought of as the wealth process of a non-self-financing trading strategy, in which the investor withdraws $\Delta C(t)$ at each trading date $t = 1, \dots, T$ for consumption purposes.

We proceed with an alternative description of the set of EMMs. For this, let us first recall that a measure $\mathbb{Q}$ on $(\Omega, \mathfrak{A})$ is equivalent to $\mathbb{P}$ if and only if there exists a uniquely determined strictly positive random variable Z with

$$\mathbb{Q}[\{\omega\}] = Z(\omega)\mathbb{P}[\{\omega\}], \qquad \omega \in \Omega.$$

In this case, $\mathbb{Q}$ is a probability measure if and only if $\mathbb{E}[Z] = 1$ since

$$\mathbb{Q}[\Omega] = \mathbb{E}^{\mathbb{Q}}[\mathbb{1}_\Omega] = \mathbb{E}[\mathbb{1}_\Omega Z] = \mathbb{E}[Z].$$

Fixing $t = 0, 1, \dots, T$, we observe that $\mathbb{E}_t[Z]$ is then the unique density of $\mathbb{Q}$ with respect to $\mathbb{P}$ restricted to $\mathfrak{F}(t)$ as clearly, $\mathbb{E}_t[Z]$ is $\mathfrak{F}(t)$-measurable and

$$\mathbb{Q}[F] = \mathbb{E}^{\mathbb{Q}}[\mathbb{1}_F] = \mathbb{E}[\mathbb{1}_F Z] = \mathbb{E}\big[\mathbb{1}_F \mathbb{E}_t[Z]\big], \qquad F \in \mathfrak{F}(t).$$

For each $F \in \mathfrak{F}(t)$ and any random variable Y, we see that

$$\mathbb{E}[\mathbb{1}_F ZY] = \mathbb{E}^{\mathbb{Q}}[\mathbb{1}_F Y] = \mathbb{E}^{\mathbb{Q}}\big[\mathbb{1}_F \mathbb{E}_t^{\mathbb{Q}}[Y]\big] = \mathbb{E}\big[\mathbb{1}_F Z \mathbb{E}_t^{\mathbb{Q}}[Y]\big] = \mathbb{E}\big[\mathbb{1}_F \mathbb{E}_t[Z]\mathbb{E}_t^{\mathbb{Q}}[Y]\big],$$

which, by definition of conditional expectation, shows that

$$\mathbb{E}_t[ZY] = \mathbb{E}_t[Z]\mathbb{E}_t^{\mathbb{Q}}[Y] \quad \text{or, equivalently,} \quad \mathbb{E}_t^{\mathbb{Q}}[Y] = \mathbb{E}_t\Big[Y \frac{Z}{\mathbb{E}_t[Z]}\Big]. \tag{3.13}$$

We may think of the family of densities $\{\mathbb{E}_t[Z]\}_{t=0,1,\dots,T}$ as a stochastic process, which we refer to as the density process of $\mathbb{Q}$ (with respect to $\mathbb{P}$). The following result now characterizes the set of density process corresponding to equivalent martingale measures. For this, we introduce a set $\mathcal{D}(S)$ consisting of all strictly positive adapted processes $D = \{D(t)\}_{t=0,1,\dots,T}$ with

$$D(0) = 1 \qquad \text{and} \qquad \mathbb{E}_{t-1}[D(t)] = 1, \qquad t = 1, \dots, T, \tag{3.14}$$

and

$$\mathbb{E}_{t-1}\big[\overline{X}^{\varphi}(t)D(t)\big] = \overline{X}^{\varphi}(t-1), \qquad t = 1, \dots, T, \tag{3.15}$$

for any self-financing trading strategy φ. We claim that there is a one-to-one correspondence between $\mathcal{D}(S)$ and $\mathcal{M}_e(S)$.

Lemma 3.18 (State Price Deflators). *If $\mathbb{Q} \in \mathcal{M}_e(S)$, there exists $D_\mathbb{Q} \in \mathcal{D}(S)$ such that the density of $\mathbb{Q}$ with respect to $\mathbb{P}$ on $\mathfrak{F}(t)$ is given by*

$$\frac{d\mathbb{Q}}{d\mathbb{P}}\bigg|_{\mathfrak{F}(t)} = \prod_{\tau=0}^{t} D_\mathbb{Q}(\tau), \qquad t = 0, 1, \dots, T.$$

Conversely, if $D \in \mathcal{D}(S)$, then the measure $\mathbb{Q}_D$ defined via

$$\mathbb{Q}_D[\{\omega\}] \triangleq \mathbb{P}[\{\omega\}] \prod_{t=0}^{T} D(t, \omega), \qquad \omega \in \Omega,$$

is an EMM. In any case, if $(\mathbb{Q}, D_\mathbb{Q})$ is such a pair, then

$$\mathbb{E}_t^\mathbb{Q}[Y] = \mathbb{E}_t\Big[Y \textstyle\prod_{\tau=t+1}^{T} D_\mathbb{Q}(\tau)\Big], \qquad t = 0, 1, \dots, T, \tag{3.16}$$

for any arbitrary random variable Y and

$$\mathbb{E}_t^\mathbb{Q}[Y] = \mathbb{E}_t\big[Y D_\mathbb{Q}(t+1)\big] \tag{3.17}$$

in the special case of Y being $\mathfrak{F}(t+1)$-measurable. ◇

Proof. Step 1. Let $\mathbb{Q}$ be an EMM and denote by Z its density with respect to $\mathbb{P}$, so that $\mathbb{E}_t[Z]$ is the density of $\mathbb{Q}$ with respect to $\mathbb{P}$ on $\mathfrak{F}(t)$. We define the process $D_\mathbb{Q} = \{D_\mathbb{Q}(t)\}_{t=0,1,\dots,T}$ by

$$D_\mathbb{Q}(0) \triangleq 1 \qquad \text{and} \qquad D_\mathbb{Q}(t) \triangleq \frac{\mathbb{E}_t[Z]}{\mathbb{E}_{t-1}[Z]} > 0, \qquad t = 1, \dots, T.$$

Then, since $\mathbb{E}_0[Z] = \mathbb{E}[Z] = 1$, it follows that

$$\mathbb{E}_t[Z] = \prod_{\tau=1}^{t} \frac{\mathbb{E}_\tau[Z]}{\mathbb{E}_{\tau-1}[Z]} = \prod_{\tau=0}^{t} D_\mathbb{Q}(\tau), \qquad t = 0, 1, \dots, T, \tag{3.18}$$

and we have

$$\mathbb{E}_{t-1}[D_\mathbb{Q}(t)] = \mathbb{E}_{t-1}\left[\frac{\mathbb{E}_t[Z]}{\mathbb{E}_{t-1}[Z]}\right] = \frac{\mathbb{E}_{t-1}\big[\mathbb{E}_t[Z]\big]}{\mathbb{E}_{t-1}[Z]} = 1, \qquad t = 1, \dots, T. \tag{3.19}$$

Let now Y be an arbitrary random variable and $t = 0, 1, \dots, T$. Using Equation (3.18) in Equation (3.13), it follows that

$$\mathbb{E}_t^\mathbb{Q}[Y] = \mathbb{E}_t\left[Y \frac{Z}{\mathbb{E}_t[Z]}\right] = \mathbb{E}_t\left[Y \frac{\mathbb{E}_T[Z]}{\mathbb{E}_t[Z]}\right] = \mathbb{E}_t\Big[Y \textstyle\prod_{\tau=t+1}^{T} D_\mathbb{Q}(\tau)\Big].$$

Moreover, if Y is $\mathfrak{F}(t+1)$-measurable, we can simplify this by iteratively conditioning and applying Equation (3.19) as follows:

$$\begin{aligned}\mathbb{E}_t^{\mathbb{Q}}[Y] &= \mathbb{E}_t\Big[Y\textstyle\prod_{\tau=t+1}^{T} D_{\mathbb{Q}}(\tau)\Big]\\ &= \mathbb{E}_t\Big[Y\textstyle\prod_{\tau=t+1}^{T-1} D_{\mathbb{Q}}(\tau)\mathbb{E}_{T-1}[D_{\mathbb{Q}}(T)]\Big]\\ &= \mathbb{E}_t\Big[Y\textstyle\prod_{\tau=t+1}^{T-1} D_{\mathbb{Q}}(\tau)\Big] = \ldots = \mathbb{E}_t\Big[Y D_{\mathbb{Q}}(t+1)\Big].\end{aligned} \tag{3.20}$$

Let now φ be a self-financing trading strategy. Using that $\mathbb{Q}$ is an EMM, it follows that $\overline{X}^{\varphi}$ is a $\mathbb{Q}$-martingale and hence, by Equation (3.20),

$$\overline{X}^{\varphi}(t-1) = \mathbb{E}_{t-1}^{\mathbb{Q}}[\overline{X}^{\varphi}(t)] = \mathbb{E}_{t-1}[\overline{X}^{\varphi}(t)D_{\mathbb{Q}}(t)], \qquad t = 1, \ldots, T.$$

We have therefore argued that $\mathbb{Q} \in \mathcal{M}_{\mathrm{e}}(S)$ implies $D_{\mathbb{Q}} \in \mathcal{D}(S)$ and that, in this case, the pair $(\mathbb{Q}, D_{\mathbb{Q}})$ satisfies Equation (3.16) and Equation (3.17).

Step 2. Let $D \in \mathcal{D}(S)$ and define $Z \triangleq \prod_{t=0}^{T} D(t) > 0$. It follows from Equation (3.14) (applied iteratively) that

$$\begin{aligned}\mathbb{E}_t[Z] = \mathbb{E}_t\Big[\prod_{\tau=0}^{T} D(\tau)\Big] &= \prod_{\tau=0}^{t} D(\tau)\mathbb{E}_t\Big[\prod_{\tau=t+1}^{T-1} D(\tau)\mathbb{E}_{T-1}[D(T)]\Big]\\ &= \prod_{\tau=0}^{t} D(\tau)\mathbb{E}_t\Big[\prod_{\tau=t+1}^{T-1} D(\tau)\Big] = \ldots = \prod_{\tau=0}^{t} D(\tau)\end{aligned} \tag{3.21}$$

for all $t = 0, 1, \ldots, T$. In particular, we have $\mathbb{E}[Z] = \mathbb{E}_0[Z] = D(0) = 1$, so that the measure $\mathbb{Q}_D$ defined via

$$\mathbb{Q}_D[\{\omega\}] = Z(\omega)\mathbb{P}[\{\omega\}], \qquad \omega \in \Omega,$$

is a probability measure equivalent to $\mathbb{P}$. Now let φ be a trading strategy with initial wealth zero and fix $t = 1, \ldots, T$. Applying Equation (3.13) with $Y = \overline{X}^{\varphi}(t)$ followed by an application of Equation (3.21) and yet another iterative application of Equation (3.14) shows that

$$\begin{aligned}\mathbb{E}_{t-1}^{\mathbb{Q}_D}\big[\overline{X}^{\varphi}(t)\big] &= \mathbb{E}_{t-1}\Big[\overline{X}^{\varphi}(t)\frac{Z}{\mathbb{E}_{t-1}[Z]}\Big]\\ &= \mathbb{E}_{t-1}\Big[\overline{X}^{\varphi}(t)\textstyle\prod_{\tau=t}^{T} D(\tau)\Big]\\ &= \mathbb{E}_{t-1}\Big[\overline{X}^{\varphi}(t)\textstyle\prod_{\tau=t}^{T-1} D(\tau)\mathbb{E}_{T-1}[D(T)]\Big]\end{aligned}$$

$$= \mathbb{E}_{t-1}\Big[\overline{X}^{\varphi}(t)\prod_{\tau=t}^{T-1} D(\tau)\Big]$$
$$= \ldots = \mathbb{E}_{t-1}\big[\overline{X}^{\varphi}(t)D(t)\big] = \overline{X}^{\varphi}(t-1),$$

where the last equality holds by Equation (3.15). But then $\overline{X}^{\varphi}$ is a $\mathbb{Q}$-martingale, hence $\mathbb{E}^{\mathbb{Q}}[\overline{X}^{\varphi}(T)] = 0$, and thus $\mathbb{Q}$ is an EMM by Lemma 3.4 (Characterization of EMMs). This concludes the proof. □

We can now turn to the characterization of the superheding price by showing that it coincides with the upper price. This result is referred to as the **superhedging duality**, as it states that the superhedging price can be computed both by **minimization** (over the initial wealth of all superhedging strategies) and by **maximization** (over all arbitrage free prices at time zero). Observe that this implies in particular that the superhedging price of an unattainable option leads to arbitrage by Theorem 3.16 (Set of Arbitrage Free Prices).

Theorem 3.19 (Superhedging Duality)**.** *In an arbitrage free market S, let ξ be an unattainable option with superhedging price $\hat{P}$ and upper price $\overline{\pi}_{\xi}$. Then*

$$\hat{P}(0) = \overline{\pi}_{\xi}$$

and there exists a superhedging strategy with initial wealth $\hat{P}(0) = \overline{\pi}_{\xi}$, i.e. the infimum in the definition of the superhedging price is attained. ◇

Proof. Step 1: $\overline{\pi}_{\xi} \leq \hat{P}(0)$. Let φ be a superhedging strategy for ξ.[5] Then Theorem 3.7 (Law of One Price) shows that

$$\overline{\pi}_{\xi} = \sup_{\mathbb{Q}\in\mathcal{M}_e(S)} \mathbb{E}^{\mathbb{Q}}\Big[\frac{\xi}{S^0(T)}\Big] \leq \sup_{\mathbb{Q}\in\mathcal{M}_e(S)} \mathbb{E}^{\mathbb{Q}}\big[\overline{X}^{\varphi}(T)\big] = X^{\varphi}(0).$$

Now minimize over all superhedging strategies φ to conclude that

$$\overline{\pi}_{\xi} \leq \inf\big\{X^{\varphi}(0) : \varphi \text{ superhedging strategy for } \xi\big\} = \hat{P}(0).$$

[5] φ exists since ξ is bounded: Let $\varphi \equiv 0$ and choose an initial wealth $x \in \mathbb{R}$ with $x \geq \xi$.

Step 2: $\overline{\pi}_\xi \geq \hat{P}(0)$. We introduce a dynamic version of the upper price via

$$X(t) \triangleq \sup_{\mathbb{Q}\in\mathcal{M}_e(S)} \mathbb{E}_t^{\mathbb{Q}}\Big[\frac{\xi}{S^0(T)}\Big], \qquad t = 0, 1, \dots, T.$$

Note that $X(0) = \overline{\pi}_\xi$ and $X(T) = \xi/S^0(T)$. In the next step, we show that $X = \{X(t)\}_{t=0,1,\dots,T}$ is a supermartingale under each $\mathbb{Q} \in \mathcal{M}_e(S)$. Hence Theorem 3.17 (Optional Decomposition) implies the existence of a self-financing trading strategy φ with initial wealth $X(0)$ and an increasing adapted process $C = \{C(t)\}_{t=0,1\dots,T}$ with $C(0) = 0$ such that

$$X(t) = X(0) + \varphi \bullet \overline{S}(t) - C(t) = \overline{X}^\varphi(t) - C(t), \qquad t = 0, 1, \dots, T.$$

Since $C(T) \geq 0$, it follows that

$$\overline{X}^\varphi(T) = X(T) + C(T) \geq X(T) = \frac{\xi}{S^0(T)},$$

i.e. φ is a superhedging strategy for ξ and thus

$$\overline{\pi}_\xi = X(0) \geq \hat{P}(0).$$

But then $\overline{\pi}_\xi = X(0) = \hat{P}(0)$ by step 1. Moreover, we note that φ is a superhedging strategy with initial wealth $X(0) = \hat{P}(0)$, which implies that the infimum in the definition of the superhedging price is indeed attained.

Step 3: Supermartingale property. Let $\mathbb{Q} \in \mathcal{M}_e(S)$. We have to show that the process X defined in the previous step is a $\mathbb{Q}$-supermartingale. For this, fix $\varepsilon > 0$, let $t = 1, \dots, T$, and choose $\bar{\mathbb{Q}} \in \mathcal{M}_e(S)$ such that

$$\sup_{\mathbb{Q}\in\mathcal{M}_e(S)} \mathbb{E}_t^{\mathbb{Q}}\Big[\frac{\xi}{S^0(T)}\Big] \leq \mathbb{E}_t^{\bar{\mathbb{Q}}}\Big[\frac{\xi}{S^0(T)}\Big] + \varepsilon, \tag{3.22}$$

which is possible since Ω is finite. Now denote by $D, \bar{D} \in \mathcal{D}(S)$ the processes corresponding to the $\mathbb{Q}, \bar{\mathbb{Q}} \in \mathcal{M}_e(S)$ implied by Lemma 3.18 (State Price Deflators). Applying Equation (3.17) for $(\mathbb{Q}, D)$ with $Y = X(t)$ followed by the definition of $X(t)$ and the inequality in Equation (3.22), it follows that

$$\begin{aligned}\mathbb{E}_{t-1}^{\mathbb{Q}}\big[X(t)\big] &= \mathbb{E}_{t-1}\big[D(t)X(t)\big]\\ &= \mathbb{E}_{t-1}\Big[D(t) \sup_{\mathbb{Q}\in\mathcal{M}_e(S)} \mathbb{E}_t^{\mathbb{Q}}\Big[\frac{\xi}{S^0(T)}\Big]\Big]\end{aligned}$$

$$
\begin{aligned}
&\leq \mathbb{E}_{t-1}\Big[D(t)\mathbb{E}_t^{\bar{\mathbb{Q}}}\Big[\frac{\xi}{S^0(T)}\Big]\Big] + \varepsilon\mathbb{E}_{t-1}[D(t)] \\
&= \mathbb{E}_{t-1}\Big[D(t)\mathbb{E}_t\Big[\frac{\xi}{S^0(T)}\prod_{\tau=t+1}^{T}\bar{D}(\tau)\Big]\Big] + \varepsilon \\
&= \mathbb{E}_{t-1}\Big[D(t)\frac{\xi}{S^0(T)}\prod_{\tau=t+1}^{T}\bar{D}(\tau)\Big] + \varepsilon,
\end{aligned}
$$

where we have used Equation (3.16) for $(\bar{\mathbb{Q}}, \bar{D})$ with $Y = \xi/S^0(T)$ and Equation (3.14) for D to obtain the second to last equality. Now define a new process $D^* = \{D^*(t)\}_{t=0,1,\ldots,T}$ by

$$
D^*(\tau) \triangleq \begin{cases} D(\tau), & \text{for } \tau = 0, 1, \ldots, t \\ \bar{D}(\tau) & \text{for } t = t+1, \ldots, T. \end{cases}
$$

Then $D^* \in \mathcal{D}(S)$, and hence there exists a corresponding EMM $\mathbb{Q}^*$ by Lemma 3.18 (State Price Deflator). But then we can apply Equation (3.16) for $(\mathbb{Q}^*, D^*)$ to obtain

$$
\begin{aligned}
\mathbb{E}_{t-1}^{\mathbb{Q}}\big[X(t)\big] &\leq \mathbb{E}_{t-1}\Big[D(t)\frac{\xi}{S^0(T)}\prod_{\tau=t+1}^{T}\bar{D}(\tau)\Big] + \varepsilon \\
&= \mathbb{E}_{t-1}\Big[\frac{\xi}{S^0(T)}\prod_{\tau=t}^{T}D^*(\tau)\Big] + \varepsilon \\
&= \mathbb{E}_{t-1}^{\mathbb{Q}^*}\Big[\frac{\xi}{S^0(T)}\Big] + \varepsilon \leq \sup_{\mathbb{Q}\in\mathcal{M}_e(S)} \mathbb{E}_{t-1}^{\mathbb{Q}}\Big[\frac{\xi}{S^0(T)}\Big] + \varepsilon = X(t-1) + \varepsilon.
\end{aligned}
$$

Sending $\varepsilon \downarrow 0$ shows that

$$
\mathbb{E}_{t-1}^{\mathbb{Q}}\big[X(t)\big] \leq X(t-1),
$$

and hence X is a $\mathbb{Q}$-supermartingale and the proof is complete. □

In the previous proof, we have argued in step 3 that

$$
\sup_{\mathbb{Q}\in\mathcal{M}_e(S)} \mathbb{E}_{t-1}^{\mathbb{Q}}\Big[\frac{\xi}{S^0(T)}\Big] \geq \mathbb{E}_{t-1}^{\bar{\mathbb{Q}}}\Big[\sup_{\mathbb{Q}\in\mathcal{M}_e(S)} \mathbb{E}_t^{\mathbb{Q}}\Big[\frac{\xi}{S^0(T)}\Big]\Big], \qquad \bar{\mathbb{Q}} \in \mathcal{M}_e(S),
$$

for all $t = 1, \ldots, T$, implying that

$$
\sup_{\mathbb{Q}\in\mathcal{M}_e(S)} \mathbb{E}_{t-1}^{\mathbb{Q}}\Big[\frac{\xi}{S^0(T)}\Big] \geq \sup_{\bar{\mathbb{Q}}\in\mathcal{M}_e(S)} \mathbb{E}_{t-1}^{\bar{\mathbb{Q}}}\Big[\sup_{\mathbb{Q}\in\mathcal{M}_e(S)} \mathbb{E}_t^{\mathbb{Q}}\Big[\frac{\xi}{S^0(T)}\Big]\Big], \quad t = 1, \ldots, T.
$$

But choosing $\mathbb{Q} = \bar{\mathbb{Q}}$ to estimate the second supremum on the right hand side implies that

$$\sup_{\bar{\mathbb{Q}} \in \mathcal{M}_e(S)} \mathbb{E}_{t-1}^{\bar{\mathbb{Q}}}\left[\sup_{\mathbb{Q} \in \mathcal{M}_e(S)} \mathbb{E}_t^{\mathbb{Q}}\left[\frac{\xi}{S^0(T)} \right] \right] \geq \sup_{\bar{\mathbb{Q}} \in \mathcal{M}_e(S)} \mathbb{E}_{t-1}^{\bar{\mathbb{Q}}}\left[\frac{\xi}{S^0(T)} \right], \quad t = 1, \dots, T.$$

In combination, the following result obtains, which may be thought of as a kind of **time consistency** of the dynamic version of the upper price.

Corollary 3.20 (Dynamic Programming)**.** *For any arbitrage free market S, we have*

$$\sup_{\mathbb{Q} \in \mathcal{M}_e(S)} \mathbb{E}_{t-1}^{\mathbb{Q}}\left[\sup_{\bar{\mathbb{Q}} \in \mathcal{M}_e(S)} \mathbb{E}_t^{\bar{\mathbb{Q}}}\left[\frac{\xi}{S^0(T)} \right] \right] = \sup_{\mathbb{Q} \in \mathcal{M}_e(S)} \mathbb{E}_{t-1}^{\mathbb{Q}}\left[\frac{\xi}{S^0(T)} \right], \quad t = 1, \dots, T,$$

for all options ξ. ◇

For now, this is all there is to say about arbitrage and European options in finite financial market models. Before moving on to American options, let us briefly repeat what we have learned so far:

- ▷ Financial market models are **free of arbitrage** if and only if there exists at least one equivalent martingale measure.
- ▷ All arbitrage free prices of options are determined by the equivalent martingale measures via the **risk neutral pricing** formula given by Equation (3.6).
- ▷ If an option is attainable, its arbitrage free price is **unique** and given by the wealth process under the corresponding replication strategy. Moreover, in this situation, the arbitrage free price coincides with the superhedging price.
- ▷ An arbitrage free market model is **complete**, i.e. every option is attainable, if and only if the equivalent martingale measure is unique.
- ▷ If an option is not attainable, the set of arbitrage free prices is an entire **interval** with boundary points given by the lower and upper price, respectively. Moreover, the upper price coincides with the superhedging price.

- ▷ There always is a **superhedging strategy** which attains the superhedging price. Moreover, the superhedging price of an unattainable option leads to arbitrage since the interval of arbitrage free prices is open.

Chapter 4

American Options and Optimal Stopping

Let us now turn to the pricing of American options. In contrast to their European counterparts, American options can be exercised at any trading date $t = 0, 1, \ldots, T$. This leads to the following definition.

Definition 4.1 (American Option)**.** Any $\mathbb{R}_+$-valued adapted stochastic process $\xi = \{\xi(t)\}_{t=0,1,\ldots,T}$ is referred to as an **American option**. ◇

Here, we think of $\xi(t)$ as the payoff the owner receives at time t if the option is exercised at time t. The time-t payoff is of course known to the owner of the option, which is why we assume ξ to be adapted. As an example, the American put option on S^i with strike K is obtained by setting

$$\xi(t) \triangleq \big(K - S^i(t)\big)_+, \qquad t = 0, 1, \ldots, T.$$

The value of an American option depends critically on the exercise date chosen by the owner. Assuming that the owner decides to exercise the option at time τ^*, the payoff of the option is $\xi(\tau^*)$, and risk neutral pricing hence suggests that the price $P_{\mathsf{a}}(0)$ of the option (at time zero) should be

$$P_{\mathsf{a}}(0) = \mathbb{E}^{\mathbb{Q}}[\overline{\xi}(\tau^*)],$$

where $\mathbb{Q}$ denotes some EMM. But how does the owner of the option choose the exercise date? A rational choice is the date which maximizes the value of the option, i.e. the owner aims to find an exercise date τ^* such that

$$\mathbb{E}^{\mathbb{Q}}[\overline{\xi}(\tau^*)] = \sup_{\tau} \mathbb{E}^{\mathbb{Q}}[\overline{\xi}(\tau)]. \tag{4.1}$$

Note that, in general, we need not restrict this maximization problem to deterministic times $\tau \in \{0, 1, \ldots, T\}$, but we could (and should!) allow the owner to choose the timing dynamically depending on the information available at any point in time. In other words, we should allow τ to be a (finite) **stopping time**, i.e. a mapping

$$\tau : \Omega \to \{0, 1, \ldots, T\}, \qquad \omega \mapsto \tau(\omega),$$

such that

$$\{\tau \leq t\} \in \mathfrak{F}(t), \qquad t = 0, 1, \ldots, T.$$

Since we maximize over stopping times in Equation (4.1), we refer to this optimization problem as an **optimal stopping problem**. If you are unfamiliar with stopping times or need a reminder, now is a good opportunity to have a look at Appendix A.

4.1 Optimal Stopping Problems and the Snell Envelope

The objective of this section is to study the **optimal stopping problem**

$$V(0) \triangleq \sup_{\tau \in \mathbb{T}} \mathbb{E}[\xi(\tau)],$$

where $\mathbb{T}$ denotes the set of all finite, i.e. $\{0, 1, \ldots, T\}$-valued, stopping times and $\xi = \{\xi(t)\}_{t=0,1,\ldots,T}$ is an arbitrary adapted stochastic process taking values in $\mathbb{R}$. More generally, we can consider the time-t conditional optimal stopping problems

$$V(t) \triangleq \sup_{\tau \in \mathbb{T}_t} \mathbb{E}_t[\xi(\tau)], \qquad t = 0, 1, \ldots, T, \tag{4.2}$$

where we write $\mathbb{T}_t$ for the set of all stopping times taking values in the truncated time index set $\mathcal{T}_t \triangleq \{t, t+1, \ldots, T\}$.

In Equation (4.2), we note that $V(t)$ is $\mathfrak{F}(t)$-measurable by finiteness of Ω and $\mathbb{T}_t$. In particular, $V = \{V(t)\}_{t=0,1,\ldots,T}$ is an adapted real-valued process which we refer to as the **value process**. We note that on infinite probability spaces, this need in general not be true and one would have to replace the supremum by the essential supremum to guarantee measurability of $V(t)$ and adaptedness of V.

Given $t = 0, 1, \ldots, T$, we say that a stopping time $\tau^* \in \mathbb{T}_t$ is **optimal** for $V(t)$ if it satisfies

$$V(t) = \sup_{\tau \in \mathbb{T}_t} \mathbb{E}_t[\xi(\tau)] = \mathbb{E}_t[\xi(\tau^*)].$$

This is particularly simple for the time-T conditional optimal stopping problem as the only $\mathcal{T}_T = \{T\}$-valued stopping time is the constant stopping time T. This is to say that if we have waited up to time T, we have no other possibility than to stop and receive the reward $\xi(T)$. In particular, we have

$$V(T) = \sup_{\tau \in \mathbb{T}_T} \mathbb{E}_T[\xi(\tau)] = \mathbb{E}_T[\xi(T)] = \xi(T).$$

While obvious, this is a good starting point: We can now consider what to do at time $T-1$ and step by step work our way **backward in time** until we arrive at $t = 0$. At each time $t = 0, 1, \ldots, T-1$, we have to decide if we want to stop immediately and receive $\xi(t)$ or if we want to postpone our decision until time $t+1$ and receive $V(t+1)$ then (at best, i.e. provided we act optimally from time $t+1$ onward). The time-t expected reward in this case is thus $\mathbb{E}_t[V(t+1)]$.

Now, if we have to decide what to do at time t, we only have to compare the two rewards $\xi(t)$ (stop immediately) and $\mathbb{E}_t[V(t+1)]$ (do not stop at time t and act optimally from $t+1$ onward). In particular, we expect[1] that the value process satisfies the recursive relation[2]

$$V(t) = \xi(t) \vee \mathbb{E}_t[V(t+1)], \qquad t = 0, 1, \ldots, T-1, \qquad V(T) = \xi(T), \tag{4.3}$$

and that it is optimal to stop at time t if $V(t) = \xi(t)$. The recursive representation of the value process is often referred to as the **Bellman principle** or **dynamic programming principle**. If the Bellman principle is indeed correct, the optimal stopping problem is essentially solved: First we compute V via the backward recursion in Equation (4.3). A candidate for an optimal stopping time is then constructed in a second step and taken to be the first time that V coincides with ξ, i.e. we expect that for each $t = 0, 1, \ldots, T$, the stopping time[3]

$$\tau_t^* \triangleq \inf\big\{s \in \mathcal{T}_t : V(s) = \xi(s)\big\}$$

[1] Careful! This is just a heuristic argument and by no means a rigorous proof.

[2] We use the symbol $\vee$ for the pointwise maximum, i.e. $a \vee b \triangleq \max\{a, b\}$ for all $a, b \in \mathbb{R}$.

[3] τ_t^* is indeed a stopping time by Lemma A.7 (Hitting Times with Finite Time Index Sets) since it can be represented as a hitting time: $\tau_t^* = \inf\{s \in \mathcal{T}_t : V(s) - \xi(s) \in \{0\}\}$.

to be optimal for $V(t)$. The remainder of this section is dedicated to turning this **heuristic idea** into a rigorous mathematical argument. The starting point is the study of the right hand side of the recursion in Equation (4.3).

Definition 4.2 (Snell Envelope)**.** The adapted process $Z = \{Z(t)\}_{t=0,1,\dots,T}$ defined by

$$Z(T) \triangleq \xi(T), \qquad Z(t) = \xi(t) \vee \mathbb{E}_t[Z(t+1)], \qquad t = 0, 1, \dots, T-1,$$

is referred to as the **Snell envelope** of ξ. ◇

In light of the above discussion, we expect that the value process and the Snell envelope are identical and that we can use the Snell envelope to construct optimal stopping times. The first step is to analyze some properties of the Snell envelope.

Proposition 4.3 (Characterization of the Snell Envelope)**.** *The Snell envelope Z of ξ is the smallest supermartingale which dominates ξ. That is, Z is a supermartingale which satisfies*

$$Z(t) \geq \xi(t), \qquad t = 0, 1, \dots, T,$$

and if $Y = \{Y(t)\}_{t=0,1,\dots,T}$ is another supermartingale with this property, it holds that $Y(t) \geq Z(t)$ for all $t = 0, 1, \dots, T$. ◇

Proof. From the definition of the Snell envelope, we immediately find that

$$Z(t) \geq \xi(t) \qquad \text{and} \qquad Z(t) \geq \mathbb{E}_t[Z(t+1)], \qquad t = 0, 1, \dots, T,$$

i.e. Z is indeed a supermartingale dominating ξ. If $Y = \{Y(t)\}_{t=0,1,\dots,T}$ is another supermartingale dominating ξ, it follows that

$$Y(T) \geq \xi(T) = Z(T).$$

Now continue inductively: We fix $t = 0, 1, \dots, T-1$, suppose that we have already argued that $Y(t+1) \geq Z(t+1)$, and proceed to show that this implies $Y(t) \geq Z(t)$. Using the definition of the Snell envelope, followed by the fact that $Y(t) \geq \xi(t)$ and the assumption that $Y(t+1) \geq Z(t+1)$, and finally the supermartingale property of Y, we find that

$$Z(t) = \xi(t) \vee \mathbb{E}_t[Z(t+1)] \leq Y(t) \vee \mathbb{E}_t[Y(t+1)] \leq Y(t),$$

which concludes the proof. □

With this, we are now ready to construct optimal stopping times and hence solve the optimal stopping problem.

Theorem 4.4 (Optimal Stopping I)**.** *For $t = 0, 1, \ldots, T$, define*

$$\tau_t^* \triangleq \inf\big\{s \in \mathcal{T}_t : Z(s) = \xi(s)\big\}. \tag{4.4}$$

Then $\tau_t^* \in \mathbb{T}_t$ and

$$Z(t) = \mathbb{E}_t\big[\xi\big(\tau_t^*\big)\big] = \sup_{\tau \in \mathbb{T}_t} \mathbb{E}_t[\xi(\tau)] = V(t). \tag{4.5}$$

In other words, τ_t^* is optimal for $V(t)$ and the Snell envelope coincides with the value process. ◇

Proof. It is immediate that $\tau_t^* \in \mathbb{T}_t$ since $Z(T) = \xi(T)$ and by Lemma A.7 (Hitting Times with Finite Time Index Sets). We prove Equation (4.5) by backward induction over t. For $t = T$, there is nothing to show since $\tau_T^* = T$ is the only stopping time in $\mathbb{T}_T$ and $Z(T) = \xi(T) = V(T)$. Now fix an arbitrary $t = 0, 1, \ldots, T-1$ and suppose that we have already established Equation (4.5) for $t+1$, i.e.

$$Z(t+1) = \mathbb{E}_{t+1}\big[\xi\big(\tau_{t+1}^*\big)\big] = \sup_{\tau \in \mathbb{T}_{t+1}} \mathbb{E}_{t+1}[\xi(\tau)] = V(t+1). \tag{4.6}$$

We have to verify the same equation for t.

Step 1: Let $\tau \in \mathbb{T}_t$ be an arbitrary stopping time and define $\sigma \triangleq \tau \vee (t+1)$, which is an element of $\mathbb{T}_{t+1}$ by Lemma A.3 (Maxima and Minima of Stopping Times). We observe that

$$\{\tau < t+1\} = \{\tau = t\} \in \mathfrak{F}(t) \qquad \text{and} \qquad \{\tau \geq t+1\} = \{\tau = t\}^c \in \mathfrak{F}(t).$$

With this, it follows that

$$\begin{aligned}\mathbb{E}_t[\xi(\tau)] &= \mathbb{E}_t\big[\xi(\tau)\mathbb{1}_{\{\tau<t+1\}}\big] + \mathbb{E}_t\big[\xi(\tau)\mathbb{1}_{\{\tau\geq t+1\}}\big] \\ &= \mathbb{E}_t\big[\xi(t)\mathbb{1}_{\{\tau<t+1\}}\big] + \mathbb{E}_t\big[\xi(\sigma)\mathbb{1}_{\{\tau\geq t+1\}}\big] \\ &= \xi(t)\mathbb{1}_{\{\tau<t+1\}} + \mathbb{E}_t\big[\xi(\sigma)\big]\mathbb{1}_{\{\tau\geq t+1\}} \\ &\leq Z(t)\mathbb{1}_{\{\tau<t+1\}} + \mathbb{E}_t\big[\mathbb{E}_{t+1}[\xi(\sigma)]\big]\mathbb{1}_{\{\tau\geq t+1\}}, \end{aligned} \tag{4.7}$$

where we have used the fact that the Snell envelope Z dominates ξ and the tower property of conditional expectation to obtain the last line. Since $\sigma \in \mathbb{T}_{t+1}$, it follows from Equation (4.6) and the supermartingale property of Z that we have

$$\mathbb{E}_t\big[\mathbb{E}_{t+1}[\xi(\sigma)]\big]\mathbb{1}_{\{\tau\geq t+1\}} \leq \mathbb{E}_t[Z(t+1)]\mathbb{1}_{\{\tau\geq t+1\}} \leq Z(t)\mathbb{1}_{\{\tau\geq t+1\}}. \tag{4.8}$$

Plugging this into Equation (4.7) yields[4]

$$\begin{aligned}\mathbb{E}_t[\xi(\tau)] \leq Z(t)\mathbb{1}_{\{\tau<t+1\}} &+ \mathbb{E}_t\big[\mathbb{E}_{t+1}[\xi(\sigma)]\big]\mathbb{1}_{\{\tau\geq t+1\}} \\ &\leq Z(t)\mathbb{1}_{\{\tau<t+1\}} + Z(t)\mathbb{1}_{\{\tau\geq t+1\}} = Z(t),\end{aligned}$$

and since $\tau \in \mathbb{T}_t$ was chosen arbitrarily, we find that

$$V(t) = \sup_{\tau\in\mathbb{T}_t} \mathbb{E}_t[\xi(\tau)] \leq Z(t).$$

Step 2: Since $\tau_t^* \in \mathbb{T}_t$, it follows that

$$\mathbb{E}_t\big[\xi\big(\tau_t^*\big)\big] \leq \sup_{\tau\in\mathbb{T}_t} \mathbb{E}_t[\xi(\tau)] = V(t),$$

and hence we are left with showing that

$$Z(t) = \mathbb{E}_t\big[\xi\big(\tau_t^*\big)\big].$$

For this, it suffices show that we have equality everywhere in step 1 if we replace τ by τ_t^*, which amounts to showing that Equation (4.7) and Equation (4.8) hold with equality for τ_t^*. For this, let us first observe that we have $\xi(t) = Z(t)$ on the event $\{\tau_t^* < t+1\} = \{\tau_t^* = t\}$ by definition of τ_t^*, which proves that we have equality in Equation (4.7). On the other

[4]Note that we would have equality everywhere if and only if Equation (4.7) and Equation (4.8) hold with equality.

hand, on the event $\{\tau_t^* \geq t+1\} \in \mathfrak{F}(t)$, we must have $Z(t) > \xi(t)$ and thus $\tau_t^* \vee (t+1) = \tau_t^* = \tau_{t+1}^*$ (by definition of τ_t^* and τ_{t+1}^*) and $Z(t) = \mathbb{E}_t[Z(t+1)]$ (by definition of Z). But then it follows from Equation (4.6) that

$$\begin{aligned}\mathbb{E}_t\big[\mathbb{E}_{t+1}[\xi(\tau_t^* \vee (t+1))]\big]\mathbb{1}_{\{\tau_t^* \geq t+1\}} &= \mathbb{E}_t\big[\mathbb{E}_{t+1}[\xi(\tau_t^* \vee (t+1))\mathbb{1}_{\{\tau_t^* \geq t+1\}}]\big] \\ &= \mathbb{E}_t\big[\mathbb{E}_{t+1}[\xi(\tau_{t+1}^*)\mathbb{1}_{\{\tau_t^* \geq t+1\}}]\big] \\ &= \mathbb{E}_t\big[\mathbb{E}_{t+1}[\xi(\tau_{t+1}^*)]\big]\mathbb{1}_{\{\tau_t^* \geq t+1\}} \\ &= \mathbb{E}_t\big[Z(t+1)\big]\mathbb{1}_{\{\tau_t^* \geq t+1\}} = Z(t)\mathbb{1}_{\{\tau_t^* \geq t+1\}},\end{aligned}$$

i.e. we also have equality in Equation (4.8) if we replace τ by τ_t^* and hence the proof is complete. □

The previous theorem implies that any optimal stopping problem on a finite probability space with a finite time index set can be solved, i.e. **there always exists an optimal stopping time**. The optimal stopping time need not be unique, but we know at least how to construct one: First compute the Snell envelope by backward induction, then, looking forward in time, stop as soon as the Snell envelope coincides with the reward process ξ for the first time.

We will discuss how to construct the **largest optimal stopping** time shortly and characterize the **set of optimal stopping times** as well. Before we can do so, however, we need some theoretical results on supermartingales first.

Theorem 4.5 (Doob Decomposition). *For any adapted stochastic process $Y = \{Y(t)\}_{t=0,1,\dots,T}$, there exists a martingale $M = \{M(t)\}_{t=0,1,\dots,T}$ and a predictable[5] process $A = \{A(t)\}_{t=0,1,\dots,T}$ with $M(0) = A(0) = 0$ such that*

$$Y(t) = Y(0) + M(t) + A(t), \qquad t = 0, 1, \dots, T.$$

Moreover, M and A are uniquely determined and given by

$$M(t) \triangleq \sum_{s=1}^{t}\Big[Y(s) - \mathbb{E}_{s-1}[Y(s)]\Big] \quad \text{and} \quad A(t) \triangleq \sum_{s=1}^{t}\Big[\mathbb{E}_{s-1}[Y(s)] - Y(s-1)\Big]$$

for all $t = 0, 1, \dots, T$. Finally, Y is a supermartingale if and only if A is decreasing. ◇

[5]By which we mean that $\{A(t)\}_{t=1,\dots,T}$ is predictable and $A(0)$ is $\mathfrak{F}(0)$-measurable, i.e. constant, to be consistent with Definition 2.9.

Proof. Step 1. The decomposition. Clearly, if we choose M and A as suggested by the theorem, we have

$$Y(0) + M(t) + A(t) = Y(0) + \sum_{s=1}^{t} \big[Y(s) - Y(s-1)\big] = Y(t).$$

Moreover, it is clear that M is adapted and A is predictable. To see that M is a martingale, we fix $t = 1, \ldots, T$ and simply observe that

$$\begin{aligned} \mathbb{E}_{t-1}\big[M(t)\big] &= \mathbb{E}_{t-1}\big[M(t-1) + Y(t) - \mathbb{E}_{t-1}[Y(t)]\big] \\ &= M(t-1) + \mathbb{E}_{t-1}[Y(t)] - \mathbb{E}_{t-1}[Y(t)] = M(t-1). \end{aligned}$$

Step 2. Uniqueness. Let $(\hat{M}, \hat{A})$ be another decomposition of Y with the same properties as (M, A). In particular,

$$Y(0) + M(t) + A(t) = Y(t) = Y(0) + \hat{M}(t) + \hat{A}(t), \qquad t = 0, 1, \ldots, T.$$

Rearranging this equation shows that

$$X(t) \triangleq M(t) - \hat{M}(t) = \hat{A}(t) - A(t), \qquad t = 0, 1, \ldots, T,$$

is a predictable martingale with $X(0) = 0$. But then, using first the predictability and then the martingale property,

$$X(t) = \mathbb{E}_{t-1}[X(t)] = X(t-1) = \ldots = X(0) = 0, \qquad t = 0, 1, \ldots, T.$$

But this is only possible if $M = \hat{M}$ and $A = \hat{A}$ by definition of X.

Step 3. The supermartingale case. We observe that, by the martingale property of M and predictability of A,

$$\mathbb{E}_{t-1}\big[\Delta Y(t)\big] = \mathbb{E}_{t-1}\big[\Delta M(t)\big] + \mathbb{E}_{t-1}\big[\Delta A(t)\big] = \Delta A(t), \qquad t = 1, \ldots, T.$$

From this, it follows immediately that Y is a supermartingale if and only if A is decreasing. □

Clearly, the Doob decomposition of a supermartingale is closely related the the optional decomposition in Theorem 3.17 as both yield decompositions into a sum of a martingale and a decreasing process. The main difference is that, in the Doob decomposition, the monotone process is **predictable**, whereas it is only **adapted** in the optional decomposition. On the other hand, the optional decomposition gives a more detailed characterization of the martingale component by identifying it as a stochastic integral, which is a martingale under **any** equivalent martingale measure.

The next result shows that the supermartingale property holds in between stopping times as well. A **warning** is in order though: Additional conditions are required for the next theorem to hold as soon as we go beyond finite probability spaces or finite time index sets.

We also note that we subsequently write $\mathbb{E}_\tau \triangleq \mathbb{E}[\cdot|\mathfrak{F}(\tau)]$ for any finite, i.e. $\{0, 1, \ldots, T\}$-valued, stopping time τ.

Theorem 4.6 (Optional Stopping (Discrete)). *Let $Y = \{Y(t)\}_{t=0,1,\ldots,T}$ be a supermartingale and let σ and τ be finite stopping times with $\sigma \leq \tau$. Then*

$$Y(\sigma) \geq \mathbb{E}_\sigma\big[Y(\tau)\big]$$

If Y is a martingale, this relation holds with equality. ◇

Proof. If we can prove the result for supermartingales, the martingale case follows immediately since then both Y and $-Y$ are supermartingales. Let us write

$$Y(\sigma) = Y(0) + \sum_{s=1}^{T} \Delta Y(s) \mathbb{1}_{\{s \leq \sigma\}} \quad \text{and} \quad Y(\tau) = Y(0) + \sum_{s=1}^{T} \Delta Y(s) \mathbb{1}_{\{s \leq \tau\}}.$$

From this representation it is easily seen that

$$Y(\tau) - Y(\sigma) = \sum_{s=1}^{T} \Delta Y(s) \mathbb{1}_{\{\sigma < s \leq \tau\}}.$$

Let us now fix $F \in \mathfrak{F}(\sigma)$ and $s \in \{1, \ldots, T\}$. Since

$$F \cap \{\sigma < s \leq \tau\} = \big(F \cap \{\sigma \leq s-1\}\big) \cap \{\tau \leq s-1\}^c \in \mathfrak{F}(s-1),$$

we see that

$$\mathbb{E}\big[\mathbb{1}_F \mathbb{1}_{\{\sigma < s \leq \tau\}} \Delta Y(s)\big] = \mathbb{E}\Big[\mathbb{1}_F \mathbb{1}_{\{\sigma < s \leq \tau\}} \mathbb{E}_{s-1}\big[\Delta Y(s)\big]\Big] \leq 0$$

by the supermartingale property of Y. But then

$$\mathbb{E}\Big[\mathbb{1}_F\mathbb{E}_\sigma\big[Y(\tau)-Y(\sigma)\big]\Big] = \mathbb{E}\big[\mathbb{1}_F\big(Y(\tau)-Y(\sigma)\big)\big]$$
$$= \sum_{s=1}^{T}\mathbb{E}\big[\mathbb{1}_F\mathbb{1}_{\{\sigma<s\leq\tau\}}\Delta Y(s)\big] \leq 0.$$

Choosing $F \triangleq \{\mathbb{E}_\sigma[Y(\tau)-Y(\sigma)]>0\}$ implies $\mathbb{P}[F]=0$ and hence

$$\mathbb{E}_\sigma\big[Y(\tau)-Y(\sigma)\big] \leq 0.$$

We conclude since $Y(\sigma)$ is $\mathfrak{F}(\sigma)$-measurable by Lemma A.11 (Stopped Process with Finite Time Index Set). □

The previous theorem allows us to characterize the set of supermartingales. Again, it should be noted that the statement of the following theorem has to be modified on infinite probability spaces and for infinite time index sets.

Theorem 4.7 (Optional Sampling (Discrete)). *Let $Y=\{Y(t)\}_{t=0,1,\dots,T}$ be an adapted process. Then Y is a supermartingale if and only if, for all finite stopping times σ,τ with $\sigma\leq\tau$, it holds that*

$$\mathbb{E}\big[Y(\sigma)\big] \geq \mathbb{E}\big[Y(\tau)\big].$$

In particular, if τ is an arbitrary stopping time and Y is a supermartingale, the stopped process $Y^\tau = Y(\cdot\wedge\tau)=\{Y(t\wedge\tau)\}_{t=0,1,\dots,T}$ is also a supermartingale. ◇

Proof. Let σ and τ be finite stopping times with $\sigma\leq\tau$. If Y is a supermartingale, Theorem 4.6 (Optional Stopping) shows that $Y(\sigma)\geq\mathbb{E}_\sigma[Y(\tau)]$, and hence

$$\mathbb{E}\big[Y(\sigma)\big] \geq \mathbb{E}\big[\mathbb{E}_\sigma[Y(\tau)]\big] = \mathbb{E}\big[Y(\tau)\big].$$

Now suppose that $\mathbb{E}[Y(\sigma)]\leq\mathbb{E}[Y(\tau)]$ for all finite stopping times σ,τ with $\sigma\leq\tau$. Given $t=1,\dots,T$ set $F\triangleq\{\mathbb{E}_{t-1}[\Delta Y(t)]>0\}\in\mathfrak{F}(t-1)$ and define

$$\sigma\triangleq t-1 \qquad \text{and} \qquad \tau\triangleq t\mathbb{1}_F+(t-1)\mathbb{1}_{F^c}.$$

The $\mathfrak{F}(t-1)$-measurability of F implies that τ is a stopping time since, for any $s = 0, 1, \dots, T$, we have $\{\tau \leq s\} = \emptyset \in \mathfrak{F}(s)$ if $s < t-1$, $\{\tau \leq s\} = \Omega \in \mathfrak{F}(s)$ if $s \geq t$, and

$$\{\tau \leq s\} = \{\tau \leq t-1\} = F^c \in \mathfrak{F}(t-1) \qquad \text{if } s = t-1.$$

Thus σ and τ are finite stopping times with $\sigma \leq \tau$ and hence

$$\begin{aligned}\mathbb{E}\big[\mathbb{1}_F \mathbb{E}_{t-1}[\Delta Y(t)]\big] &= \mathbb{E}\big[\mathbb{1}_F \Delta Y(t)\big] \\ &= \mathbb{E}\big[\mathbb{1}_F\big(Y(\tau) - Y(\sigma)\big)\big] = \mathbb{E}\big[Y(\tau) - Y(\sigma)\big] \leq 0,\end{aligned}$$

where we have used that $Y(\tau)\mathbb{1}_{F^c} = Y(t-1)\mathbb{1}_{F^c} = Y(\sigma)\mathbb{1}_{F^c}$ to obtain the third equality. But then $\mathbb{P}[F] = 0$, i.e. $\mathbb{E}_{t-1}[\Delta Y(t)] \leq 0$, and hence Y is a supermartingale. □

In the martingale case, the optional sampling theorem can be stated in a slightly different manner.

Corollary 4.8 (Optional Sampling for Martingales). *Consider an adapted process $M = \{M(t)\}_{t=0,1,\dots,T}$. Then M is a martingale if and only if*

$$\mathbb{E}\big[M(\tau)\big] = \mathbb{E}\big[M(T)\big] \qquad \textit{for all } \tau \in \mathbb{T}.$$

In particular, if τ is an arbitrary stopping time and M is a martingale, the stopped process $M^\tau = M(\cdot \wedge \tau) = \{M(t \wedge \tau)\}_{t=0,1,\dots,T}$ is also a martingale. ◇

Proof. Since both M and $-M$ are supermartingales if M is a martingale, it follows from optional sampling for supermartingales that M is a martingale if and only if

$$\mathbb{E}\big[M(\sigma)\big] = \mathbb{E}\big[M(\tau)\big] \qquad \text{for all } \tau, \sigma \in \mathbb{T} \text{ with } \sigma \leq \tau.$$

But since we have equality here, this is equivalent to

$$\mathbb{E}\big[M(T)\big] = \mathbb{E}\big[M(\tau)\big] \qquad \text{for all } \tau \in \mathbb{T}.$$ □

Let us now return to the optimal stopping problem

$$V(t) = \sup_{\tau \in \mathbb{T}_t} \mathbb{E}_t\big[\xi(\tau)\big], \qquad t = 0, 1, \dots, T.$$

We recall that V coincides with the Snell envelope Z of ξ, i.e. V satisfies

$$V(T) = \xi, \qquad V(t) = \xi(t) \vee \mathbb{E}_t\big[V(t+1)\big], \qquad t = 0, 1, \ldots, T-1, \tag{4.9}$$

and is the smallest supermartingale dominating ξ. The theoretical results on supermartingales can now be used to characterize the set of optimal stopping times.

Proposition 4.9 (Characterization of Optimal Stopping Times). *Fix a time index $t \in \{0, 1, \ldots, T\}$. A stopping time $\tau^* \in \mathbb{T}_t$ is optimal for $V(t)$ if and only if the truncated stopped value process $V^* = V(\cdot \wedge \tau^*) = \{V(s \wedge \tau^*)\}_{s \in \mathcal{T}_t}$ is a martingale and $V(\tau^*) = \xi(\tau^*)$.* ◇

Proof. Step 1. Suppose that $\tau^* \in \mathbb{T}_t$ is such that V^* is a martingale and $V(\tau^*) = \xi(\tau^*)$. Observe that $V(\tau^*) = V(T \wedge \tau^*) = V^*(T)$ and $V^*(t) = V(t \wedge \tau^*) = V(t)$. But then, by the martingale property of V^*,

$$\mathbb{E}_t\big[\xi(\tau^*)\big] = \mathbb{E}_t\big[V(\tau^*)\big] = \mathbb{E}_t\big[V^*(T)\big] = V^*(t) = V(t) = \sup_{\tau \in \mathbb{T}} \mathbb{E}\big[\xi(\tau)\big],$$

i.e. τ^* is indeed optimal for $V(t)$.

Step 2. Suppose that $\tau^* \in \mathbb{T}_t$ is optimal for $V(t)$. Then optimality, the fact that V dominates ξ, and finally Theorem 4.6 (Optional Stopping) imply

$$V(t) = \mathbb{E}_t\big[\xi(\tau^*)\big] \le \mathbb{E}_t\big[V(\tau^*)\big] \le V(t),$$

so that equality must hold everywhere. Since $\xi(\tau^*) \le V(\tau^*)$ pointwise, this is only possible if $\xi(\tau^*) = V(\tau^*)$ pointwise. To see that the truncated and stopped process $V^* = V(\cdot \wedge \tau^*) = \{V(s \wedge \tau^*)\}_{s \in \mathcal{T}_t}$ is a martingale, we first observe that Theorem 4.7 (Optional Sampling) implies that V^* is a supermartingale. Now fix $\sigma \in \mathbb{T}_t$ and observe that the supermartingale property (together with optional stopping/sampling) and the equality $V^*(t) = V(t) = \mathbb{E}_t[V(\tau^*)] = \mathbb{E}_t[V^*(T)]$ imply that

$$\mathbb{E}\big[V^*(\sigma)\big] \le \mathbb{E}\big[V^*(t)\big] = \mathbb{E}\big[\mathbb{E}_t[V^*(T)]\big] = \mathbb{E}\big[V^*(T)\big] \le \mathbb{E}\big[V^*(\sigma)\big].$$

Since this is only possible if equality holds everywhere, i.e.

$$\mathbb{E}\big[V^*(\sigma)\big] = \mathbb{E}\big[V^*(T)\big] \qquad \text{for all } \sigma \in \mathbb{T}_t,$$

it follows that V^* is a martingale by Corollary 4.8 (Optional Sampling for Martingales) and the proof is complete. □

As noted above, optimal stopping times need not be unique. The construction of one particular stopping time via the Snell envelope has already been discussed in Theorem 4.4 (Optimal Stopping I). We now present a second construction based on the Doob decomposition of the value process.

Theorem 4.10 (Optimal Stopping II). *For $t = 0, 1, \dots, T$, let us denote by $(M^t, A^t) = \{(M^t(s), A^t(s))\}_{s \in \mathcal{T}_t}$ the Doob decomposition of the truncated value process $V = \{V(s)\}_{s \in \mathcal{T}_t}$. Then*[6]

$$\rho_t^* \triangleq \inf\big\{s \in \mathcal{T}_t \setminus \{T\} : A^t(s+1) < 0\big\} \wedge T$$

defines an optimal stopping time for $V(t)$. ◇

Proof. By Theorem 4.5 (Doob Decomposition), A^t is a predictable decreasing process with $A^t(t) = 0$ since V^t is a supermartingale. Note that $A^t(s+1)$ is $\mathfrak{F}(s)$-measurable for each $s \in \mathcal{T}_t$, which guarantees that ρ_t^* is indeed a (finite) stopping time. Indeed, we have

$$\begin{aligned} \{\rho_t^* = s\} &= \{A^t(s) = 0\} \cap \{A^t(s+1) < 0\} \in \mathfrak{F}(s), \qquad s \in \mathcal{T}_t \setminus \{T\}, \\ \{\rho_t^* = T\} &= \{A^t(T-1) = 0\} \in \mathfrak{F}(T). \end{aligned}$$

For any $s \in \mathcal{T}_t \setminus \{T\}$, we observe that

$$0 \geq \mathbb{E}_s\big[\Delta V(s+1)\big] = \mathbb{E}_s\big[\Delta M^t(s+1)\big] + \mathbb{E}_s\big[\Delta A^t(s+1)\big] = \Delta A^t(s+1)$$

by the supermartingale property of V, the martingale property of M^t, and predictability of A^t. Hence

$$\big\{\Delta A^t(s+1) < 0\big\} = \big\{\mathbb{E}_s\big[\Delta V(s+1)\big] < 0\big\} = \big\{V(s) > \mathbb{E}_s\big[V(s+1)\big]\big\}.$$

But then, as V satisfies the Bellman principle in Equation (4.9), we find that

$$\big\{\Delta A^t(s+1) < 0\big\} = \big\{V(s) > \mathbb{E}_s\big[V(s+1)\big]\big\} \subseteq \big\{V(s) = \xi(s)\big\}.$$

This implies that $V(\rho_t^*) = \xi(\rho_t^*)$. Moreover, we obviously have

$$V\big(s \wedge \rho_t^*\big) = M\big(s \wedge \rho_t^*\big), \qquad s \in \mathcal{T}_t,$$

which implies that the truncated stopped process $\{V(s \wedge \rho_t^*)\}_{s \in \mathcal{T}_t}$ is a martingale. But then Proposition 4.9 (Characterization of Optimal Stopping Times) yields the optimality of ρ_t^* for $V(t)$. □

[6] By convention, we assume that $\inf \emptyset \triangleq +\infty$.

With Theorem 4.4 (Optimal Stopping I) and Theorem 4.10 (Optimal Stopping II), we now have two procedures to construct optimal stopping times. The construction in Theorem 4.4 is based on the idea of stopping as soon as the value process coincides with the reward ξ for the first time, whereas the construction in Theorem 4.10 is based on the idea of stopping just before the value process loses its martingale property and turns into an honest supermartingale. Since by Proposition 4.9 (Characterization of Optimal Stopping Times) these two properties characterize the set of all optimal stopping times, it should not come as a surprise that the stopping time in Theorem 4.4 is the **smallest** optimal stopping time, whereas the one in Theorem 4.10 is the **largest** optimal stopping time.

Corollary 4.11 (Smallest and Largest Optimal Stopping Times)**.** *Fix a time index $t = 0, 1, \dots, T$ and define two optimal stopping times for $V(t)$ via*

$$\tau_t^* \triangleq \inf\big\{s \in \mathcal{T}_t : V(s) = \xi(s)\big\}$$

and

$$\rho_t^* \triangleq \inf\big\{s \in \mathcal{T}_t \setminus \{T\} : A^t(s+1) < 0\big\} \wedge T,$$

respectively, where (M^t, A^t) denotes the Doob decomposition of $\{V(s)\}_{s\in\mathcal{T}_t}$. Then τ_t^* is the smallest and ρ_t^* the largest optimal stopping time for $V(t)$, i.e. if σ_t^* is optimal for $V(t)$, then $\tau_t^* \leq \sigma_t^* \leq \rho_t^*$. ◇

Proof. Let σ_t^* be optimal for $V(t)$. Then $\mathbb{P}[\sigma_t^* < \tau_t^*] = 0$ since otherwise $V(\sigma_t^*) \neq \xi(\sigma_t^*)$ with positive probability (by definition of τ_t^*), which would imply that σ_t^* is not optimal. Moreover, we have

$$V(s \wedge \sigma_t^*) = M^t(s \wedge \sigma_t^*) + A^t(s \wedge \sigma_t^*), \qquad s \in \mathcal{T}_t.$$

Since σ_t^* is optimal, $V(\cdot \wedge \sigma_t^*)$ must be a martingale and hence $A^t(\sigma_t^*) = 0$. But by definition of ρ_t^* this is only possible if $\sigma_t^* \leq \rho_t^*$. □

We conclude this section with an example showcasing that, in general, τ_t^* and ρ_t^* are not identical. To see this, it suffices to consider the case in which ξ is a **martingale**. Then optional stopping states that

$$\xi(t) = \mathbb{E}_t\big[\xi(\tau)\big] \qquad \text{for all } t = 0, 1, \dots, T \text{ and } \tau \in \mathbb{T}_t,$$

showing that in this case **any stopping time is optimal**. In particular, $\tau_t^* = t$ and $\rho_t^* = T$, and it holds that $V = \xi$.

4.2 Risk Neutral Pricing of American Options

Let us now apply the abstract optimal stopping results to compute prices of American options. Throughout this section, we take as given a finite financial market S which is free of arbitrage. To simplify the exposition, we restrict ourselves to complete markets, i.e. we assume that any European option can be replicated. We denote by $\mathbb{Q}$ the unique EMM.

Since we are in a complete market setting, we expect that the arbitrage free price of an American option $\xi = \{\xi(t)\}_{t=0,1,\dots,T}$ should be **uniquely determined**. From the discussion at the beginning of this chapter, we expect this price to be given by $P_{\mathsf{a}} = \{P_{\mathsf{a}}(t)\}_{t=0,1,\dots,T}$ given by

$$P_{\mathsf{a}}(t) = \sup_{\tau \in \mathbb{T}_t} S^0(t)\mathbb{E}_t^{\mathbb{Q}}\left[\frac{\xi(\tau)}{S^0(\tau)}\right], \qquad t = 0, 1, \dots, T. \tag{4.10}$$

Dividing both side by $S^0(t)$ shows that this is equivalent to

$$\overline{P}_{\mathsf{a}}(t) = \sup_{\tau \in \mathbb{T}_t} \mathbb{E}_t^{\mathbb{Q}}\left[\overline{\xi}(\tau)\right], \qquad t = 0, 1, \dots, T,$$

i.e. the discounted price $\overline{P}_{\mathsf{a}}$ is expected to be the **value process** arising from optimally stopping the discounted American option $\overline{\xi}$, or, equivalently, $\overline{P}_{\mathsf{a}}$ is the **Snell envelope** of $\overline{\xi}$. The aim of this section is to make this rigorous.

Our guess that $\overline{P}_{\mathsf{a}}$ is a Snell envelope should cause some mild **discomfort**. One the one hand, we have learned in the previous chapter that discounted arbitrage free prices are martingales, whereas the results of the previous section state that the Snell envelope is, in general, only a supermartingale. How can this possibly be? The answer to this question is simple, as the Snell envelope is a martingale up to the largest optimal stopping time and so it only turns into an honest supermartingale if the buyer of the option misses to **exercise optimally**. In particular, buying the option at a later date means there are less points in time available to exercise and hence the buyer of the option may not be willing to pay as much as at initial time zero.

To show that the unique arbitrage free price of an American option is given by Equation (4.10), we take on **two different perspectives**: the perspective of the buyer and the seller, respectively. Let us first argue that whenever the price of an American option is strictly smaller then the quantity on the right hand side of Equation (4.10), the **buyer** of the option is able to make a riskless profit.

Proposition 4.12 (Pricing American Options (Lower Bound)). *Let ξ be an American option with corresponding price process $P_{\mathsf{a}} = \{P_{\mathsf{a}}(t)\}_{t=0,1,\dots,T}$. If the extended market (S, P_{a}) is free of arbitrage opportunities, then we must have*

$$P_{\mathsf{a}}(t) \geq \sup_{\tau \in \mathbb{T}_t} S^0(t) \mathbb{E}_t^{\mathbb{Q}}\Big[\frac{\xi(\tau)}{S^0(\tau)}\Big], \qquad t = 0, 1, \dots, T. \qquad \diamond$$

Proof. For $t = 0, 1, \dots, T$ fixed, let us assume by contradiction that there exists an event $F \in \mathfrak{F}(t)$ with $\mathbb{P}[F] > 0$ such that

$$\overline{P}_{\mathsf{a}}(t) < \sup_{\tau \in \mathbb{T}_t} \mathbb{E}_t^{\mathbb{Q}}\big[\overline{\xi}(\tau)\big] \qquad \text{on } F.$$

Let us denote by τ^* an optimal stopping time for the optimal stopping problem on the right hand side, i.e. $\tau^* \in \mathbb{T}_t$ and

$$\sup_{\tau \in \mathbb{T}_t} \mathbb{E}_t^{\mathbb{Q}}\big[\overline{\xi}(\tau)\big] = \mathbb{E}_t^{\mathbb{Q}}\big[\overline{\xi}(\tau^*)\big].$$

For each $s = t, t+1, \dots, T$, completeness of the market implies that there exists a replication strategy $\varphi_s = \{\varphi_s(t)\}_{t=0,1,\dots,T}$ for the artificial European option with payoff $\xi(s)\mathbb{1}_{\{\tau^*=s\}}$, i.e. such that

$$\overline{X}^{\varphi_s}(T) = \overline{\xi}(s)\mathbb{1}_{\{\tau^*=s\}}.$$

Using that $\overline{X}^{\varphi_s}$ is a $\mathbb{Q}$-martingale, we note that

$$\overline{X}^{\varphi_s}(s) = \mathbb{E}_s^{\mathbb{Q}}\big[\overline{X}^{\varphi_s}(T)\big] = \mathbb{E}_s^{\mathbb{Q}}\big[\overline{\xi}(s)\mathbb{1}_{\{\tau^*=s\}}\big] = \overline{\xi}(s)\mathbb{1}_{\{\tau^*=s\}}$$

and therefore

$$\sum_{s=t}^{T} \overline{X}^{\varphi_s}(t) = \mathbb{E}_t^{\mathbb{Q}}\Big[\sum_{s=t}^{T} \overline{X}^{\varphi_s}(s)\Big] = \mathbb{E}_t^{\mathbb{Q}}\Big[\sum_{s=t}^{T} \overline{\xi}(s)\mathbb{1}_{\{\tau^*=s\}}\Big] = \mathbb{E}_t^{\mathbb{Q}}\big[\overline{\xi}(\tau^*)\big],$$

which is strictly bigger than $\overline{P}_{\mathsf{a}}(t)$ on the event F. Now consider the following trading strategy in the extended market (S, P_{a}):

- ▷ buy $\mathbb{1}_F$ units of the American option at time t; exercise this option at time τ^*;
- ▷ for each $s = t, t+1, \ldots, T$, follow the strategy $-\mathbb{1}_F\varphi_s$ from time t to s;
- ▷ buy $[\sum_{s=t}^{T} \overline{X}^{\varphi_s}(t) - \overline{P}_{\mathsf{a}}(t)]\mathbb{1}_F \geq 0$ units of S^0 at time t.

At any time $s = t, t+1, \ldots, T$, the investor makes a profit of $\mathbb{1}_F \xi(s) \mathbb{1}_{\{\tau^*=s\}}$ from exercising the option, which is exactly covers the losses due to liquidation of φ_s. Moreover, at time t, the strategy yields a profit of

$$\Big[\sum_{s=t}^{T} X^{\varphi_s}(t) - P_{\mathsf{a}}(t)\Big]\mathbb{1}_F = S^0(t)\Big[\sum_{s=t}^{T} \overline{X}^{\varphi_s}(t) - \overline{P}_{\mathsf{a}}(t)\Big]\mathbb{1}_F$$

from entering into the strategies $-\varphi_s$, $s = t, t+1, \ldots, T$, and buying the American option. Note that this profit is exactly the amount of money required to make the purchases of S^0 at time t. This shows that the trading strategy is self-financing. Moreover, it is clear that the strategy can be set up with an initial wealth of zero and the wealth at time T is given by $S^0(T)[\sum_{s=t}^{T} \overline{X}^{\varphi_s}(t) - \overline{P}_{\mathsf{a}}(t)]\mathbb{1}_F$ from the position in S^0, which is strictly positive on F. We have therefore constructed an arbitrage opportunity. □

To obtain an upper bound on arbitrage free prices of the American option ξ, we take on the perspective of the **seller of the option**, who has the task of hedging against all possible payoffs, i.e. any choice of exercise date. In line with the results on European options, a tight upper bound can be obtained in terms of the superhedging price.

Definition 4.13 ((American) Superhedging Strategy; Superhedging Price). Let $\xi = \{\xi(t)\}_{t=0,1,\ldots,T}$ be an American option and fix $t = 0, 1, \ldots, T$. A self-financing trading strategy φ is called time-t **superhedging strategy** for ξ provided that

$$X^{\varphi}(s) \geq \xi(s), \qquad s = t, t+1, \ldots, T.$$

With this, we refer to the process $\hat{P}_{\mathsf{a}} = \{\hat{P}_{\mathsf{a}}(t)\}_{t=0,1,\ldots,T}$ defined by

$$\hat{P}_{\mathsf{a}}(t) \triangleq \inf\big\{X^{\varphi}(t) : \varphi \text{ is a time-}t\text{ superhedging strategy for } \xi\big\}$$

for $t = 0, 1, \ldots, T$ as the superhedging price of ξ. ◇

We have to introduce the notion of time-t conditional superhedging strategies to guarantee that the superhedging price can indeed be attained. It can make a huge difference if an American option is to be superhedged on the the entire time period $\{0, 1, \ldots, T\}$ or just on a subset. Consider for example the American option ξ given by

$$\xi(0) = 1 \qquad \text{and} \qquad \xi(t) = 0, \qquad t = 1, \ldots, T.$$

The superhedging price of this option is clearly $\xi = \overline{\xi}$ itself, but despite the fact that we are in a complete market, there cannot be a self-financing trading strategy φ with $\overline{X}^{\varphi}(t) = \overline{\xi}(t)$ for all $t = 0, 1, \ldots, T$ since $\overline{\xi}$ is not a martingale under $\mathbb{Q}$.

Nevertheless, just as in the European case, the **superhedging price** provides an upper bound for the arbitrage free prices of American options and there is a time-t superhedging strategy which attains the superhedging price. For European options, this result was based on the optional decomposition of supermartingales. In the present setting, we can argue in a similar fashion using the **Doob decomposition**.

Theorem 4.14 (Superhedging of American Options)**.** *Let ξ be an American option with superhedging price $\hat{P}_{\mathsf{a}} = \{\hat{P}_{\mathsf{a}}(t)\}_{t=0,1,\ldots,T}$. For each $t = 0, 1, \ldots, T$, there exists a time-t superhedging strategy φ_t such that*

$$\hat{P}_{\mathsf{a}}(t) = X^{\varphi_t}(t) = \sup_{\tau \in \mathbb{T}_t} S^0(t) \mathbb{E}_t^{\mathbb{Q}}\Big[\frac{\xi(\tau)}{S^0(\tau)}\Big].$$

In other words, for each $t = 0, 1, \ldots, T$, there exists a time-t superhedging strategy which attains the superhedging price and the discounted superhedging price coincides with the value process of optimally stopping the discounted American option. ◇

Proof. Step 1. Consider the optimal stopping problem

$$V(t) \triangleq \sup_{\tau \in \mathbb{T}_t} \mathbb{E}_t^{\mathbb{Q}}\big[\overline{\xi}(\tau)\big], \qquad t = 0, 1, \ldots, T.$$

If φ is an arbitrary time-t superhedging strategy, we know that $\overline{X}^{\varphi}$ is a $\mathbb{Q}$-martingale (and hence also a $\mathbb{Q}$-supermartingale) which dominates $\overline{\xi}$ on

$\{t, t+1, \ldots, T\}$. But V is the smallest $\mathbb{Q}$-supermartingale which dominates $\overline{\xi}$ everywhere, and hence it follows as in Proposition 4.3 (Characterization of the Snell Envelope) that

$$\overline{X}^{\varphi}(t) \geq V(t) = \sup_{\tau \in \mathbb{T}_t} \mathbb{E}_t^{\mathbb{Q}}\big[\overline{\xi}(\tau)\big].$$

Multiplying both sides by $S^0(t)$ and using that φ was chosen arbitrarily, it follows that

$$\hat{P}_{\mathsf{a}}(t) \geq \sup_{\tau \in \mathbb{T}_t} S^0(t)\mathbb{E}_t^{\mathbb{Q}}\Big[\frac{\xi(\tau)}{S^0(\tau)}\Big].$$

Step 2. Fix $t = 0, 1, \ldots, T$ and denote by $(M^t, A^t) = \{(M^t(s), A^t(s))\}_{s \in \mathcal{T}_t}$ the Doob decomposition (under the measure $\mathbb{Q}$) associated with the trunctated value process $\{V(s)\}_{s \in \mathcal{T}_t}$. Since the market is complete, there exists a self-financing strategy $\varphi_t = \{\varphi_t(s)\}_{s=1,\ldots,T}$ which replicates the payoff

$$\big[V(t) + M^t(T)\big]S^0(T), \qquad \text{i.e.} \qquad \overline{X}^{\varphi_t}(T) = V(t) + M^t(T).$$

Let now $s = t, t+1, \ldots, T$ and observe that the $\mathbb{Q}$-martingale property of both $\overline{X}^{\varphi_t}$ and M^t guarantee that

$$\overline{X}^{\varphi_t}(s) = \mathbb{E}_s^{\mathbb{Q}}\big[\overline{X}^{\varphi_t}(T)\big] = V(t) + \mathbb{E}_s^{\mathbb{Q}}\big[M^t(T)\big] = V(t) + M^t(s).$$

Since V is a $\mathbb{Q}$-supermartingale, the process A^t is decreasing (and hence negative as $A^t(t) = 0$) and thus

$$\overline{X}^{\varphi_t}(s) = V(t) + M^t(s) \geq V(t) + M^t(s) + A^t(s) = V(s) \geq \overline{\xi}(s),$$

showing that φ_t is a time-t superhedging strategy against ξ. On the other hand, we have $M^t(t) = 0$ and hence

$$\overline{X}^{\varphi_t}(t) = V(t) + M^t(t) = V(t) = \sup_{\tau \in \mathbb{T}_t} \mathbb{E}_t^{\mathbb{Q}}\Big[\frac{\xi(\tau)}{S^0(\tau)}\Big].$$

In particular, since φ_t is a time-t superhedging strategy, it follows that

$$\hat{P}_{\mathsf{a}}(t) \leq X^{\varphi_t}(t) = \sup_{\tau \in \mathbb{T}_t} S^0(t)\mathbb{E}_t^{\mathbb{Q}}\Big[\frac{\xi(\tau)}{S^0(\tau)}\Big],$$

which completes the proof. □

Since the superhedging price can be attained, it is now straightforward to show that it is an upper bound on the the set of arbitrage free prices, which yields the following counterpart to Proposition 4.12 (Pricing American Options (Lower Bound)).

Proposition 4.15 (Pricing American Options (Upper Bound)). *Let ξ be an American option with corresponding price process $P_{\mathsf{a}} = \{P_{\mathsf{a}}(t)\}_{t=0,1,\dots,T}$. If the extended market (S, P_{a}) is free of arbitrage opportunities, then we must have*

$$P_{\mathsf{a}}(t) \leq \hat{P}_{\mathsf{a}}(t) = \sup_{\tau \in \mathbb{T}_t} S^0(t)\mathbb{E}_t^{\mathbb{Q}}\Big[\frac{\xi(\tau)}{S^0(\tau)}\Big], \qquad t = 0, 1, \dots, T. \qquad \diamond$$

Proof. We argue by contradiction and assume that there exists $t \in \{0, 1, \dots, T\}$ and an event $F \in \mathfrak{F}(t)$ with $\mathbb{P}[F] > 0$ such that

$$P_{\mathsf{a}}(t) > \hat{P}_{\mathsf{a}}(t) \qquad \text{on } F.$$

Denote by $\varphi_t = \{\varphi_t(s)\}_{s=1,\dots,T}$ the time-t superhedging strategy which attains the superhedging price, i.e. which satisfies $X^{\varphi_t}(t) = \hat{P}_{\mathsf{a}}(t)$. In the extended financial market (S, P_{a}), consider the trading strategy with initial wealth zero which

- ▷ sells $\mathbb{1}_F$ units of the American option at time t,
- ▷ follows the strategy $\mathbb{1}_F\varphi_t$ from time t up to the time $\tau \in \mathbb{T}_t$ at which the buyer of the American option decides to exercise,
- ▷ buys $\mathbb{1}_F[X^{\varphi_t}(\tau) - \xi(\tau)]/S^0(\tau)$ units of S^0 at time τ and holds this position up to time T, and
- ▷ buys $\mathbb{1}_F[P_{\mathsf{a}}(t) - \hat{P}_{\mathsf{a}}(t)]/S^0(t)$ units of S^0 at time t and holds this position up to time T.

The net cash flow (obtained from selling the option, entering φ_t, and buying S^0) at time t is given by

$$\mathbb{1}_F\Big[P_{\mathsf{a}}(t) - X^{\varphi_t}(t) - \big[P_{\mathsf{a}}(t) - \hat{P}_{\mathsf{a}}(t)\big]\frac{S^0(t)}{S^0(t)}\Big] = 0,$$

where we use that entering the strategy $\mathbb{1}_F\varphi_t$ at time t costs $X^{\varphi_t}(t)\mathbb{1}_F = \hat{P}_{\mathsf{a}}(t)\mathbb{1}_F$ units of money. At the exercise date of the American option, the net cash flow (obtained from delivering the payoff of the option to its buyer, exiting φ_t, and buying S^0) is given by

$$\mathbb{1}_F\Big[-\xi(\tau) + X^{\varphi_t}(\tau) - \big[X^{\varphi_t}(\tau) - \xi(\tau)\big]\frac{S^0(\tau)}{S^0(\tau)}\Big] = 0,$$

where we use that exiting the strategy $\mathbb{1}_F\varphi_t$ yields a profit of $X^{\varphi_t}(\tau)\mathbb{1}_F$. Finally, at time T, the investor is only left with a positive position in S^0, implying that the wealth at time T is given by

$$\mathbb{1}_F\big[X^{\varphi_t}(\tau) - \xi(\tau)\big]\frac{S^0(T)}{S^0(\tau)} + \mathbb{1}_F\big[P_{\mathsf{a}}(t) - \hat{P}_{\mathsf{a}}(t)\big]\frac{S^0(T)}{S^0(t)} \geq 0,$$

which is strictly positive on the event F. Indeed, the first summand is positive as φ is a time-t superhedging strategy and thus $X^{\varphi_t}(\tau) \geq \xi(\tau)$, whereas the second summand is strictly positive on F by assumption. We have thus constructed an arbitrage opportunity, which is the desired contradiction and completes the proof. □

Combining Proposition 4.12 and Proposition 4.15, we have thus argued that the only **possible** arbitrage free price P_{a} of an American option ξ in a complete market is given by

$$P_{\mathsf{a}}(t) = \hat{P}_{\mathsf{a}}(t) = \sup_{\tau\in\mathbb{T}_t} S^0(t)\mathbb{E}_t^{\mathbb{Q}}\Big[\frac{\xi(\tau)}{S^0(\tau)}\Big], \qquad t = 0, 1, \ldots, T.$$

However, we have **not argued yet** that this price is arbitrage free. Before we state the result, however, let us first have a closer look at how trading in markets with American options works.

Definition 4.16 ((Rational) Execution Strategy)**.** For $t = 0, 1, \ldots, T$, we say that a triple (η, τ, E) is a time-t **execution strategy** for an American option ξ if $\tau \in \mathbb{T}_t$, η is an $\mathfrak{F}(t)$-measurable and $\mathbb{R}$-valued random variable, and $E \in \mathfrak{F}(\tau)$. We say that (η, τ, E) is **rational** if there exists an optimal stopping time τ_t^* of

$$V(t) \triangleq \sup_{\sigma\in\mathbb{T}_t} \mathbb{E}_t^{\mathbb{Q}}\big[\overline{\xi}(\sigma)\big]$$

such that $\tau\mathbb{1}_E = \tau_t^*\mathbb{1}_E$ and $\tau \leq \rho_t^*$, where ρ_t^* is the largest optimal stopping time for $V(t)$. ◇

In the above definition, we think of η as the number of shares of ξ bought at time t. The stopping time τ corresponds to the time at which the position is liquidated. The event E is the event on which the buyer (which need not be the trader itself) of the η shares of the option liquidates the position by **exercising** the option, whereas E^c corresponds to the event on which the position is liquidated via **trading** the option in the market. Rational execution strategies correspond to those in which the buyer acts optimally in the sense of maximizing the value of the option. Note that no rational buyer will hold the option longer than ρ_t^* as this would decrease the value of the option. With this, we can now define the notion of simple dynamic trading strategies and corresponding simple arbitrage opportunities in markets with American options.

Definition 4.17 (Simple American Trading Strategies and Arbitrage)**.** For $t = 0, 1, \ldots, T$, let (η, τ, E) be a rational time-t execution strategy of an American option ξ and fix a predictable process $\varphi = \{\varphi(s)\}_{s=1,\ldots,T}$ taking values in $\mathbb{R}^{d+1}$. Denote by $\hat{P}_{\mathsf{a}} = \{\hat{P}_{\mathsf{a}}(s)\}_{s=0,1,\ldots,T}$ the superhedging price of ξ. We say that (φ, η, τ, E) is a **simple American trading strategy** provided that $(\varphi, \varphi_{\eta,\tau}) = \{(\varphi(s), \varphi_{\eta,\tau}(s))\}_{s=1,\ldots,T}$ is a self-financing trading strategy in the sense of Definition 2.10 in the extended market $(S, \hat{P}_{\mathsf{a}})$, where

$$\varphi_{\eta,\tau}(s) = \eta \mathbb{1}_{\{t+1 \leq s \leq \tau\}}, \qquad s = 1, \ldots, T. \tag{4.11}$$

We refer to (φ, η, τ, E) as a **simple American arbitrage** if $(\varphi, \varphi_{\eta,\tau})$ is an arbitrage opportunity in the sense of Definition 3.1 in the extended financial market $(S, \hat{P}_{\mathsf{a}})$. ◇

A simple American trading strategy (φ, η, τ, E) hence corresponds to a self-financing trading strategy in the extended financial market $(S, \hat{P}_{\mathsf{a}})$, in which the trader buys/sells η units the American option ξ at time t and liquidates this position at time τ. Observe that the liquidation of the position in the American option must necessarily occur at the price $\hat{P}_{\mathsf{a}}(\tau)$ since $\tau \mathbb{1}_E = \tau_t^* \mathbb{1}_E$ for an optimal stopping time τ_t^*, in which case we know that

$$\hat{P}_{\mathsf{a}}(\tau)\mathbb{1}_E = \hat{P}_{\mathsf{a}}(\tau_t^*)\mathbb{1}_E = \xi(\tau_t^*)\mathbb{1}_E$$

by Proposition 4.9 (Characterization of Optimal Stopping Times). On the other hand, on the event E^c, the option is traded in the market at at price of $\hat{P}_{\mathsf{a}}(\tau)$ and hence we **do not have to distinguish** between liquidation via trading and exercising.

Note that, with this, we **do not allow** any further dynamic trading in the American option except for liquidation of the initially acquired position η, i.e. after time t it is no longer possible to purchase or sell new shares of the American option. However, by taking linear combinations of finitely many simple trading strategies $(\varphi, \varphi_{\eta,\tau})$, we can clearly describe the entire range of dynamic trading strategies. The advantage of restricting to simple strategies and arbitrage, however, is that these are more **tractable** concepts as we are able to keep track of when the trader enters into a position when it is liquidated.

Theorem 4.18 (Risk Neutral Pricing of American Options)**.** *Let S be a complete market with equivalent martingale measure $\mathbb{Q}$. Let moreover ξ be an American option with superhedging price $\hat{P}_{\mathsf{a}} = \{\hat{P}_{\mathsf{a}}(t)\}_{t=0,1,\dots,T}$ and denote by $P_{\mathsf{a}} = \{P_{\mathsf{a}}(t)\}_{t=0,1,\dots,T}$ the price at which the option is traded. Then there is no simple American arbitrage in the extended market (S, P_{a}) if and only if*

$$P_{\mathsf{a}}(t) = \hat{P}_{\mathsf{a}}(t) = \sup_{\tau \in \mathbb{T}_t} S^0(t)\mathbb{E}_t^{\mathbb{Q}}\Big[\frac{\xi(\tau)}{S^0(\tau)}\Big], \qquad t = 0, 1, \dots, T. \tag{4.12}$$

Moreover, for each $t = 0, 1, \dots, T$, there is a time-t superhedging strategy $\varphi_t = \{\varphi_t(s)\}_{s=1,\dots,T}$ against ξ satisfying

$$P_{\mathsf{a}}(t) = \hat{P}_{\mathsf{a}}(t) = \sup_{\tau \in \mathbb{T}_t} S^0(t)\mathbb{E}_t^{\mathbb{Q}}\Big[\frac{\xi(\tau)}{S^0(\tau)}\Big] = X^{\varphi_t}(t). \qquad \diamond$$

Proof. By Proposition 4.12 (Pricing American Options (Lower Bound)) and Proposition 4.15 (Pricing American Options (Upper Bound)), we only have to show that if the price of ξ is given by Equation (4.12), then the extended market $(S, P_{\mathsf{a}}) = (S, \hat{P}_{\mathsf{a}})$ is free of simple American arbitrage. For this, we argue by contradiction and assume that for some $t = 0, 1, \dots, T$, there exists a rational time-t execution strategy (η, τ, E) and a predictable process φ such that (φ, η, τ, E) is a simple American arbitrage. This simply means that $(\varphi, \varphi_{\eta,\tau})$ with $\varphi_{\eta,\tau}$ given by Equation (4.11) is an arbitrage opportunity in the classical sense of Definition 3.1 in the extended financial market $(S, \hat{P}_{\mathsf{a}})$. The assumption that (η, τ, E) be rational corresponds to the assumption that the buyer acts optimally. In particular, by Proposition 4.9 (Characterization of Optimal Stopping Times) and writing

$$\overline{P}_{\mathsf{a}}(t) \triangleq \frac{\hat{P}_{\mathsf{a}}(t)}{S^0(t)} = \sup_{\sigma\in\mathbb{T}_t} \mathbb{E}_t^{\mathbb{Q}}\big[\overline{\xi}(\sigma)\big], \qquad t = 0, 1, \dots, T,$$

we find that $\overline{P}_{\mathsf{a}}(\cdot \wedge \tau) = \{\overline{P}_{\mathsf{a}}(s \wedge \tau)\}_{s=t,t+1,\dots,T}$ is a $\mathbb{Q}$-martingale since $\tau \leq \rho_t^*$. Recall that at time τ it does not matter if the buyer of the option exercises the option or sells it as $\overline{\xi}(\tau)\mathbb{1}_E = \overline{P}_{\mathsf{a}}(\tau)\mathbb{1}_E$ and hence both possibilities lead to the same profits. Moreover, we observe that

$$\begin{aligned}\overline{X}^{(\varphi,\varphi_{\eta,\tau})}(s) &= \varphi \bullet \overline{S}(s) + \varphi_{\eta,\tau} \bullet \overline{P}_{\mathsf{a}}(s)\\ &= \varphi \bullet \overline{S}(s) + \eta\mathbb{1}_{\{t+1\leq s\leq\tau\}} \bullet \overline{P}_{\mathsf{a}}(s), \qquad s = 0, 1, \dots, T,\end{aligned}$$

and since η is $\mathfrak{F}(t)$-measurable, it follows that $\overline{X}^{(\varphi,\varphi_{\eta,\tau})}$ is a $\mathbb{Q}$-martingale. In particular, we must have

$$0 = \overline{X}^{(\varphi,\varphi_{\eta,\tau})}(0) = \mathbb{E}^{\mathbb{Q}}\big[\overline{X}^{(\varphi,\varphi_{\eta,\tau})}(T)\big].$$

On the other hand, $(\varphi, \varphi_{\eta,\tau})$ is an arbitrage opportunity in $(S, \hat{P}_{\mathsf{a}}) = (S, P_{\mathsf{a}})$ and hence $\overline{X}^{(\varphi,\varphi_{\eta,\tau})}(T) \geq 0$ and $\mathbb{Q}[\overline{X}^{(\varphi,\varphi_{\eta,\tau})}(T) > 0] > 0$, leading to the contradiction

$$\mathbb{E}^{\mathbb{Q}}\big[\overline{X}^{(\varphi,\varphi_{\eta,\tau})}(T)\big] > 0. \qquad \square$$

We close this chapter on American options with an example.

Example 4.19 (Markovian American Options in the CRR Model)**.** Let $S = (S^0, S^1)$ be the CRR model with $-1 < d < r < u$. We already know that S is a complete market for these parameters. For simplicity, we assume that $r = 0$ so that discounted and undiscounted prices coincide. We consider an American option $\xi = \{\xi(t)\}_{t=0,1,\dots,T}$ of the form

$$\xi(t) \triangleq g\big(S^1(t)\big), \qquad t = 0, 1, \dots, T, \qquad \text{where } g : \mathbb{R} \to \mathbb{R}.$$

Our aim is to derive an algorithm to compute the unique arbitrage free price of ξ and to derive a superhedging strategy. For this, the first step is to compute the Snell envelope (under $\mathbb{Q}$) of the (discounted) American option

$\overline{\xi} = \xi$. Recall that the Snell envelope of $\overline{\xi}$ coincides with the value process of the optimal stopping problem for $\overline{\xi}$, which corresponds to the unique discounted arbitrage free price $\overline{P}_{\mathsf{a}}$ of ξ; see Theorem 4.18 (Risk Neutral Pricing of American Options). The Snell envelope can be computed via the recursion in Equation (4.9). For this, let us introduce a family of functions $\{p_t\}_{t=0,1,\ldots,T}$ defined recursively by $p_T \triangleq g$ and, for each $t = 0, 1, \ldots, T-1$,

$$p_t : \mathbb{R} \to \mathbb{R}, \quad s \mapsto p_t(s) \triangleq g(s) \vee \Big[q p_{t+1}\big((1+u)s\big) + (1-q)p_{t+1}\big((1+d)s\big)\Big].$$

We claim that

$$\overline{P}_{\mathsf{a}}(t) = p_t\big(S^1(t)\big), \qquad t = 0, 1, \ldots, T.$$

Indeed, we first observe that

$$\overline{P}_{\mathsf{a}}(T) = \overline{\xi}(T) = \frac{1}{S^0(T)} g\big(S^1(T)\big) = g\big(S^1(T)\big) = p_T\big(S^1(T)\big).$$

and, for each $t = 0, 1, \ldots, T-1$, assuming the result is valid for $t+1$, we find that

$$\begin{aligned}
\overline{P}_{\mathsf{a}}(t) &= \overline{\xi}(t) \vee \mathbb{E}_t^{\mathbb{Q}}\big[\overline{P}_{\mathsf{a}}(t+1)\big] \\
&= g\big(S^1(t)\big) \vee \mathbb{E}_t^{\mathbb{Q}}\big[p_{t+1}\big(S^1(t+1)\big)\big] \\
&= g\big(S^1(t)\big) \vee \Big[q p_{t+1}\big((1+u)S^1(t)\big) + (1-q)p_{t+1}\big((1+u)S^1(t)\big)\Big] \\
&= p_t\big(S^1(t)\big).
\end{aligned}$$

This yields a recursive algorithm to compute

$$P_{\mathsf{a}}(0) = \overline{P}_{\mathsf{a}}(0) = p_0(S^1(0)) = p_0(s).$$

An optimal stopping time corresponding to $\overline{P}_{\mathsf{a}}(0)$ is given by the first time that $\overline{P}_{\mathsf{a}}$ coincides with $\overline{\xi}$, which can be written as

$$\tau_0^* = \inf\big\{t \in \{0, 1, \ldots, T\} : g\big(S^1(t)\big) = p_t\big(S^1(t)\big)\big\}.$$

To compute a time-0 superhedging strategy, we first need to compute the Doob decomposition

$$\overline{P}_{\mathsf{a}}(t) = \overline{P}_{\mathsf{a}}(0) + M(t) + A(t), \qquad t = 0, 1, \ldots, T,$$

of the $\mathbb{Q}$-supermartingale $\overline{P}_{\mathsf{a}}$. As in the proof of Theorem 4.14 (Superhedging of American Options), a superhedging strategy can then be found by finding a replication strategy for the artificial payoff $\overline{P}_{\mathsf{a}}(0) + M(T)$. Recall from Theorem 4.5 (Doob Decomposition) that the martingale component $M = \{M(t)\}_{t=0,1,\ldots,T}$ in the Doob decomposition is given by $M(0) = 0$ and

$$\begin{aligned}\Delta M(t) &= \overline{P}_{\mathsf{a}}(t) - \mathbb{E}^{\mathbb{Q}}_{t-1}\big[\overline{P}_{\mathsf{a}}(t)\big] \\ &= p_t\big(S^1(t)\big) - \Big[q p_t\big((1+u)S^1(t-1)\big) + (1-q)p_t\big((1+d)S^1(t-1)\big)\Big] \\ &= \Big[p_t\big(S^1(t)\big) - p_t\big((1+d)S^1(t-1)\big)\Big]\frac{u}{u-d} \\ &\qquad\qquad + \Big[p_t\big((1+u)S^1(t-1)\big) - p_t\big(S^1(t)\big)\Big]\frac{d}{u-d} \\ &= \Big[p_t\big((1+u)S^1(t-1)\big) - p_t\big((1+d)S^1(t-1)\big)\Big]\frac{R_t}{u-d}\end{aligned}$$

for all $t = 0, 1, \ldots, T$. The superhedging strategy is now found by finding an $\mathfrak{F}(t-1)$-measurable random variable $\varphi(t)$ such that

$$\Delta M(t) = \varphi(t)\big[S^1(t) - S^1(t-1)\big],$$

since then $\overline{P}_{\mathsf{a}}(0) + M = \overline{P}_{\mathsf{a}}(0) + \varphi \bullet S^1 = \overline{P}_{\mathsf{a}}(0) + \varphi \bullet \overline{S}^1$. But choosing

$$\varphi(t) \triangleq \frac{p_t\big((1+u)S^1(t-1)\big) + p_t\big((1+d)S^1(t-1)\big)}{(u-d)S^1(t-1)}$$

and using that $S^1(t) - S^1(t-1) = R_t S^1(t-1)$ yields

$$\begin{aligned}&\varphi(t)\big[S^1(t) - S^1(t-1)\big] \\ &\qquad = \Big[p_t\big((1+u)S^1(t-1)\big) + p_t\big((1+d)S^1(t-1)\big)\Big]\frac{R_t}{u-d} = \Delta M(t). \qquad \diamond\end{aligned}$$

Chapter 5

Risk-Neutral Pricing in the Black-Scholes Model

It is now time to leave the world of finite probability spaces behind and turn towards more **complex models**. In particular, we shall eventually work on uncountable probability spaces and allow for trading at any time $t \in [0, T]$.

To keep the presentation simple, we restrict our analysis to one particularly tractable model: The famous **Black-Scholes model**, which is probably the most well-known and commonly used market models. To set up the model, we shall furthermore have the need to develop some **stochastic analysis**. In particular, we will learn about **Brownian motion**, which is one of the most important stochastic processes in continuous time, and we have to concern ourselves with **stochastic integration**.

Before we tackle these ambitious projects, however, let us begin with some motivation by studying the CRR model and the prices of European vanillas in the **continuous time limit**. We shall see that by proper rescaling, the model and the prices converge as we let the number of trading periods tend to infinity.

5.1 Continuous Time Limit of the CRR Model

We return to the CRR model, which we have seen to be arbitrage free and complete provided that $-1 < d < r < u$. The unique equivalent martingale measure in the CRR model is given by

$$\mathbb{Q}[\{\omega\}] = q^{U(\omega)}(1-q)^{D(\omega)}, \qquad \omega \in \Omega, \qquad \text{where } q \triangleq \frac{r-d}{u-d} \in (0,1),$$

where $U : \Omega \to \{0, 1, \dots, T\}$ and $D : \Omega \to \{0, 1, \dots, T\}$ count the number of up and down movements in each state of the world, i.e.

$$U(\omega) \triangleq \sum_{t=1}^{T} \mathbb{1}_{\{R_t(\omega)=u\}} \qquad \text{and} \qquad D(\omega) \triangleq \sum_{t=1}^{T} \mathbb{1}_{\{R_t(\omega)=d\}}, \qquad \omega \in \Omega.$$

Our aim is to see if this model admits a **continuous time limit** by letting the number of trading periods between time 0 and T tend to infinity. Of course, this is only possible if the parameters of the model are properly scaled.

Definition 5.1 (CRR(N) Model)**.** Given $N \in \mathbb{N}$, a riskless interest rate $r \in \mathbb{R}$, a volatility $\sigma > 0$, and an initial price $s > 0$, we define the **CRR(N) model** to be the classical CRR model $S_N = (S_N^0, S_N^1) = \{(S_N^0(t_k), S_N^1(t_k))\}_{k=0,1,\dots,NT}$, where $t_k \triangleq k/N$ for all $k = 0, 1, \dots, TN$ and the parameters r^N, u^N, d^N in this classical CRR model[1] are given by

$$r^N \triangleq \frac{r}{N}, \qquad u^N \triangleq \frac{r}{N} + \frac{\sigma}{\sqrt{N}}, \qquad d^N \triangleq \frac{r}{N} - \frac{\sigma}{\sqrt{N}}. \qquad \diamond$$

In other words, in the CRR(N) model, we have

$$S_N^0(t_k) = \left(1 + r^N\right)^k = \left(1 + \frac{r}{N}\right)^{Nt_k}, \qquad k = 0, 1, \dots, NT, \tag{5.1}$$

as well as

$$S_N^1(t_k) = s \prod_{i=1}^{k} \left(1 + R_N^i\right), \qquad k = 0, 1, \dots, NT,$$

where each R_N^i, $i = 0, 1, \dots, NT$, is $\{d^N, u^N\}$-valued with

$$\mathbb{P}[R_N^i = d^N] = \mathbb{P}\Big[R_N^i = \frac{r}{N} - \frac{\sigma}{\sqrt{N}}\Big] = 1 - p,$$
$$\mathbb{P}[R_N^i = u^N] = \mathbb{P}\Big[R_N^i = \frac{r}{N} + \frac{\sigma}{\sqrt{N}}\Big] = p.$$

The risk neutral probability q^N of the CRR(N) model is readily computed as

$$q^N = \frac{r^N - d^N}{u^N - d^N} = \frac{\sigma/\sqrt{N}}{2\sigma/\sqrt{N}} = \frac{1}{2}.$$

[1] We may have $d^N \leq -1$ for some N, in which case the CRR(N) model is not well-defined. However, $d^N \to 0$ as $N \to \infty$ and we are only interest in the behavior of S_N for large N anyway.

Theorem 5.2 (Weak Convergence of the CRR(N) Model)**.** *For $N \in \mathbb{N}$, denote by $S_N = (S_N^0, S_N^1)$ the CRR(N) model with riskless interest rate $r \in \mathbb{R}$, volatility $\sigma > 0$, and initial price $s > 0$, defined on the finite probability space $(\Omega_N, \mathfrak{A}_N, \mathbb{Q}_N)$, where $\mathbb{Q}_N$ denotes the equivalent martingale measure of the model. Then there exists a probability space $(\Omega, \mathfrak{A}, \mathbb{Q})$ on which there is defined a standard normal random variable Z such that*

$$\lim_{N\to\infty} S_N^0(T) = e^{rT} \qquad \text{and} \qquad \lim_{N\to\infty} S_N^1(T) = se^{(r-\frac{1}{2}\sigma^2)T+\sigma\sqrt{T}Z},$$

where the latter convergence is in distribution. ◇

Proof. The convergence of $S_N^0(T)$ is obvious from Equation (5.1), so we only have to concern ourselves with the convergence of $S_N^1(T)$. We denote by ξ_N the random variable which counts the number of up movements in the CRR(N) model, i.e.

$$\xi_N \triangleq \sum_{k=1}^{NT} \mathbb{1}_{\{R_N^k = u^N\}}, \qquad \text{so that} \qquad NT - \xi_N = \sum_{k=1}^{NT} \mathbb{1}_{\{R_N^k = d^N\}}.$$

We note that under $\mathbb{Q}_N$, the random variable ξ_N is binomially distributed on $\{0, 1, \ldots, NT\}$ with parameter $q^N = 1/2$. By the de Moivre-Laplace theorem (central limit theorem for binomial random variables), we therefore find a probability space $(\Omega, \mathfrak{A}, \mathbb{Q})$ on which there is defined a standard normal random variable Z such that the sequence $\{Z_N\}_{N\in\mathbb{N}}$ with

$$Z_N \triangleq \frac{\xi_N - \mathbb{E}^{\mathbb{Q}_N}[\xi_N]}{\sqrt{\mathrm{Var}^{\mathbb{Q}_N}[\xi_N]}} = \frac{\xi_N - \frac{1}{2}NT}{\frac{1}{2}\sqrt{NT}} = 2\frac{\xi_N}{\sqrt{NT}} - \sqrt{NT}, \qquad N \in \mathbb{N},$$

converges in distribution to Z. Now note that

$$S_N^1(T) = s\big(1 + u^N\big)^{\xi_N}\big(1 + d^N\big)^{NT-\xi_N}.$$

Using that

$$\frac{\xi_N}{\sqrt{N}} = \frac{1}{2}Z_N\sqrt{T} + \frac{1}{2}\sqrt{N}T$$

and writing

$$U^N \triangleq \big(1 + u^N\big)^{\sqrt{N}} = \Big(1 + \frac{\sigma}{\sqrt{N}} + \frac{r}{\sqrt{N}^2}\Big)^{\sqrt{N}},$$

$$D^N \triangleq \left(1+d^N\right)^{\sqrt{N}} = \left(1 - \frac{\sigma}{\sqrt{N}} + \frac{r}{\sqrt{N}^2}\right)^{\sqrt{N}},$$

we see that

$$\begin{aligned}\left(1+u^N\right)^{\xi_N} &= \left(U^N\right)^{\xi_N/\sqrt{N}} = \left(U^N\right)^{\frac{1}{2}Z_N\sqrt{T}+\frac{1}{2}\sqrt{N}T},\\ \left(1+d^N\right)^{NT-\xi_N} &= \left(D^N\right)^{\sqrt{N}T-\xi_N/\sqrt{N}} = \left(D^N\right)^{-\frac{1}{2}Z_N\sqrt{T}+\frac{1}{2}\sqrt{N}T}.\end{aligned}$$

From this representation, we find that

$$\left(1+u^N\right)^{\xi_N}\left(1+d^N\right)^{NT-\xi_N} = \left(\frac{U^N}{D^N}\right)^{\frac{1}{2}Z_N\sqrt{T}}\left(D^N U^N\right)^{\frac{1}{2}\sqrt{N}T}.$$

Now observe that

$$\lim_{N\to\infty} U^N = e^{\sigma} \quad \text{and} \quad \lim_{N\to\infty} D^N = e^{-\sigma}, \qquad \text{so that} \qquad \lim_{N\to\infty}\frac{U^N}{D^N} = e^{2\sigma}.$$

Moreover, we clearly have

$$\begin{aligned}\lim_{N\to\infty}\left(D^N U^N\right)^{\sqrt{N}} &= \lim_{N\to\infty}\left[\left(1-\frac{\sigma}{\sqrt{N}}+\frac{r}{\sqrt{N}^2}\right)\left(1+\frac{\sigma}{\sqrt{N}}+\frac{r}{\sqrt{N}^2}\right)\right]^N\\ &= \lim_{N\to\infty}\left[1+\frac{2r-\sigma^2}{N}+\frac{r^2}{N^2}\right]^N = e^{2r-\sigma^2}.\end{aligned}$$

Putting everything together, this implies that

$$\begin{aligned}\lim_{N\to\infty} S_N^1(T) &= s\lim_{N\to\infty}\left(1+u^N\right)^{\xi_N}\left(1+d^N\right)^{NT-\xi_N}\\ &= s\left(e^{2\sigma}\right)^{\frac{1}{2}Z\sqrt{T}}\left(e^{2r-\sigma^2}\right)^{\frac{1}{2}T} = se^{(r-\frac{1}{2}\sigma^2)T+\sigma\sqrt{T}Z},\end{aligned}$$

where the convergence is in distribution. □

Knowing that the CRR(N) model converges in distribution as $N \to \infty$, it follows immediately from the **Portmanteau theorem** that the time-0 arbitrage free prices of any European option given as a continuous and bounded function of $S_N^1(T)$ converges as well. In particular, this holds in the case of a European put and, using the put-call parity, the result can be extended to the European call as well.

Theorem 5.3 (Black-Scholes Formula). *For $N \in \mathbb{N}$, in the CRR(N) model with riskless interest rate $r \in \mathbb{R}$, volatility $\sigma > 0$, and initial price $s > 0$, denote by $C_N(0)$ and $P_N(0)$ the unique arbitrage free prices of a European call and put, respectively, with strike price K and maturity T. Then*

$$\lim_{N\to\infty} C_N(0) = s\Phi(d_+) - Ke^{-rT}\Phi(d_-),$$
$$\lim_{N\to\infty} P_N(0) = Ke^{-rT}\Phi(-d_-) - s\Phi(-d_+),$$

where $\Phi : \mathbb{R} \to [0,1]$ denotes the cumulative distribution function of a standard normal random variable and

$$d_+ \triangleq \frac{\log(s/K) + (r+\sigma^2/2)T}{\sigma\sqrt{T}} \quad \text{and} \quad d_- \triangleq \frac{\log(s/K) + (r-\sigma^2/2)T}{\sigma\sqrt{T}}. \qquad \diamond$$

Proof. Step 1. We show that it suffices to prove the result for the European put. Indeed, suppose that the convergence of $P_N(0)$ has already been established. Using Theorem 1.7 (Put-Call Parity), which holds since the CRR(N) model is free of arbitrage, it follows that

$$\begin{aligned}
\lim_{N\to\infty} C_N(0) &= \lim_{N\to\infty}\Big[P_N(0) + S_N^1(0) - K\frac{1}{S_N^0(T)}\Big] \\
&= Ke^{-rT}\Phi(-d_-) - s\Phi(-d_+) + s - Ke^{-rT} \\
&= s\big[1-\Phi(-d_+)\big] - Ke^{-rT}\big[1-\Phi(-d_-)\big] \\
&= s\Phi(d_+) - Ke^{-rT}\Phi(d_-),
\end{aligned}$$

which is the claim for the Euopean call.

Step 2. Convergence of the European put price. Recall that the put price in the CRR(N) model is given by

$$P_N(0) = \mathbb{E}^{\mathbb{Q}_N}\Big[\frac{(K - S_N^1(T))_+}{S_N^0(T)}\Big] = \frac{1}{S_N^0(T)}\mathbb{E}^{\mathbb{Q}_N}\Big[\big(K - S_N^1(T)\big)_+\Big],$$

where $\mathbb{Q}_N$ is the unique equivalent martingale measure. Now, using Theorem 5.2 (Weak Convergence of the CRR(N) Model) and since $x \mapsto (K-x)_+$ is bounded and continuous, it follows from the Portmanteau theorem that

$$\lim_{N\to\infty} P_N(0) = e^{-rT}\mathbb{E}^{\mathbb{Q}}\Big[\Big(K - se^{(r-\frac{1}{2}\sigma^2)T+\sigma\sqrt{T}Z}\Big)_+\Big],$$

where Z is a standard normal random variable on a suitable probability space $(\Omega, \mathfrak{A}, \mathbb{Q})$. To compute the expectation, let us set $\alpha \triangleq (r - \frac{1}{2}\sigma^2)T$, $\beta \triangleq \sigma\sqrt{T}$, and $\gamma \triangleq K/s$. Then

$$\mathbb{E}^{\mathbb{Q}}\left[\left(K - se^{(r-\frac{1}{2}\sigma^2)T+\sigma\sqrt{T}Z}\right)_+\right] = s\mathbb{E}^{\mathbb{Q}}\left[\left(\gamma - e^{\alpha+\beta Z}\right)_+\right].$$

We write

$$A \triangleq \{\gamma - e^{\alpha+\beta Z} > 0\} = \{Z < \kappa\}, \qquad \text{where } \kappa \triangleq \frac{\log(\gamma) - \alpha}{\beta}.$$

Then

$$\begin{aligned}
\mathbb{E}^{\mathbb{Q}}\left[\left(\gamma - e^{\alpha+\beta Z}\right)_+\right] &= \mathbb{E}^{\mathbb{Q}}\left[\mathbb{1}_A \gamma\right] - \mathbb{E}^{\mathbb{Q}}\left[\mathbb{1}_A e^{\alpha+\beta Z}\right] \\
&= \gamma\Phi(\kappa) - \int_{-\infty}^{\kappa} \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}z^2} e^{\alpha+\beta z} \mathrm{d}z \\
&= \gamma\Phi(\kappa) - e^{\alpha+\frac{1}{2}\beta^2} \int_{-\infty}^{\kappa} \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}(z-\beta)^2} \mathrm{d}z \\
&= \gamma\Phi(\kappa) - e^{\alpha+\frac{1}{2}\beta^2}\Phi(\kappa - \beta).
\end{aligned}$$

But then, using the definition of α, β, γ as well as d_-, d_+, we conclude that

$$\begin{aligned}
\lim_{N\to\infty} P_N(0) &= e^{-rT}\mathbb{E}^{\mathbb{Q}}\left[\left(K - se^{(r-\frac{1}{2}\sigma^2)T+\sigma\sqrt{T}Z}\right)_+\right] \\
&= se^{-rT}\left[\gamma\Phi(\kappa) - e^{\alpha+\frac{1}{2}\beta^2}\Phi(\kappa - \beta)\right] \\
&= Ke^{-rT}\Phi(-d_-) - s\Phi(-d_+).
\end{aligned}$$

□

5.2 Brownian Motion and Stochastic Integration

At this point, the question arises if we can define a **continuous time market model** which yields exactly the European call and put prices given by Theorem 5.3 (Black-Scholes Formula). The starting point for such a model is of course Theorem 5.2 (Weak Convergence of the CRR(N) Model). A guess for such a model would be $S = (S^0, S^1) = \{(S^0(t), S^1(t))\}_{t\in[0,T]}$ with risk neutral dynamics

$$S^0(t) \triangleq e^{rt} \qquad \text{and} \qquad S^1(t) \triangleq se^{(r-\frac{1}{2}\sigma^2)t+\sigma W(t)}, \qquad t \in [0, T], \tag{5.2}$$

for a suitable process $W = \{W(t)\}_{t\in[0,T]}$. What kind of properties should this process have? Given $s, t \in [0, T]$ with $s < t$, we may repeat the argument in the proof of Theorem 5.2 (Weak Convergence of the CRR(N) Model) to find that[2]

$$\lim_{N\to\infty} \frac{S_N^1(t)}{S_N^1(s)} = e^{(r-\frac{1}{2}\sigma^2)(t-s)+\sigma\sqrt{t-s}Z_{s,t}} \qquad \text{in distribution}$$

for a standard normal random variable $Z_{s,t}$. From this and comparing with Equation (5.2), we see that the process W should have the property that

$$W(t) - W(s) \text{ is normally distributed with mean zero and variance } t - s.$$

Moreover, $S_N^1(t)/S_N^1(s)$ does not depend on the any random variable realized before time s, from which we conclude that W should be such that

$$W(t) - W(s) \text{ is independent of } \overline{\mathfrak{F}}^W(s) \triangleq \sigma\big(W(r) : r \in [0, s]\big).$$

Next, we note that $S^1(0) = s$ can only hold if $W(0) = 0$ and finally it seems reasonable to assume that $t \mapsto S^1(t)$ is continuous, which translates into the requirement that

$$t \mapsto W(t, \omega) \text{ is continuous for all } \omega \in \Omega.$$

A process with these properties is given a special name.

Definition 5.4 (Wiener Process, Brownian Motion). A stochastic process $W = \{W(t)\}_{t\in[0,T]}$ with natural filtration denoted by $\overline{\mathfrak{F}}^W = \{\overline{\mathfrak{F}}^W(t)\}_{t\in[0,T]}$ is called **Wiener process** or **Brownian motion** if

- **(W1)** $W(0) = 0$ almost surely;
- **(W2*)** for each $s, t \in [0, T]$ with $s < t$, the increment $W(t) - W(s)$ is independent of $\overline{\mathfrak{F}}^W(s)$;
- **(W3)** for each $s, t \in [0, T]$ with $s < t$, the increment $W(t) - W(s)$ is normally distributed with mean zero and variance $t - s$;
- **(W4)** the mapping $t \mapsto W(t, \omega)$ is continuous on $[0, T]$ for each $\omega \in \Omega$. ◇

[2] At least formally by considering CRR(N) models with initial time zero replaced by s and terminal time T replaced by t.

A few remarks are in order here: First of all, the probability space $(\Omega, \mathfrak{A}, \mathbb{P})$ on which Brownian motion is defined must necessarily be infinite (even uncountable) since finite probability spaces do not support sequences of independent random variables. This forces us to leave the realm of finite probability spaces. In particular, we subsequently have to be more careful about **measurability** and **integrability** of random variables and there are non-empty nullsets, i.e. there are non-empty events $N \in \mathfrak{A}$ with $\mathbb{P}[N] = 0$. We denote by $\mathcal{N}$ the collection of all such nullsets, i.e.

$$\mathcal{N} \triangleq \{N \in \mathfrak{A} : \mathbb{P}[N] = 0\}.$$

With this, we now introduce a new filtration $\mathfrak{F}^W = \{\mathfrak{F}^W(t)\}_{t\in[0,T]}$ by setting

$$\mathfrak{F}^W(t) \triangleq \bigcap_{s\in(t,T]} \sigma\big(\overline{\mathfrak{F}}^W(s) \cup \mathcal{N}\big), \quad t \in [0, T), \quad \mathfrak{F}^W(T) \triangleq \sigma\big(\overline{\mathfrak{F}}^W(T) \cup \mathcal{N}\big).$$

It is not difficult to see that $\mathfrak{F}^W$ is the smallest filtration to which W is **adapted** and such that it is **complete** in that

$$\mathcal{N} \subseteq \mathfrak{F}^W(t), \qquad t \in [0, T],$$

and **right continuous** in the sense that

$$\mathfrak{F}^W(t) = \bigcap_{s\in(t,T]} \mathfrak{F}^W(s), \qquad t \in [0, T).$$

The filtration $\mathfrak{F}^W$ is referred to as the **augmented filtration** generated by W. It is readily seen that

(W2) for each $s, t \in [0, T]$ with $s < t$, the increment $W(t) - W(s)$ is independent of $\mathfrak{F}^W(s)$;

Finally, it is not difficult to show that $\mathfrak{F}^W(0)$ is $\mathbb{P}$-trivial, i.e. $\mathbb{P}[F] \in \{0, 1\}$ for all $F \in \mathfrak{F}^W(0)$. This result is a consequence of Kolmogorov's 0/1 law and typically referred to as **Blumenthal's 0/1 law**. In particular, since $\mathcal{N} \subset \mathfrak{F}^W(0)$ by construction, it follows that $\mathfrak{F}^W(0) = \sigma(\mathcal{N})$. Notice that this implies that any $\mathfrak{F}^W(0)$-measurable random variable is **almost surely constant**.

Convention: We subsequently identify random variables which coincide almost surely, i.e. we identify two random variables X and Y defined on the same probability space provided that $\mathbb{P}[X = Y] = 1$. Practically, this means that we **drop any almost sure qualifiers** when writing down equations involving random variables. Note that if $X = \{X(t)\}_{t\in[0,T]}$ and $Y = \{Y(t)\}_{t\in[0,T]}$ are stochastic processes, we can treat them as path-valued random variables and hence $X = Y$ means that almost surely all paths of X and Y coincide. Completeness of the filtration $\mathfrak{F}^W$ guarantees that we do not run into measurability issues here in the sense that if X is $\mathfrak{F}^W$-adapted and $X = Y$, then Y must also be $\mathfrak{F}^W$-adapted as $\{X(t) \neq Y(t)\} \in \mathcal{N} \subseteq \mathfrak{F}^W(t)$ for all $t \in [0, T]$.

Warning: Before proceeding, let us highlight a very important point: It is by no means clear if a Brownian motion even exists. Indeed, proving **existence of Brownian motion** is not straightforward, mainly due to the requirement that the paths $t \mapsto W(t, \omega)$ are continuous. As the proof is out of the scope of these notes, we subsequently simply take existence of Brownian motion for granted. Note, however, that Theorem 5.2 (Weak Convergence of the CRR(N) Model) is already a strong indicator that Brownian motion really does exist and pretty much every textbook on stochastic processes can be consulted to confirm this.

Let us argue that Brownian motion is a **square integrable martingale**. For this, we subsequently fix a filtered probability space $(\Omega, \mathfrak{A}, \mathfrak{F}, \mathbb{P})$ which is sufficiently rich to support a Brownian motion $W = \{W(t)\}_{t\in[0,T]}$. We furthermore assume that $\mathfrak{F} = \mathfrak{F}^W$ is the augmented filtration generated by W and we assume that $\mathfrak{A} = \mathfrak{F}(T) = \mathfrak{F}^W(T)$.

Definition 5.5 ((Square Integrable) Martingale)**.** An adapted process $M = \{M(t)\}_{t\in[0,T]}$ is called **martingale** if $\mathbb{E}[|M(t)|] < \infty$ for all $t \in [0, T]$ and

$$M(s) = \mathbb{E}_s[M(t)], \qquad s, t \in [0, T],\ s < t.$$

M is called **square integrable** if $\mathbb{E}[|M(t)|^2] < \infty$ for all $t \in [0, T]$ and the space of continuous square integrable martingales is denoted by $\mathcal{M}^2$. ◇

On finite probability spaces, this definition is equivalent to Definition 3.2 (Martingales on Finite Probability Spaces). The definition here accounts for the fact that conditional expectations of a random variable exist if and only if the random variable is integrable. Let us also highlight that, in light of our convention, the equality $M(s) = \mathbb{E}_s[M(t)]$ is to be read as $M(s)$ coincides **almost surely** with $\mathbb{E}_s[M(t)]$. We also note that **sub- and supermartingales** are defined in an obvious analogous way.

Lemma 5.6 (Martingale Property of Brownian Motion)**.** *Brownian motion is a continuous square integrable martingale, i.e.* $W \in \mathcal{M}^2$. ◇

Proof. Let $s, t \in [0, T]$ with $s < t$. We first note that $W(t) = W(t) - W(0)$, from which we conclude that $W(t)$ is normally distributed with variance t and hence, in particular, square integrable. Moreover, W is a martingale since the properties **(W2)** (independent increments) and **(W3)** (normally distributed increments) show that

$$\mathbb{E}_s[W(t) - W(s)] = \mathbb{E}[W(t) - W(s)] = 0. \qquad \square$$

The space $\mathcal{M}^2$ of continuous square integrable martingales can be turned into a Hilbert space by introducing the **scalar product**

$$\langle \cdot | \cdot \rangle_{\mathcal{M}^2} : \mathcal{M}^2 \times \mathcal{M}^2 \to \mathbb{R}, \qquad (M, N) \mapsto \langle M | N \rangle_{\mathcal{M}^2} \triangleq \mathbb{E}\big[M(T)N(T)\big].$$

Note that $\langle \cdot | \cdot \rangle_{\mathcal{M}^2}$ can be identified with the scalar product on the Hilbert space of square intregrable $\mathfrak{A} = \mathfrak{F}(T)$-measurable random variables. Also, once again, we make use of our convention of identifying random variables which coincide almost surely to guarantee that $\langle \cdot | \cdot \rangle_{\mathcal{M}^2}$ is indeed a scalar product. The norm on $\mathcal{M}^2$ induced by this scalar product is given by

$$\| \cdot \|_{\mathcal{M}^2} : \mathcal{M}^2 \to \mathbb{R}_+, \quad M \mapsto \|M\|_{\mathcal{M}^2} \triangleq \big(\langle M | M \rangle_{\mathcal{M}^2}\big)^{1/2} = \mathbb{E}\big[|M(T)|^2\big]^{1/2}.$$

At first sight, it should come as a surprise that $(M, \| \cdot \|_{\mathcal{M}^2})$ is a Hilbert space as $\| \cdot \|_{\mathcal{M}^2}$ only takes into account the terminal values of martingales, i.e. if $\{M_k\}_{k \in \mathbb{N}}$ is a Cauchy sequence in $\mathcal{M}^2$, it is by no means clear if there exists a **continuous** square integrable martingale M such that $M_k \to M$ in $(\mathcal{M}^2, \| \cdot \|_{\mathcal{M}^2})$. It is, however, obvious that $M_k(T) \to M(T)$ in the Hilbert space of square integrable random variables. The convergence of the entire process can be argued for using Theorem B.3 (Doob's L^p-Inequality), which for $p = 2$ states that

$$\mathbb{E}\Big[\sup_{t \in [0,T]} |M(t)|^2\Big]^{1/2} \leq 2\mathbb{E}\big[|M(T)|^2\big]^{1/2} = 2\|M\|_{\mathcal{M}^2}, \qquad M \in \mathcal{M}^2.$$

If you have never seen a Doob inequality before, now is a good time to have a look at Appendix B.

Theorem 5.7 ($\mathcal{M}^2$ as a Hilbert Space)**.** *The space $\mathcal{M}^2$ equipped with the scalar product $\langle\cdot|\cdot\rangle_{\mathcal{M}^2}$ and its induced norm $\|\cdot\|_{\mathcal{M}^2}$ is a Hilbert space.* ◇

Proof. It is clear that $\mathcal{M}^2$ is a linear space and that $\langle\cdot|\cdot\rangle_{\mathcal{M}^2}$ furnishes a scalar product on $\mathcal{M}^2$. Therefore, we only have to show that $(\mathcal{M}^2, \|\cdot\|_{\mathcal{M}^2})$ is complete, i.e. if $\{M_k\}_{k\in\mathbb{N}} \subset \mathcal{M}^2$ is a Cauchy sequence with respect to $\|\cdot\|_{\mathcal{M}^2}$, there exists $M \in \mathcal{M}^2$ such that $M_k \to M$ with respect to $\|\cdot\|_{\mathcal{M}^2}$. By passing to a subsequence, we may assume that

$$\|M_{k+1} - M_k\|_{\mathcal{M}^2} \leq 2^{-(k+1)}, \qquad k \in \mathbb{N}.$$

By Theorem B.3 (Doob's L^p-Inequality), it follows that

$$\begin{aligned}\mathbb{E}\Big[\sup_{t\in[0,T]} |M_{k+1}(t) - M_k(t)|^2\Big]^{1/2} &\leq 2\mathbb{E}\big[|M_{k+1}(T) - M_k(T)|^2\big]^{1/2} \\ &= 2\|M_{k+1} - M_k\|_{\mathcal{M}^2} \leq 2^{-k}, \qquad k \in \mathbb{N}.\end{aligned}$$

By monotone convergence, this implies that

$$\begin{aligned}&\mathbb{E}\Big[\sum_{k=1}^{\infty} \sup_{t\in[0,T]} |M_{k+1}(t) - M_k(t)|^2\Big] \\ &\qquad = \lim_{n\to\infty} \mathbb{E}\Big[\sum_{k=1}^{n} \sup_{t\in[0,T]} |M_{k+1}(t) - M_k(t)|^2\Big] < \infty.\end{aligned}$$

But then, for $\mathbb{P}$-almost every $\omega \in \Omega$, the sequence $\{M_k(\cdot,\omega)\}_{k\in\mathbb{N}}$ is a Cauchy sequence in the Banach space of continuous functions equipped with the norm of uniform convergence on $[0, T]$. This implies that there exists a continuous process $M = \{M(t)\}_{t\in[0,T]}$ such that

$$\lim_{k\to\infty} \sup_{t\in[0,T]} |M_k(t) - M(t)| = 0.$$

From this and completeness of $\mathfrak{F}^W$, it follows that M is adapted. Fatou's lemma on the other hand implies that

$$\mathbb{E}\big[|M_k(t) - M(t)|^2\big]^{1/2} \leq \liminf_{n\to\infty} \mathbb{E}\big[|M_k(t) - M_n(t)|^2\big]^{1/2}$$

$$\leq \liminf_{n\to\infty} \sum_{\ell=k}^{n} \mathbb{E}\Big[|M_{\ell+1}(t) - M_\ell(t)|^2\Big]^{1/2}$$
$$\leq \liminf_{n\to\infty} \sum_{\ell=k}^{n} 2^{-\ell} = 2^{-(k-1)}$$

for every $k \in \mathbb{N}$ and $t \in [0, T]$, and hence

$$\lim_{k\to\infty} \mathbb{E}\Big[|M_k(t) - M(t)|^2\Big]^{1/2} = 0, \qquad t \in [0, T]. \tag{5.3}$$

In particular, M is square integrable and we have

$$\lim_{k\to\infty} \|M_k - M\|_{\mathcal{M}^2} = \lim_{k\to\infty} \mathbb{E}\Big[|M_k(T) - M(T)|^2\Big]^{1/2} = 0,$$

so it only remains to check if M is a martingale. For this, let us fix $s, t \in [0, T]$ with $s < t$ and let $F \in \mathfrak{F}(s)$. Then the convergence in Equation (5.3) and the martingale property of each M_k, $k \in \mathbb{N}$, imply that

$$\begin{aligned}\mathbb{E}\big[\mathbb{1}_F M(t)\big] = \lim_{k\to\infty} \mathbb{E}\big[\mathbb{1}_F M_k(t)\big] &= \lim_{k\to\infty} \mathbb{E}\big[\mathbb{1}_F \mathbb{E}_s[M_k(t)]\big] \\ &= \lim_{k\to\infty} \mathbb{E}\big[\mathbb{1}_F M_k(s)\big] = \mathbb{E}\big[\mathbb{1}_F M(s)\big],\end{aligned}$$

i.e. $M(s) = \mathbb{E}_s[M(t)]$, so that M is indeed a martingale. □

Our next aim is to define a **stochastic integral** with respect to Brownian motion which generalizes the discrete time stochastic integral in Definition 2.11. For this, we first need to generalize the notion of predictability to processes with time index set $[0, T]$.

Definition 5.8 (Predictable σ-Field; Predictable)**.** The σ-field $\mathcal{P}$ on $(0, T] \times \Omega$ which is generated by all adapted left-continuous processes $X = \{X(t)\}_{t\in[0,T]}$ treated as mappings

$$X : (0, T] \times \Omega \to \mathbb{R}, \qquad (t, \omega) \mapsto X(t, \omega),$$

is referred to as the **predictable σ-field**. With this, a real-valued process $Y = \{Y(t)\}_{t\in[0,T]}$ is called **predictable** if

$$Y : (0, T] \times \Omega \to \mathbb{R}, \qquad (t, \omega) \mapsto Y(t, \omega), \qquad \text{is } \mathcal{P}\text{-measurable}$$

and if $Y(0)$ is $\mathfrak{F}(0)$-measurable. ◇

Just as in the case of a finite time index set, one can argue that any predictable process is also **adapted**. The proof uses a monotone class argument. We also note here that the predictable σ-field depends on the choice of filtration.

Let us now turn towards stochastic integration. As usual when constructing integrals, we begin with stochastic integrals for **simple integrands**.

Definition 5.9 (Simple Integrand)**.** A process $H = \{H(t)\}_{t\in[0,T]}$ which can be written in the form

$$H(t) = \sum_{k=1}^{n} H_k \mathbb{1}_{(t_{k-1},t_k]}(t), \qquad t \in [0,T], \tag{5.4}$$

for $n \in \mathbb{N}$ and suitable constants $0 = t_0 < t_1 < \ldots < t_n = T$ and bounded $\mathfrak{F}(t_{k-1})$-measurable random variables H_k, $k = 1, \ldots, n$, is called **simple integrand**. The linear space of all simple integrands is denoted by $\mathcal{S}$. ◇

Note that every simple integrand H is **predictable** as it is left continuous and adapted. Moreover, we observe that the paths $t \mapsto H(t,\omega)$ are just **step functions** on $[0,T]$. It is thus straightforward to define the stochastic integral for simple integrands.

Definition 5.10 (Itō Integral for Simple Integrands)**.** Let $H \in \mathcal{S}$ be a simple integrand represented as in Equation (5.4). Then we define the **Itō Integral** of H with respect to W as the process $H \bullet W = \{H \bullet W(t)\}_{t\in[0,T]}$ given by

$$H \bullet W(t) \triangleq \sum_{k=1}^{n} H_k \big[W(t \wedge t_k) - W(t \wedge t_{k-1})\big], \qquad t \in [0,T]. \qquad \diamond$$

One may think of the Itō integral $H \bullet W$ as the process obtained by **scaling the increments** $W(t_k) - W(t_{k-1})$ of the Brownian motion with a factor of H_k. It is clear from the definition that the Itō integral is **linear** in that

$$\big(\alpha H + K\big) \bullet W = \alpha\big(H \bullet W\big) + K \bullet W, \qquad H, K \in \mathcal{S},\ \alpha \in \mathbb{R}.$$

The next result shows that the Itō integral is a continuous square integrable martingale.

Proposition 5.11 (Itō Integral in $\mathcal{M}^2$)**.** *For each $H \in \mathcal{S}$, the Itō integral $H \bullet W$ is a continuous square integrable martingale, i.e. $H \bullet W \in \mathcal{M}^2$.* ◇

Proof. Suppose that H is represented as in Equation (5.4). It is clear that $H \bullet W$ is adapted and continuous. Moreover, if we denote by $\alpha > 0$ a mutual upper bound of $|H_k|$, $k = 1, \dots, n$, it follows that

$$\big|H \bullet W(t)\big| = \Big|\sum_{k=1}^{n} H_k\big[W(t \wedge t_k) - W(t \wedge t_{k-1})\big]\Big| \le 2\alpha n \sup_{s\in[0,T]} |W(s)|$$

for all $t \in [0, T]$. But then Theorem B.3 (Doob's L^p-Inequality) shows that

$$\mathbb{E}\big[|H \bullet W(t)|^2\big]^{1/2} \le 2\alpha n \mathbb{E}\Big[\sup_{s\in[0,T]} |W(s)|^2\Big]^{1/2} \le 4\alpha n \mathbb{E}\big[|W(T)|^2\big]^{1/2} < \infty,$$

i.e. $H \bullet W$ is square integrable. Finally, the martingale property follows as in the proof of Lemma 3.4 (Characterization of EMMs): Let $s, t \in [0, T]$ with $s < t$. Then

$$\begin{aligned}
&\mathbb{E}_s\big[H \bullet W(t) - H \bullet W(s)\big] \\
&\qquad = \mathbb{E}_s\Big[\sum_{k=1}^{n} H_k\big[W(t_k \wedge t) - W(t_{k-1} \vee s)\big]\Big] \\
&\qquad = \sum_{k=1}^{n} \mathbb{E}_s\Big[H_k \mathbb{E}_{t_{k-1}\vee s}\big[W(t_k \wedge t) - W(t_{k-1} \vee s)\big]\Big] = 0,
\end{aligned}$$

where we have used the martingale property of W and that H_k is $\mathfrak{F}(t_{k-1})$-measurable and hence also $\mathfrak{F}(t_{k-1} \vee s)$-measurable for all $k = 1, \dots, n$. □

The fact that $H \bullet W \in \mathcal{M}^2$ is extremely important to us as this will allow us to extend the integral to more general integrands. The main question at this point is, however, what kind of integrands we may allow.

Definition 5.12 (L^2-Integrand)**.** A predictable process $H = \{H(t)\}_{t\in[0,T]}$ is called **L^2-integrand** if

$$\|H\|_{\mathcal{I}^2} \triangleq \mathbb{E}\Big[\int_0^T |H(t)|^2 \mathrm{d}t\Big]^{1/2} < \infty.$$

The linear space of all L^2-integrands is denoted by $\mathcal{I}^2$. ◇

Treating an L^2-integrand H as a mapping

$$H : [0, T] \times \Omega \to \mathbb{R}, \qquad (t, \omega) \mapsto H(t, \omega),$$

it becomes apparent that $(\mathcal{I}^2, \|\cdot\|_{\mathcal{I}^2})$ can be identified with the Hilbert space of square integrable random variables defined on the probability space $([0, T] \times \Omega, \mathcal{P}, \mathrm{d}t \otimes \mathbb{P})$. The scalar product on this space is given by

$$\langle \cdot | \cdot \rangle_{\mathcal{I}^2} : \mathcal{I}^2 \times \mathcal{I}^2 \to \mathbb{R}, \qquad (H, K) \mapsto \langle H | K \rangle_{\mathcal{I}^2} \triangleq \mathbb{E}\Big[\int_0^T H(t) K(t) \mathrm{d}t \Big].$$

Moreover, as the following result shows, the space $\mathcal{S}$ of simple integrands is dense in $\mathcal{I}^2$.

Proposition 5.13 (Approximation of L^2-Integrands). *The space $\mathcal{S}$ of simple integrands is dense in $\mathcal{I}^2$ with respect to the norm $\|\cdot\|_{\mathcal{I}^2}$.* ◇

Proof. Step 1. Suppose first that $H \in \mathcal{I}^2$ is continuous and bounded. Then we can approximate H by a sequence $\{H_k\}_{k \in \mathbb{N}} \subset \mathcal{S}$ of the form

$$H_k(t) \triangleq \sum_{\ell=1}^{k} H\Big(\frac{\ell - 1}{k} T \Big) \mathbb{1}_{\left(\frac{\ell-1}{k} T, \frac{\ell}{k} T \right]}(t), \qquad t \in [0, T], \; k \in \mathbb{N}.$$

Clearly, with this, we have

$$\lim_{k \to \infty} H_k(t, \omega) = H(t, \omega), \qquad (t, \omega) \in (0, T] \times \Omega,$$

and dominated convergence implies that

$$\lim_{k \to \infty} \|H_k - H\|_{\mathcal{I}^2}^2 = \lim_{k \to \infty} \mathbb{E}\Big[\int_0^T |H_k(t) - H(t)|^2 \mathrm{d}t \Big] = 0.$$

Step 2. Suppose that $H \in \mathcal{I}^2$ is bounded. Define

$$H_k(t) \triangleq \frac{1}{1/k} \int_{(t - 1/k) \vee 0}^{t} H(s) \mathrm{d}s, \qquad t \in [0, T], \; k \in \mathbb{N}.$$

Then each H_k, $k \in \mathbb{N}$, is bounded, continuous, and adapted. In particular, each H_k is predictable and hence $H_k \in \mathcal{I}^2$. But by the fundamental theorem of calculus for Lebesgue integrals, for each $\omega \in \Omega$, there exists a Lebesgue nullset $N_\omega \subset [0,T]$ such that

$$\lim_{k\to\infty} H_k(t,\omega) = H_k(t,\omega) \qquad \text{for all } t \in [0,T] \setminus N_\omega.$$

In other words, $H_k \to H_k$ almost everywhere with respect to $\mathbb{P} \otimes \mathrm{d}t$, and dominated convergence implies that $H_k \to H$ with respect to $\|\cdot\|_{\mathcal{I}^2}$ as in the first step.

Step 3. Let $H \in \mathcal{I}^2$ be arbitrary. Define

$$H_k(t) \triangleq H(t)\mathbb{1}_{\{|H(t)|\leq k\}}, \qquad t \in [0,T],\ k \in \mathbb{N}.$$

Then each $H_k \in \mathcal{I}^2$, $k \in \mathbb{N}$, is bounded and $H_k \to H$ with respect to $\|\cdot\|_{\mathcal{I}^2}$ by dominated convergence once again. □

By the previous result, for each $H \in \mathcal{I}^2$, there exists a sequence $\{H_k\}_{k\in\mathbb{N}} \subset \mathcal{S}$ of simple integrands such that $H_k \to H$ with respect to $\|\cdot\|_{\mathcal{I}^2}$. Moreover, for each H_k, $k \in \mathbb{N}$, we have defined the Itō integral $H_k \bullet W$, which belongs to the Hilbert space $\mathcal{M}^2$. Thus, if we can show that $\{H_k \bullet W\}_{k\in\mathbb{N}}$ is a Cauchy sequence in $\mathcal{M}^2$, we can define the Itō integral $H \bullet W$ as the limit of the Cauchy sequence $\{H_k \bullet W\}_{k\in\mathbb{N}}$.

Theorem 5.14 (Itō Isometry). *The mapping*

$$I : \mathcal{S} \subset \mathcal{I}^2 \to \mathcal{M}^2, \qquad H \mapsto I(H) \triangleq H \bullet W$$

is a Hilbert space isometry, i.e.

$$\langle H|K\rangle_{\mathcal{I}^2} = \langle H \bullet W | K \bullet W\rangle_{\mathcal{M}^2}, \qquad H, K \in \mathcal{S}. \qquad \diamond$$

Proof. Any scalar product $\langle\cdot|\cdot\rangle_{\mathcal{H}}$ on a Hilbert space $(\mathcal{H}, \|\cdot\|_{\mathcal{H}})$ which is compatible with the norm $\|\cdot\|_{\mathcal{H}}$ satisfies the polarization identity

$$\langle f|g\rangle_{\mathcal{H}} = \frac{1}{4}\big(\|f+g\|_{\mathcal{H}}^2 - \|f-g\|_{\mathcal{H}}^2\big), \qquad f, g \in \mathcal{H}.$$

From this, it follows that it suffices to argue that

$$\|H\|_{\mathcal{I}^2} = \|H \bullet W\|_{\mathcal{M}^2}, \qquad H \in \mathcal{S}.$$

Let us therefore fix $H \in \mathcal{S}$ and assume that it admits a representation of the same form as in Equation (5.4). Then

$$\begin{aligned}\|H \bullet W\|_{\mathcal{M}^2}^2 &= \mathbb{E}\Big[\big|H \bullet W(T)\big|^2\Big] \\ &= \mathbb{E}\Big[\Big(\sum_{k=1}^{n} H_k\big[W(t_k) - W(t_{k-1})\big]\Big)^2\Big] \\ &= \sum_{k=1}^{n} \mathbb{E}\Big[H_k^2\big(W(t_k) - W(t_{k-1})\big)^2\Big] \\ &\quad + 2\sum_{k=1}^{n}\sum_{\ell=1}^{k-1} \mathbb{E}\Big[H_k H_\ell\big(W(t_k) - W(t_{k-1})\big)\big(W(t_\ell) - W(t_{\ell-1})\big)\Big].\end{aligned}$$

Now fix $k = 1, \ldots, n$. Then H_k is $\mathfrak{F}(t_{k-1})$-measurable and hence independence and normal distribution of increments of Brownian motion show that

$$\begin{aligned}\mathbb{E}\Big[H_k^2\big(W(t_k) - W(t_{k-1})\big)^2\Big] &= \mathbb{E}\Big[H_k^2 \mathbb{E}_{t_{k-1}}\big[\big(W(t_k) - W(t_{k-1})\big)^2\big]\Big] \\ &= \mathbb{E}\Big[H_k^2 \mathbb{E}\big[\big(W(t_k) - W(t_{k-1})\big)^2\big]\Big] \\ &= \mathbb{E}\Big[(t_k - t_{k-1})H_k^2\Big] = \mathbb{E}\Big[\int_{t_{k-1}}^{t_k} |H(t)|^2 \mathrm{d}t\Big].\end{aligned}$$

Moreover, for each $\ell = 1, \ldots, k-1$, we observe that H_k, H_ℓ, and the increment $W(t_\ell) - W(t_{\ell-1})$ are $\mathfrak{F}(t_{k-1})$-measurable and thus, by the martingale property of Brownian motion,

$$\begin{aligned}&\mathbb{E}\Big[H_k H_\ell\big(W(t_k) - W(t_{k-1})\big)\big(W(t_\ell) - W(t_{\ell-1})\big)\Big] \\ &\qquad = \mathbb{E}\Big[H_k H_\ell\big(W(t_\ell) - W(t_{\ell-1})\big)\mathbb{E}_{t_{k-1}}\big[W(t_k) - W(t_{k-1})\big]\Big] = 0.\end{aligned}$$

Putting the pieces together, this implies that

$$\|H \bullet W\|_{\mathcal{M}^2}^2 = \sum_{k=1}^{n} \mathbb{E}\Big[\int_{t_{k-1}}^{t_k} |H(t)|^2 \mathrm{d}t\Big] = \|H\|_{\mathcal{I}^2}^2. \qquad \square$$

With the **Itō isometry** at hand, it follows that if $\{H_k\}_{k\in\mathbb{N}}$ is a Cauchy sequence in $\mathcal{S}$ converging to some $H \in \mathcal{I}^2$, then $\{H_k \bullet W\}_{k\in\mathbb{N}}$ is Cauchy in $\mathcal{M}^2$, implying that it converges to some continuous square integrable martingale which we may denote by $H \bullet W$ and refer to as the Itō integral of H with respect to W.

Definition 5.15 (Itō Integral for L^2-Integrands)**.** For each $H \in \mathcal{I}^2$, fix an approximating sequence $\{H_k\}_{k\in\mathbb{N}} \subset \mathcal{S}$ of simple integrands such that

$$\lim_{k\to\infty} \|H_k - H\|_{\mathcal{I}^2} = 0.$$

Denote by $H \bullet W \in \mathcal{M}^2$ the continuous square integrable martingale with

$$\lim_{k\to\infty} \|H_k \bullet W - H \bullet W\|_{\mathcal{M}^2} = 0.$$

Then we refer to $H \bullet W$ as the **Itō integral** of H with respect to W. ◇

Observe that the Itō integral $H \bullet W$ is **not defined pathwise** (i.e. constructed separately for each $\omega \in \Omega$), but constructed via **convergence in expectation**. In this sense, the Itō integral is quite different from the Lebesgue integral. We also note that, for now, the integral $H \bullet W$ depends on the choice of approximating sequence $\{H_k\}_{k\in\mathbb{N}} \subset \mathcal{S}$. Let us therefore show now that different approximating sequences yield the same integral.

Proposition 5.16 (Well-Posedness of the Itō Integral)**.** *Suppose that $\{\tilde{H}_k\}_{k\in\mathbb{N}}$ and $\{\hat{H}_k\}_{k\in\mathbb{N}}$ are two sequences of simple integrands converging to the same limit $H \in \mathcal{I}^2$. Denote by $\tilde{H} \bullet W$ and $\hat{H} \bullet W$ the limits of the respective sequences $\{\tilde{H}_k \bullet W\}_{k\in\mathbb{N}}$ and $\{\hat{H}_k \bullet W\}_{k\in\mathbb{N}}$ of Itō integrals in $\mathcal{M}^2$. Then*

$$\tilde{H} \bullet W = \hat{H} \bullet W.$$ ◇

Proof. We define a sequence $\{H_k\}_{k\in\mathbb{N}}$ of simple integrands by setting

$$H_{2k-1} \triangleq \tilde{H}_k \qquad \text{and} \qquad H_{2k} \triangleq \hat{H}_k, \qquad k \in \mathbb{N}.$$

Then $\{H_k\}_{k\in\mathbb{N}}$ also converges to $H \in \mathcal{I}^2$. Denote by $H \bullet W$ the limit of $\{H_k \bullet W\}_{k\in\mathbb{N}}$ in $\mathcal{M}^2$. Then any subsequence of $\{H_k \bullet W\}_{k\in\mathbb{N}}$ also converges to $H \bullet W$ in $\mathcal{M}^2$, implying that

$$\lim_{k\to\infty} \|\tilde{H}_k \bullet W - H \bullet W\|_{\mathcal{M}^2} = \lim_{k\to\infty} \|\hat{H}_k \bullet W - H \bullet W\|_{\mathcal{M}^2} = 0,$$

i.e. $\|\tilde{H} \bullet W - \hat{H} \bullet W\|_{\mathcal{M}^2} = 0$. But then Theorem B.3 (Doob's L^p-Inequality) shows that

$$\mathbb{E}\Big[\sup_{t\in[0,T]} \big|\tilde{H} \bullet W(t) - \hat{H} \bullet W(t)\big|^2\Big] \le 4\|\tilde{H} \bullet W - \hat{H} \bullet W\|_{\mathcal{M}^2} = 0,$$

which implies that $\tilde{H} \bullet W = \hat{H} \bullet W$ as claimed. □

It is possible to extend the Itō integral even further. The crucial ingredient is given by the following **stopping property**.

Lemma 5.17 (Stopping Property of the Itō Integral). *Let $H \in \mathcal{I}^2$ and let τ be a finite, i.e. $[0,T]$-valued, stopping time. Then $H\mathbb{1}_{(0,\tau]} \in \mathcal{I}^2$ and*[3]

$$\big(H\mathbb{1}_{(0,\tau]}\big) \bullet W(t) = H \bullet W(t \wedge \tau), \qquad t \in [0,T]. \qquad \diamond$$

Proof. Let us first observe that $\mathbb{1}_{(0,\tau]}$ is left continuous and adapted, thus predictable. From this, we see that $H\mathbb{1}_{(0,\tau]} \in \mathcal{I}^2$.

Step 1: Assume that $H \in \mathcal{S}$ is a simple integrand and τ takes only finitely many values. Then it is immediate that $\mathbb{1}_{(0,\tau]}$ and hence also$H\mathbb{1}_{(0,\tau]}$ are simple integrands as well. With this, the result is obtained by an elementary calculation using the explicit formula for the Itō integral for simple integrands in Definition 5.10.

Step 2: Suppose that H is simple, but τ is an arbitrary finite stopping time. Then there exists a sequence $\{\tau_k\}_{k\in\mathbb{N}}$ of stopping times converging to τ such that each τ_k, $k \in \mathbb{N}$, takes only finitely many values and $\tau_k \geq \tau$. For example, a possible choice for τ_k is given by

$$\tau_k \triangleq T \wedge \Big(\sum_{\ell=1}^{k2^k} \ell 2^{-k} \mathbb{1}_{\{\tau \in [(\ell-1)2^{-k}, \ell 2^{-k})\}} + \infty \mathbb{1}_{\{\tau \geq k\}} \Big);$$

compare also with Proposition A.12 (Approximation of Stopping Times). Observe that, by dominated convergence,

$$\lim_{k\to\infty} \|H\mathbb{1}_{(0,\tau]} - H\mathbb{1}_{(0,\tau_k]}\|^2_{\mathcal{I}^2} = \lim_{k\to\infty} \mathbb{E}\Big[\int_0^T |H(t)|^2 \mathbb{1}_{(\tau,\tau_k]}(t)\mathrm{d}t\Big] = 0,$$

and therefore, by the Itō isometry,

$$\lim_{k\to\infty} \big\| \big(H\mathbb{1}_{(0,\tau]}\big) \bullet W - \big(H\mathbb{1}_{(0,\tau_k]}\big) \bullet W \big\|_{\mathcal{M}^2} = 0.$$

[3] Observe that the stopped Itō integral $H \bullet W(\cdot \wedge \tau)$ is still adapted by Proposition A.13 (Measurability of Stopped Processes).

As we have seen before, by Doob's L^p inequality and after passing to a subsequence if necessary, this entails

$$\lim_{k\to\infty} \big(H\mathbb{1}_{(0,\tau_k]}\big) \bullet W(t) = \big(H\mathbb{1}_{(0,\tau]}\big) \bullet W(t), \qquad t \in [0,T].$$

On the other hand, by continuity of the paths of $H \bullet W$ and step 1, we have

$$\begin{aligned} H \bullet W(t\wedge\tau) &= \lim_{k\to\infty} H \bullet W(t \wedge \tau_k) \\ &= \lim_{k\to\infty} \big(H\mathbb{1}_{(0,\tau_k]}\big) \bullet W(t) = \big(H\mathbb{1}_{(0,\tau]}\big) \bullet W(t), \qquad t \in [0,T]. \end{aligned}$$

Step 3: In the general case of $H \in \mathcal{I}^2$, we approximate H by a sequence $\{H_k\}_{k\in\mathbb{N}} \subset \mathcal{S}$ of simple integrands. Then, by step 2,

$$H_k \bullet W(t\wedge\tau) = \big(H_k\mathbb{1}_{(0,\tau]}\big) \bullet W(t), \qquad t \in [0,T],\ k \in \mathbb{N},$$

and the result follows as in step 2 since $H_k\mathbb{1}_{(0,\tau]} \to H\mathbb{1}_{(0,\tau]}$ for $k \to \infty$ in $\mathcal{I}^2$ by dominated convergence. □

The stopping property allows us to extend the Itō integral even further.

Definition 5.18 (W-Integrand)**.** We say that a predictable stochastic process $H = \{H(t)\}_{t\in[0,T]}$ is ***W*-integrable** if

$$\int_0^T |H(t)|^2 \mathrm{d}t < \infty. \qquad \diamond$$

Note that the space of W-integrable integrands is **larger** than $\mathcal{I}^2$. However, if H is W-integrable, we may define hitting times

$$\tau_k \triangleq \inf\Big\{t \in [0,T] : \int_0^t |H(s)|^2 \mathrm{d}s \geq k\Big\}, \qquad k \in \mathbb{N}.$$

It follows from Proposition A.9 (Hitting Times of Continuous Time Processes) that each τ_k is even a stopping time. But then $H\mathbb{1}_{(0,\tau_k]} \in \mathcal{I}^2$ since

$$\|H\mathbb{1}_{(0,\tau_k]}\|^2_{\mathcal{I}^2} = \mathbb{E}\Big[\int_0^T |H(t)|^2 \mathbb{1}_{(0,\tau_k]} \mathrm{d}t\Big] = \mathbb{E}\Big[\int_0^{\tau_k\wedge T} |H(t)|^2 \mathrm{d}t\Big] \leq k.$$

In particular, the Itō integral $(H\mathbb{1}_{(0,\tau_k]})\bullet W$ is well-defined and an application of Lemma 5.17 (Stopping Property of the Itō Integral) yields

$$(H\mathbb{1}_{(0,\tau_k]})\bullet W(\cdot\wedge\tau_k) = (H\mathbb{1}_{(0,\tau_{k+1}]})\bullet W(\cdot\wedge\tau_k), \qquad k\in\mathbb{N}.$$

As continuity of $t\mapsto\int_0^t |H(s)|^2\mathrm{d}s$ implies that

$$\mathbb{P}\Big[\lim_{k\to\infty}\tau_k > T\Big] = 1,$$

we find that there exists a continuous adapted process $Y=\{Y(t)\}_{t\in[0,T]}$ with

$$Y(t\wedge\tau_k) = (H\mathbb{1}_{(0,\tau_k]})\bullet W(t\wedge\tau_k), \qquad t\in[0,T],\ k\in\mathbb{N}.$$

This shows that $Y(\cdot\wedge\tau_k)\in\mathcal{M}^2$ and it also guarantees that the process Y is uniquely determined. We therefore arrive at the following definition.

Definition 5.19 (Itō Integral)**.** Let $H=\{H(t)\}_{t\in[0,T]}$ be W-integrable. Then we define the **Itō Integral** of H with respect to W as the unique continuous adapted process $H\bullet W=\{H\bullet W(t)\}_{t\in[0,T]}$ such that

$$H\bullet W(\cdot\wedge\tau) = \big(H\mathbb{1}_{(0,\tau]}\big)\bullet W$$

for every stopping time τ such that $H\mathbb{1}_{(0,\tau]}\in\mathcal{I}^2$. $\diamond$

This definition concludes the construction of the Itō integral. We make the crucial observation that, in general, $H\bullet W$ is **not a martingale**. This is true only if $H\in\mathcal{I}^2$. We also observe that the general Itō integral inherits the **linearity** and the **stopping property** from its simpler versions. More precisely, if $\alpha\in\mathbb{R}$ and H and K are both W-integrable, then $\alpha H+K$ is W-integrable as well and it holds that

$$\big(\alpha H+K\big)\bullet W = \alpha\big(H\bullet W\big)+K\bullet W.$$

The stopping property, on the other hand, just means that if τ is a stopping time, then

$$\big(H\mathbb{1}_{(0,\tau]}\big)\bullet W = H\bullet W(\cdot\wedge\tau).$$

Our next aim is to prove a version of the Fundamental Theorem of Calculus for Itō integrals which is typically referred to as **Itō's formula**.

Definition 5.20 (Itō Process, Characteristic Pair)**.** We say that a continuous adapted process $X = \{X(t)\}_{t\in[0,T]}$ is an **Itō process** if there exist predictable processes $\mu = \{\mu(t)\}_{t\in[0,T]}$ and $\sigma = \{\sigma(t)\}_{t\in[0,T]}$ with

$$\int_0^T |\mu(t)|\mathrm{d}t < \infty \qquad \text{and} \qquad \int_0^T |\sigma(t)|^2\mathrm{d}t < \infty$$

such that X can be written as

$$X(t) = X(0) + \int_0^t \mu(s)\mathrm{d}s + \sigma \bullet W(t), \qquad t \in [0,T]. \tag{5.5}$$

The pair (μ, σ) is referred to as the **characteristic pair** of X. ◇

We note that the characteristic pair of an Itō process X is **uniquely determined** almost everywhere.

Lemma 5.21 (Uniqueness of the Characteristic Pair)**.** *Suppose that X is an Itō process with two characteristic pairs (μ, σ) and $(\tilde{\mu}, \tilde{\sigma})$. Then*

$$\mu(t,\omega) = \tilde{\mu}(t,\omega) \qquad \textit{and} \qquad \sigma(t,\omega) = \tilde{\sigma}(t,\omega)$$

for $\mathrm{d}t \otimes \mathbb{P}$-almost every $(t,\omega) \in [0,T] \times \mathbb{P}$. ◇

Proof. By Equation (5.5), we observe that

$$0 = X(t) - X(t) = \int_0^t \big(\mu(s) - \tilde{\mu}(s)\big)\mathrm{d}s + (\sigma - \tilde{\sigma}) \bullet W(t), \qquad t \in [0,T].$$

Rearranging this equation shows that the process $Y = \{Y(t)\}_{t\in[0,T]}$ given by

$$Y(t) \triangleq -\int_0^t \big(\mu(s) - \tilde{\mu}(s)\big)\mathrm{d}s = (\sigma - \tilde{\sigma}) \bullet W(t), \qquad t \in [0,T],$$

is of finite variation since it can be decomposed into the difference of two increasing processes as follows:

$$Y(t) = \int_0^t \big(\mu(s) - \tilde{\mu}(s)\big)_- \mathrm{d}s - \int_0^t \big(\mu(s) - \tilde{\mu}(s)\big)_+ \mathrm{d}s, \qquad t \in [0,T].$$

On the other hand, by the construction of the Itō integral, we can find an increasing sequence $\{\tau_k\}_{k\in\mathbb{N}}$ of stopping times converging to infinity such that

$$Y(\cdot \wedge \tau_k) = (\sigma - \tilde{\sigma}) \bullet W(\cdot \wedge \tau_k) \in \mathcal{M}^2, \qquad k \in \mathbb{N}.$$

But then $Y(\cdot \wedge \tau_k)$ is a continuous martingale of finite variation and hence

$$Y(T \wedge \tau_k) = Y(t \wedge \tau_k) = Y(0) = 0, \qquad t \in [0,T],\ k \in \mathbb{N},$$

by Theorem B.5 (Continuous Martingales of Finite Variation). The Itō isometry then implies that

$$\begin{aligned} 0 = \|Y(\cdot \wedge \tau_k)\|^2_{\mathcal{M}^2} &= \|\sigma(\cdot \wedge \tau_k) - \tilde{\sigma}(\cdot \wedge \tau_k)\|^2_{\mathcal{I}^2} \\ &= \mathbb{E}\Big[\int_0^T |\sigma(t \wedge \tau_k) - \tilde{\sigma}(t \wedge \tau_k)|^2 \mathrm{d}t\Big], \end{aligned}$$

which is only possible if $\sigma(\cdot \wedge \tau_k) - \tilde{\sigma}(\cdot \wedge \tau_k) = 0$ for all $k \in \mathbb{N}$ and hence $\sigma = \tilde{\sigma}$ almost everywhere with respect to $\mathrm{d}t \otimes \mathbb{P}$. Since Y is therefore equal to zero everywhere, we obtain

$$0 = Y(t) = -\int_0^t \big(\mu(s) - \tilde{\mu}(s)\big)\mathrm{d}s, \qquad t \in [0,T],$$

which, after taking the derivative with respect to t and using the Fundamental Theorem of Calculus for the Lebesgue integral, is only possible if $\mu = \tilde{\mu}$ almost everywhere with respect to $\mathrm{d}t \otimes \mathbb{P}$. □

Having defined an integral with respect to Brownian motion, it is straightforward to define an **integral with respect to Itō processes**.

Definition 5.22 (X-Integrand, Itō Integral for Itō Processes)**.** Let X be an Itō process with characteristic pair (μ, σ). Let moreover $Y = \{Y(t)\}_{t\in[0,T]}$ be predictable. Then we say that Y is ***X*-integrable** if

$$\int_0^T |\mu(t)Y(t)|\mathrm{d}t < \infty \qquad \text{and} \qquad \int_0^T |\sigma(t)Y(t)|^2\mathrm{d}t < \infty.$$

In this case, we define the **Itō integral** of Y with respect to X as the continuous adapted process $Y \bullet X = \{Y \bullet X(t)\}$ given by

$$Y \bullet X(t) \triangleq \int_0^t \mu(s)Y(s)\mathrm{d}s + \sigma Y \bullet W(t), \qquad t \in [0,T].$$ ◇

At this point, it is convenient to introduce the **stochastic differential notation**. Given a process $X = \{X(t)\}_{t\in[0,T]}$, we subsequently write

$$\mathrm{d}X(t) \triangleq X(t) - X(0), \qquad t \in [0,T]. \tag{5.6}$$

Moreover, if X is an Itō process with characteristic pair (μ, σ) and Y is X-integrable, we set

$$Y(t)\mathrm{d}X(t) \triangleq Y(t)\mu(t)\mathrm{d}t + \sigma(t)Y(t)\mathrm{d}W(t), \qquad t \in [0,T]. \tag{5.7}$$

In particular, choosing $Y \equiv 1$, it follows that we may subsequently express the Itō process X as

$$\mathrm{d}X(t) = \mu(t)\mathrm{d}t + \sigma(t)\mathrm{d}W(t), \qquad t \in [0,T].$$

If $\mu \equiv 0$, we have $X = X(0) + \sigma \bullet W$, which implies that

$$\sigma(t)\mathrm{d}W(t) = \mathrm{d}X(t) = \mathrm{d}\sigma \bullet W(t) = \sigma \bullet W(t), \quad t \in [0,T],$$

whereas if $\sigma \equiv 0$, we have $X = \int_0^\cdot \mu(t)\mathrm{d}t$ and our convention reads

$$\mu(t)\mathrm{d}t = \mathrm{d}X(t) = \mathrm{d}\int_0^t \mu(s)\mathrm{d}s = \int_0^t \mu(s)\mathrm{d}s, \quad t \in [0,T].$$

In either case, this shows that the notation in Equation (5.7) is consistent with the notation in Equation (5.6).

Another useful concept worthy of introduction at this point is the so-called **quadratic (co-)variation** of Itō processes.

Definition 5.23 (Quadratic Variation, Quadratic Covariation)**.** Let X and Y be two Itō processes with characteristic pairs (μ_X, σ_X) and (μ_Y, σ_Y), respectively. Then we refer to the process $\langle X, Y\rangle = \{\langle X, Y\rangle(t)\}_{t\in[0,T]}$ given by

$$\langle X, Y\rangle(t) \triangleq \int_0^t \sigma_X(s)\sigma_Y(s)\mathrm{d}s, \qquad t \in [0,T],$$

as the **quadratic covariation** of X and Y. Moreover, we define the **quadratic variation** of X to be the process $\langle X\rangle = \{\langle X\rangle(t)\}_{t\in[0,T]}$ given by

$$\langle X\rangle(t) \triangleq \langle X, X\rangle(t) = \int_0^t \sigma_X(s)^2\mathrm{d}s, \qquad t \in [0,T]. \qquad \diamond$$

The importance of these two concepts will become clear immediately in the following theorem. Before stating the result, however, we extend the stochastic differential notation by setting

$$\mathrm{d}X(t)\mathrm{d}Y(t) \triangleq \mathrm{d}\langle X, Y\rangle(t) = \sigma_X(t)\sigma_Y(t)\mathrm{d}t, \qquad t \in [0,T]$$

as well as

$$\mathrm{d}X(t)\mathrm{d}X(t) \triangleq \mathrm{d}\langle X\rangle(t) = \sigma_X(t)^2\mathrm{d}t, \qquad t \in [0,T],$$

in case of two Itō processes with characteristic pairs (μ_X, σ_X) and (μ_Y, σ_Y), respectively. Note that, with this convention, we have

$$\mathrm{d}W(t)\mathrm{d}W(t) = \mathrm{d}t \quad \text{and} \quad \mathrm{d}t\mathrm{d}t = \mathrm{d}t\mathrm{d}W(t) = \mathrm{d}W(t)\mathrm{d}t = 0 \tag{5.8}$$

for all $t \in [0,T]$, which obtains by choosing $\sigma_X \equiv 1$ and $\sigma_Y \equiv 0$.

With this, we can now state **Itō's formula**. Since the proof is rather technical and lengthy, we only present a sketch. If you are interested in the full proof for a much more general version of this formula, you are invited to attend the course on **Mathematical Finance II**.

Theorem 5.24 (Itō Formula)**.** *For $n \in \mathbb{N}$, let $X_1, \ldots, X_n$ be Itō processes with characteristic pairs $(\mu_1, \sigma_1), \ldots, (\mu_n, \sigma_n)$, respectively. Let $f : \mathbb{R}^n \to \mathbb{R}$ be twice continuously differentiable and define a new process $Y = \{Y(t)\}_{t\in[0,T]}$ by*

$$Y(t) \triangleq f\big(X(t)\big), \qquad t \in [0,T], \qquad \text{where } X(t) \triangleq \big(X_1(t), \ldots, X_n(t)\big).$$

Then Y is an Itō process with characteristic pair (μ_f, σ_f) given by[4]

$$\mu_f(t) \triangleq \sum_{i=1}^{n} \mu_i(t) f_{x_i}\big(X(t)\big) + \frac{1}{2}\sum_{i,j=1}^{n} \sigma_i(t)\sigma_j(t) f_{x_i x_j}\big(X(t)\big), \qquad t \in [0,T],$$

$$\sigma_f(t) \triangleq \sum_{i=1}^{n} \sigma_i(t) f_{x_i}\big(X(t)\big), \qquad t \in [0,T],$$

which is to say that $Y = f(X)$ can be written in differential form as

$$\mathrm{d}f\big(X(t)\big) = \sum_{i=1}^{n} f_{x_i}\big(X(t)\big)\mathrm{d}X_i(t) + \frac{1}{2}\sum_{i,j=1}^{n} f_{x_i,x_j}\big(X(t)\big)\mathrm{d}X_i(t)\mathrm{d}X_j(t)$$

for all $t \in [0,T]$. ◇

[4] Here, f_{x_i} denotes the first-oder partial derivative of f with respect to the i-th variable and $f_{x_i x_j}$ the second-order partial derivative of f with respect to the i-th and j-th variable.

Sketch of the Proof. We only present a very rough sketch of the proof in the special case of $n = 1$ with additional boundedness and integrability assumptions. For this, let X be a bounded Itō process (uniformly in $t \in [0, T]$ and $\omega \in \Omega$) with characteristic pair (μ, σ) such that $\sigma \in \mathcal{I}^2$. Moreover, let $f : \mathbb{R} \to \mathbb{R}$ be twice continuously differentiable with first- and second-order derivatives denoted by f_x and f_{xx}, respectively. Since X is bounded (hence never leaving a sufficiently large compact set) and by continuity we may then also assume that f, f_x, and f_{xx} are bounded and uniformly continuous. We have to prove

$$Y(t) - Y(0) = \int_0^t \Big[\mu(s) f_x\big(X(s)\big) + \frac{1}{2}\sigma(s)^2 f_{xx}\big(X(s)\big)\Big] \mathrm{d}s + \big(\sigma f_x(X)\big) \bullet W(t)$$

for all $t \in [0, T]$. To get there, we first note that a Taylor approximation of f yields

$$f(x + h) - f(x) = f_x(x)h + \frac{1}{2} f_{xx}(x)h^2 + R(x, h)|h|^2, \qquad x, h \in \mathbb{R},$$

where the error function $R : \mathbb{R} \times \mathbb{R}_+ \to \mathbb{R}$ satisfies

$$\lim_{h \downarrow 0} \sup_{x \in \mathbb{R}} |R(x, h)| = 0.$$

Now fix $t \in [0, T]$ and let $0 = t_0 < t_1 < \ldots < t_n = t$. We write $\pi_n \triangleq \max_{k=1,\ldots,n} |t_k - t_{k-1}|$ and assume that $\lim_{n\to\infty} |\pi_n| = 0$. Using a telescope sum and the Taylor approximation of f, it follows that

$$\begin{aligned}
Y(t) - Y(0) &= f\big(X(t)\big) - f\big(X(0)\big) \\
&= \sum_{i=1}^{n} f\big(X(t_k)\big) - f\big(X(t_{k-1})\big) \\
&= \sum_{k=1}^{n} f\big(X(t_{k-1}) + X(t_k) - X(t_{k-1})\big) - f\big(X(t_{k-1})\big) \\
&= J_1^n + J_2^n + J_3^n,
\end{aligned}$$

where

$$J_1^n \triangleq \sum_{k=1}^{n} f_x\big(X(t_{k-1})\big)\big[X(t_k) - X(t_{k-1})\big],$$

$$J_2^n \triangleq \frac{1}{2}\sum_{k=1}^{n} f_{xx}\big(X(t_{k-1})\big)\big[X(t_k) - X(t_{k-1})\big]^2,$$
$$J_3^n \triangleq \sum_{k=1}^{n} R\big(X(t_{k-1}), X(t_k) - X(t_{k-1})\big)\big[X(t_k) - X(t_{k-1})\big]^2.$$

It is then possible to show that

$$\lim_{n\to\infty} J_1^n = f_x\big(X(t)\big)\mathrm{d}X(t) = \int_0^t \mu(s) f_x\big(X(s)\big)\mathrm{d}s + \big(\sigma f_x(X)\big) \bullet W(t),$$
$$\lim_{n\to\infty} J_2^n = \frac{1}{2} f_{xx}\big(X(t)\big)\mathrm{d}X(t)\mathrm{d}X(t) = \frac{1}{2}\int_0^t \sigma(s)^2 f_{xx}\big(X(s)\big)\mathrm{d}s,$$
$$\lim_{n\to\infty} J_3^n = 0,$$

where the convergence is in probability. Thus

$$\begin{aligned} Y(t) - Y(0) &= J_1^n + J_2^n + J_3^n \\ &\to \int_0^t \Big[\mu(s) f_x\big(X(s)\big) + \frac{1}{2}\sigma(s)^2 f_{xx}\big(X(s)\big)\Big]\mathrm{d}s \\ &\qquad + \big(\sigma f_x(X)\big) \bullet W(t) \end{aligned}$$

in probability as $n \to \infty$, which concludes the proof by almost sure uniqueness of limits arising from convergence in probability. □

The importance of the Itō formula cannot be stressed enough and it is hence worthwhile to have a closer look. Fist of all, let us note that Itō's formula highlights the convenience of our stochastic differential notation: If X is an Itō process with characteristic pair (μ, σ) and $f : \mathbb{R} \to \mathbb{R}$ is twice continuously differentiable, Itō's formula in differential notation states that

$$\mathrm{d}f\big(X(t)\big) = f_x\big(X(t)\big)\mathrm{d}X(t) + \frac{1}{2} f_{xx}\big(X(t)\big)\mathrm{d}X(t)\mathrm{d}X(t).$$

Using $\mathrm{d}X(t) = \mu(t)\mathrm{d}t + \sigma(t)\mathrm{d}W(t)$, **formal computations** lead to

$$\begin{aligned} \mathrm{d}f\big(X(t)\big) &= f_x\big(X(t)\big)\mathrm{d}X(t) + \frac{1}{2} f_{xx}\big(X(t)\big)\mathrm{d}X(t)\mathrm{d}X(t) \\ &= f_x\big(X(t)\big)\big[\mu(t)\mathrm{d}t + \sigma(t)\mathrm{d}W(t)\big] \\ &\qquad + \frac{1}{2} f_{xx}\big(X(t)\big)\big[\mu(t)\mathrm{d}t + \sigma(t)\mathrm{d}W(t)\big]^2 \end{aligned}$$

$$
\begin{aligned}
&= \mu(t) f_x\big(X(t)\big)\mathrm{d}t + \sigma(t) f_x\big(X(t)\big)\mathrm{d}W(t) \\
&\quad + \frac{1}{2} f_{xx}\big(X(t)\big)\Big[\mu(t)^2 \mathrm{d}t\mathrm{d}t + \sigma(t)^2 \mathrm{d}W(t)\mathrm{d}W(t) \\
&\qquad + 2\mu(t)\sigma(t)\mathrm{d}t\mathrm{d}W(t)\Big].
\end{aligned}
$$

Using Equation (5.8) implies that

$$
\mu(t)^2 \mathrm{d}t\mathrm{d}t + \sigma(t)^2 \mathrm{d}W(t)\mathrm{d}W(t) + 2\mu(t)\sigma(t)\mathrm{d}t\mathrm{d}W(t) = 0 + \sigma(t)^2\mathrm{d}t + 0,
$$

and hence, after grouping $\mathrm{d}t$ and $\mathrm{d}W(t)$ terms,

$$
\begin{aligned}
\mathrm{d}f\big(X(t)\big) &= \Big[\mu(t) f_x\big(X(t)\big) + \frac{1}{2}\sigma(t)^2 f_{xx}\big(X(t)\big)\Big]\mathrm{d}t + \sigma(t) f_x\big(X(t)\big)\mathrm{d}W(t) \\
&= \mu_f(t)\mathrm{d}t + \sigma_f(t)\mathrm{d}W(t)
\end{aligned}
$$

as expected. Hence, instead of having to memorize the complicated expressions for μ_f and σ_f, one can instead memorize the arguably easier differential form of Itō's formula and recover μ_f and σ_f using formal computations with stochastic differentials using the rules in Equation (5.8).

Next, let us consider the very special Itō process X with $X(t) = t$ for all $t \in [0, T]$, i.e. the Itō process with characteristic pair $(1, 0)$. Applying the Itō formula to a twice continuously differentiable function $f : \mathbb{R} \to \mathbb{R}$ yields

$$
f(t) = f\big(X(t)\big) = f(0) + \int_0^t f_x(s)\mathrm{d}s, \qquad t \in [0, T].
$$

From this, we see that the Itō formula extends the **Fundamental Theorem of Calculus** to Itō integrals. In differential form, this rewrites as

$$
\mathrm{d}f\big(X(t)\big) = \mathrm{d}f(t) = f_x(t)\mathrm{d}t = f_x\big(X(t)\big)\mathrm{d}X(t), \qquad t \in [0, T],
$$

whereas for general Itō processes with nonvanishing σ the equation reads

$$
\mathrm{d}f\big(X(t)\big) = f_x\big(X(t)\big)\mathrm{d}X(t) + \frac{1}{2} f_{xx}\big(X(t)\big)\mathrm{d}\langle X\rangle(t), \qquad t \in [0, T].
$$

We may therefore think of quadratic (co-)variations as **correction terms** which make sure that the Fundamental Theorem of Calculus can be extended to Itō integrals/processes.

We should also take note here that in the special case of $X(t) = t$ it actually suffices that f admits continuous first-order derivatives[5] instead of being twice continuously differentiable. More generally, if $X_1, \dots, X_n$ are Itō processes and there exists $i \in \{1, \dots, n\}$ such that X_i has characteristic pair $(\mu, 0)$, i.e. the Itō integral in its representation vanishes, then it suffices that $f : \mathbb{R}^n \to \mathbb{R}$ is **once continuously differentiable** in its i-th component. We shall often make use of this observation in the case of $n = 2$ with an Itō process X with characteristic pair (μ, σ) and a function $f : [0,T] \times \mathbb{R} \to \mathbb{R}$. If f is once continuously differentiable in its first variable and twice continuously differentiable in its second variable, this generalization of Itō's formula yields

$$\mathrm{d}f\big(t, X(t)\big) = \Big[f_t\big(t, X(t)\big) + \mu(t) f_x\big(t, X(t)\big) + \frac{1}{2}\sigma(t)^2 f_{xx}\big(t, X(t)\big)\Big]\mathrm{d}t \\ + \sigma(t) f_x\big(t, X(t)\big)\mathrm{d}W(t)$$

for all $t \in [0,T]$, where f_t denotes the time derivative and f_x and f_{xx} the spatial derivatives of f.

Another very useful special case is the situation in which

$$f : \mathbb{R}^2 \to \mathbb{R}, \qquad (x_1, x_2) \mapsto f(x_1, x_2) = x_1 x_2.$$

In this case, Itō's formula states that the **product of two Itō processes is an Itō process**. This is used so often that it deserves to be perpetuated in a corollary.

Corollary 5.25 (Itō's Product Rule)**.** *Let X and Y be two Itō processes with characteristic pairs (μ_X, σ_X) and (μ_Y, σ_Y), respectively. Then the product $XY = \{X(t)Y(t)\}_{t\in[0,T]}$ is again an Itō process of the form*

$$\begin{aligned} \mathrm{d}\big(XY\big)(t) &= Y(t)\mathrm{d}X(t) + X(t)\mathrm{d}Y(t) + \mathrm{d}X(t)\mathrm{d}Y(t) \\ &= \Big[\mu_X(t)Y(t) + \mu_Y(t)X(t) + \sigma_X(t)\sigma_Y(t)\Big]\mathrm{d}t \\ &\qquad + \Big[\sigma_Y(t)X(t) + \sigma_X(t)Y(t)\Big]\mathrm{d}W(t) \end{aligned}$$

for all $t \in [0,T]$. ◇

[5] Or, more generally, in this case it is sufficient that f is absolutely continuous.

With this, we now have finally all tools available to start doing mathematical finance in continuous time! We wrap up this section by gathering the **formal rules of calculus** for stochastic differentials:

$$
\begin{aligned}
\mathrm{d}X &= \mu \mathrm{d}t + \sigma \mathrm{d}W, \\
[H + \alpha K]\mathrm{d}X &= H\mathrm{d}X + \alpha K \mathrm{d}X, \\
H\mathrm{d}[X + \alpha Y] &= H\mathrm{d}X + \alpha H \mathrm{d}Y, \\
[H\mathrm{d}X][K\mathrm{d}Y(t)] &= HK\mathrm{d}X\mathrm{d}Y, \\
\mathrm{d}f(X) &= f_x(X)\mathrm{d}X + \tfrac{1}{2} f_{xx}(X)\mathrm{d}X\mathrm{d}X, \\
\mathrm{d}t\mathrm{d}t &= \mathrm{d}t\mathrm{d}W = \mathrm{d}W\mathrm{d}t = 0, \\
\mathrm{d}W\mathrm{d}W &= \mathrm{d}t.
\end{aligned}
$$

It is advisable to learn these rules by heart as they will come in handy later!

5.3 Option Pricing in the Black-Scholes Model

Let us now return to mathematical finance. Inspired by Theorem 5.2 (Weak Convergence of the CRR(N) Model), we shall subsequently introduce a **continuous time** financial market model. For this, we fix a probability space $(\Omega, \mathfrak{A}, \mathbb{P})$ which is sufficiently rich to support a Brownian motion $W = \{W(t)\}_{t\in[0,T]}$ and we denote by $\mathfrak{F} = \mathfrak{F}^W = \{\mathfrak{F}^W(t)\}_{t\in[0,T]}$ the augmented filtration generated by W.

Definition 5.26 (Black-Scholes Model)**.** For an initial price $s > 0$, riskless interest rate $r \in \mathbb{R}$, volatility $\sigma > 0$, and excess return $\lambda \in \mathbb{R}$, we define the **Black-Scholes model** to be the two-dimensional process $S = (S^0, S^1) = \{(S^0(t), S^1(t))\}_{t\in[0,T]}$ given by

$$S^0(t) \triangleq e^{rt} \qquad \text{and} \qquad S^1(t) \triangleq s e^{(r+\lambda-\frac{1}{2}\sigma^2)t+\sigma W(t)}, \qquad t \in [0,T]. \qquad \diamond$$

We observe that S^0 and S^1 are **Itō processes**. Indeed, consider the Itō processes

$$\mathrm{d}R^0(t) = r\mathrm{d}t, \qquad t \in [0,T], \qquad R^0(0) = 0,$$

$$\mathrm{d}R^1(t) = \big[r + \lambda - \frac{1}{2}\sigma^2\big]\mathrm{d}t + \sigma \mathrm{d}W(t), \qquad t \in [0,T], \qquad R^1(0) = 0,$$

i.e.

$$R^0(t) = rt \qquad \text{and} \qquad R^1(t) = \big[r + \lambda - \frac{1}{2}\sigma^2\big]t + \sigma W(t), \qquad t \in [0,T].$$

Now apply Itō's formula with $f : \mathbb{R} \to \mathbb{R}$, $x \mapsto f(x) \triangleq e^x$ and use that $f = f_x = f_{xx}$ to obtain

$$\begin{aligned}\mathrm{d}S^0(t) = \mathrm{d}f\big(R^0(t)\big) &= f\big(R^0(t)\big)\mathrm{d}R^0(t) + \frac{1}{2}f\big(R^0(t)\big)\mathrm{d}R^0(t)\mathrm{d}R^0(t)\\ &= rf\big(R^0(t)\big)\mathrm{d}t + \frac{1}{2}r^2 f\big(R^0(t)\big)\mathrm{d}t\mathrm{d}t = rS^0(t)\mathrm{d}t\end{aligned}$$

as well as

$$\begin{aligned}\mathrm{d}S^1(t) = s\mathrm{d}f\big(R^1(t)\big) &= sf\big(R^1(t)\big)\mathrm{d}R^1(t) + \frac{1}{2}sf\big(R^1(t)\big)\mathrm{d}R^1(t)\mathrm{d}R^1(t)\\ &= \big[r + \lambda - \frac{1}{2}\sigma^2\big]S^1(t)\mathrm{d}t + \sigma S^1(t)\mathrm{d}W(t) + \frac{1}{2}\sigma^2 S^1(t)\mathrm{d}t\\ &= (r+\lambda)S^1(t)\mathrm{d}t + \sigma S^1(t)\mathrm{d}W(t)\end{aligned}$$

for all $t \in [0,T]$. The latter representation of S^1 is the first instance of a **stochastic differential equation**. Note, however, that this name is misleading in that it is actually a **stochastic integral equation**.

Observe that S^0 is a **numéraire** in the sense of Definition 2.4, so that we can think of S^0 as the evolution of one unit of money in time. S^1, on the other hand, can be thought of as the price of a risky asset such as a share of a stock. The **discounted risky asset price** is clearly given by

$$\overline{S}^1(t) = \frac{S^1(t)}{S^0(t)} = se^{(\lambda - \frac{1}{2}\sigma^2)t + \sigma W(t)}, \qquad t \in [0,T],$$

or, in differential notation,

$$\mathrm{d}\overline{S}^1(t) = \lambda \overline{S}^1(t)\mathrm{d}t + \sigma \overline{S}^1(t)\mathrm{d}W(t), \qquad t \in [0,T].$$

From this representation we see that $\overline{S}^1$ can only be a **martingale** if $\lambda = 0$. If $\lambda > 0$, however, we expect $\overline{S}^1$ to be a **submartingale** whereas we expect $\overline{S}^1$ to be a **supermartingale** if the excess return λ is negative.

Let us now turn to **trading strategies**. Just as in discrete time, trading strategies are identified with the set of processes which can be integrated with respect to S.

Definition 5.27 ((Self-Financing) Trading Strategy; Wealth Process)**.** An $\mathbb{R}^2$-valued predictable process $\varphi = (\varphi_0, \varphi_1) = \{(\varphi_0(t), \varphi_1(t))\}_{t\in[0,T]}$ is called **trading strategy** if φ_0 is S^0-integrable and φ_1 is S^1-integrable. The **wealth process** $X^\varphi = \{X^\varphi(t)\}_{t\in[0,T]}$ associated with φ is defined as

$$X^\varphi(t) \triangleq \langle \varphi(t)|S(t)\rangle = \varphi_0(t)S^0(t) + \varphi_1(t)S^1(t), \qquad t \in [0,T],$$

and we say that φ is **self-financing** with respect to S provided that

$$\mathrm{d}X^\varphi(t) = \varphi_0(t)\mathrm{d}S^0(t) + \varphi_1(t)\mathrm{d}S^1(t), \qquad t \in [0,T]. \qquad \diamond$$

Observe the subtle difference in the definition of self-financing trading strategies. In discrete-time, the self-financing condition in Equation (2.1) was argued to be equivalent to the condition that the wealth process is equal to initial wealth plus **profits and losses from trading**; see Lemma 2.12 (Characterization of Self-Financing Trading Strategies). In continuous time, on the other hand, we take the latter characterization as the **defining property**. Note that, for this to make sense, we need that φ is integrable with respect to S. We also note that the differential form of the wealth process can easily be simplified to obtain

$$\begin{aligned}\mathrm{d}X^\varphi(t) &= \varphi_0(t)\mathrm{d}S^0(t) + \varphi_1(t)\mathrm{d}S^1(t)\\ &= \varphi_0(t)rS^0(t)\mathrm{d}t + \varphi_1(t)\big[(r+\lambda)S^1(t)\mathrm{d}t + \sigma S^1(t)\mathrm{d}W(t)\big]\\ &= \big[rX^\varphi(t) + \lambda\varphi_1(t)S^1(t)\big]\mathrm{d}t + \sigma\varphi_1(t)S^1(t)\mathrm{d}W(t), \qquad t \in [0,T].\end{aligned}$$

Just as in discrete time, a trading strategy is self-financing with respect to S if and only if it is self-financing with respect to the **discounted market** $\overline{S}$; compare with Lemma 2.12 (Characterization of Self-Financing Strategies).

Lemma 5.28 (Numéraire Invariance)**.** *A trading strategy φ is self-financing with respect to S if and only if it is self-financing with respect to the discounted market $\overline{S}$.* $\diamond$

Proof. Let $i = 0, 1$ and $t \in [0, T]$. We first observe that, since S^0 is bounded, φ_i is S^i-integrable if and only if φ_i is $\overline{S}^i$-integrable. Now let $\beta = \{\beta(t)\}_{t\in[0,T]}$ be an arbitrary Itō process. Then Itō's product rule shows that

$$\mathrm{d}\big(\beta S^i\big)(t) = \beta(t)\mathrm{d}S^i(t) + S^i(t)\mathrm{d}\beta(t) + \mathrm{d}\beta(t)\mathrm{d}S^i(t) \tag{5.9}$$

as well as

$$\mathrm{d}\big(\beta X^\varphi\big)(t) = \beta(t)\mathrm{d}X^\varphi(t) + X^\varphi(t)\mathrm{d}\beta(t) + \mathrm{d}\beta(t)\mathrm{d}X^\varphi(t).$$

If φ is self-financing, we can use $\mathrm{d}X^\varphi(t) = \varphi_0(t)\mathrm{d}S^0(t) + \varphi_1(t)\mathrm{d}S^1(t)$ and the general fact that $X^\varphi(t) = \varphi_0(t)S^0(t) + \varphi_1(t)S^1(t)$ to arrive at

$$\begin{aligned}
\mathrm{d}\big(\beta X^\varphi\big)(t) &= \beta(t)\mathrm{d}X^\varphi(t) + X^\varphi(t)\mathrm{d}\beta(t) + \mathrm{d}\beta(t)\mathrm{d}X^\varphi(t) \\
&= \beta(t)\big[\varphi_0(t)\mathrm{d}S^0(t) + \varphi_1(t)\mathrm{d}S^1(t)\big] \\
&\qquad\qquad + \big[\varphi_0(t)S^0(t) + \varphi_1(t)S^1(t)\big]\mathrm{d}\beta(t) \\
&\qquad\qquad\qquad\qquad + \mathrm{d}\beta(t)\big[\varphi_0(t)\mathrm{d}S^0(t) + \varphi_1(t)\mathrm{d}S^1(t)\big] \\
&= \varphi_0(t)\big[\beta(t)\mathrm{d}S^0(t) + S^0(t)\mathrm{d}\beta(t) + \mathrm{d}\beta(t)\mathrm{d}S^0(t)\big] \\
&\qquad\qquad + \varphi_1(t)\big[\beta(t)\mathrm{d}S^1(t) + S^1(t)\mathrm{d}\beta(t) + \mathrm{d}\beta(t)\mathrm{d}S^1(t)\big].
\end{aligned}$$

But if we then use Equation (5.9), it follows that

$$\mathrm{d}\big(\beta X^\varphi\big)(t) = \varphi_0(t)\mathrm{d}\big(\beta S^0\big)(t) + \varphi_1(t)\mathrm{d}\big(\beta S^1\big)(t),$$

and upon choosing $\beta \triangleq 1/S^0$, this rewrites as

$$\mathrm{d}\overline{X}^\varphi(t) = \varphi_0(t)\mathrm{d}\overline{S}^0(t) + \varphi_1(t)\mathrm{d}\overline{S}^1(t),$$

i.e. φ is self-financing with respect to $\overline{S}$. Conversely, if φ is self-financing with respect to $\overline{S}$, repeating the same argument with $\overline{S}$ is place of S and choosing $\beta \triangleq S^0$, one find that

$$\mathrm{d}X^\varphi(t) = \varphi_0(t)\mathrm{d}S^0(t) + \varphi_1(t)\mathrm{d}S^1(t),$$

i.e. φ is self-financing with respect to S. □

Again in analogy with the discrete time case, we can discard φ_0 provided we are **given the initial wealth**; compare with Lemma 2.13 (Construction of Self-Financing Strategies).

Lemma 5.29 (Trading Strategies with Fixed Initial Wealth). *Let $x \in \mathbb{R}$ and φ_1 be S^1-integrable. Then $\varphi_0 = \{\varphi_0(t)\}_{t\in[0,T]}$ given by*

$$\varphi_0(t) \triangleq x + \varphi_1 \bullet \overline{S}^1(t) - \varphi_1(t)\overline{S}^1(t), \qquad t \in [0,T]$$

is the unique S^0-integrable process such that $\varphi \triangleq (\varphi_0, \varphi_1)$ is a self-financing trading strategy with initial wealth $X^\varphi(0) = x$. ◇

Proof. We first observe that φ_0 is clearly predictable. To see that φ_0 is S^0-integrable, it suffices to argue that

$$\begin{aligned}\int_0^T \left|\varphi_0(t) r S^0(t)\right| \mathrm{d}t &= \int_0^T \left|\big(x + \varphi_1 \bullet \overline{S}^1(t) - \varphi_1(t)\overline{S}^1(t)\big) r S^0(t)\right| \mathrm{d}t \\ &\leq \int_0^T \left|r x S^0(t)\right| \mathrm{d}t \\ &\qquad + \int_0^T \left|\varphi_1 \bullet \overline{S}^1(t)\right|\left|r S^0(t)\right| \mathrm{d}t \\ &\qquad\qquad + \int_0^T \left|r\varphi_1(t)\overline{S}^1(t) S^0(t)\right| \mathrm{d}t < \infty.\end{aligned}$$

Here, the finiteness of the integrals of $|rxS^0|$ and $|\varphi_1 \bullet \overline{S}^1||rS^0|$ follows simply from the fact that these integrands are continuous, hence bounded on $[0,T]$ for each $\omega \in \Omega$. The finiteness of the integral of $|r\varphi_1 \overline{S}^1 S^0| = |r\varphi_1 S^1|$, on the other hand, follows directly from S^1-integrability of φ_1. Thus $\varphi = (\varphi_0, \varphi_1)$ is indeed a trading strategy. But then by Lemma 5.28 (Numéraire Invariance), φ is self-financing with $X^\varphi(0) = x$ if and only if

$$\overline{X}^\varphi(t) = \overline{X}^\varphi(0) + \varphi_0 \bullet \overline{S}^0(t) + \varphi_1 \bullet \overline{S}^1(t) = x + \varphi_1 \bullet \overline{S}^1(t), \qquad t \in [0,T].$$

But $\overline{X}^\varphi(t) = \varphi_0(t)\overline{S}^0(t) + \varphi_1(t)\overline{S}^1(t) = \varphi_0(t) + \varphi_1(t)\overline{S}^1(t)$ for all $t \in [0,T]$, and hence solving for φ_0 yields that φ is self-financing with $X^\varphi(0) = x$ if and only if

$$\varphi_0(t) = x + \varphi_1 \bullet \overline{S}^1(t) - \varphi_1(t)\overline{S}^1(t), \qquad t \in [0,T].$$ □

Still mirroring our approach in discrete time, we can now **identify** self-financing trading strategies in the sense of Definition 5.27 with pairs (x, φ_1), where $x \in \mathbb{R}$ denotes the initial wealth and φ_1 is S^1-integrable and models the number of shares of S^1 held by the investor. As before, we switch between these two notions whenever we see fit.

Definition 5.30. A pair (x, φ_1) with $x \in \mathbb{R}$ and φ_1 being S^1-integrable is called **trading strategy with initial wealth** x. ◇

We can now move on to the topic of **arbitrage**. In contrast to discrete time, our definition in the Black-Scholes model is sightly more restrictive.

Definition 5.31 (Admissible Trading Strategy; Arbitrage)**.** A self-financing trading strategy φ is called **admissible** if the corresponding wealth process is lower bounded, i.e. if there exists some $\alpha \geq 0$ such that

$$X^{\varphi}(t) \geq -\alpha, \qquad t \in [0, T].$$

Moreover, an admissible trading strategy φ is referred to as an **arbitrage opportunity** provided that

$$X^{\varphi}(0) = 0, \qquad X^{\varphi}(T) \geq 0, \qquad \text{and} \qquad \mathbb{P}[X^{\varphi}(T) > 0] > 0. \qquad ◇$$

Admissible strategies are exactly those self-financing strategies which can be financed with a **finite credit line**. If we were not to restrict to admissible strategies in the definition of arbitrage, then the Black-Scholes model would not be free of arbitrage. From an **economic point of view**, it is a very reasonable assumption to restrict to strategies with finite credit lines as traders would never find a counterparty willing to provide unrestricted credit.

In discrete time, the existence of arbitrage was directly linked to the existence of an **equivalent martingale measure**, i.e. the existence of a probability measure $\mathbb{Q} \sim \mathbb{P}$ such that $\overline{S}$ (or, equivalently, $\overline{S}^1$) is a martingale under $\mathbb{Q}$. It should not come as a surprise that the same is true in the Black-Scholes model. Before we construct an EMM in the Black-Scholes model, however, let us first state a useful result on exponential martingales. For this, we recall that the **moment generating function** of a normally distributed random variable X with mean μ and variance ϑ^2 is given by

$$M_X : \mathbb{R} \to \mathbb{R}_+, \qquad u \mapsto M_X(u) \triangleq \mathbb{E}\big[e^{uX}\big] = e^{u\mu + \frac{1}{2}u^2\vartheta^2}.$$

Moreover, we recall that moment generating functions **characterize** the distribution of random variables **uniquely**, provided of course they exist. In particular, this is the case for the normal distribution.

Lemma 5.32 (Exponential Martingales). *For any $\vartheta \in \mathbb{R}$, the continuous adapted process $Z = \{Z(t)\}_{t\in[0,T]}$ given by*

$$Z(t) \triangleq e^{-\frac{1}{2}\vartheta^2 t + \vartheta W(t)}, \qquad t \in [0,T],$$

is a martingale. ◇

Proof. Let $s, t \in [0,T]$ with $s \le t$. Then

$$\begin{aligned}\mathbb{E}_s\big[Z(t)\big] = Z(s)\mathbb{E}_s\left[\frac{Z(t)}{Z(s)}\right] &= Z(s)\mathbb{E}_s\left[e^{-\frac{1}{2}\vartheta^2(t-s)+\vartheta(W(t)-W(s))}\right] \\ &= Z(s)e^{-\frac{1}{2}\vartheta^2(t-s)}\mathbb{E}\left[e^{\vartheta(W(t)-W(s))}\right],\end{aligned}$$

where we have used independence of $W(t) - W(s)$ and $\mathfrak{F}(s)$ to switch to the unconditional expectation in the last line. But then we may recognize the expectation as the moment generating function of the normally distributed random variable $W(t) - W(s)$ (with mean zero and variance $t - s$). Thus

$$\mathbb{E}\left[e^{\vartheta(W(t)-W(s))}\right] = M_{W(t)-W(s)}(\vartheta) = e^{\frac{1}{2}\vartheta^2(t-s)}.$$

It therefore follows that

$$\mathbb{E}_s\big[Z(t)\big] = Z(s)e^{-\frac{1}{2}\vartheta^2(t-s)}\mathbb{E}\left[e^{\vartheta(W(t)-W(s))}\right] = Z(s),$$

i.e. Z is indeed a martingale. □

With this result on exponential martingales at hand, it is now straightforward to construct an **equivalent martingale measure** in the Black-Scholes model.

Proposition 5.33 (EMM in the Black-Scholes Model). *Consider the adapted continuous process $Z = \{Z(t)\}_{t\in[0,T]}$ given by*

$$Z(t) \triangleq e^{-\frac{1}{2}\theta^2 t - \theta W(t)}, \qquad t \in [0,T],$$

where $\theta \triangleq \lambda/\sigma$. Then there exists an equivalent martingale measure $\mathbb{Q}$ such that the density of $\mathbb{Q}$ with respect to $\mathbb{P}$ on $(\Omega, \mathfrak{F}(t))$ is given by

$$\left.\frac{\mathrm{d}\mathbb{Q}}{\mathrm{d}\mathbb{P}}\right|_{\mathfrak{F}(t)} : \Omega \to \mathbb{R}_+, \qquad \omega \mapsto \left.\frac{\mathrm{d}\mathbb{Q}}{\mathrm{d}\mathbb{P}}\right|_{\mathfrak{F}(t)}(\omega) = Z(t,\omega), \qquad t \in [0,T]. \quad ◇$$

Proof. By Lemma 5.32 (Exponential Martingales) with $\vartheta \triangleq -\theta$, it follows that Z is a strictly positive martingale with $\mathbb{E}[Z(T)] = \mathbb{E}[Z(0)] = 1$. Hence, defining $\mathbb{Q}$ in terms of the density process Z indeed defines a probability measure equivalent to $\mathbb{P}$. It remains to check if $\overline{S}$ is a martingale under $\mathbb{Q}$. Since $\overline{S}^0 \equiv 1$, we only have to check the martingale property of $\overline{S}^1$. Let us therefore fix $s, t \in [0, T]$ with $s \leq t$. By Equation (3.13) we find that

$$\mathbb{E}_s^{\mathbb{Q}}\left[\overline{S}^1(t)\right] = \mathbb{E}_s\left[\overline{S}^1(t)\frac{Z(t)}{\mathbb{E}_s[Z(t)]}\right] = \mathbb{E}_s\left[\overline{S}^1(t)\frac{Z(t)}{Z(s)}\right] = \frac{1}{Z(s)}\mathbb{E}_s\left[\overline{S}^1(t)Z(t)\right].$$

Now, by definition of $\overline{S}^1$ and Z and using that $\lambda = \sigma\theta$, we see that

$$\overline{S}^1(t)Z(t) = se^{(\lambda-\frac{1}{2}\sigma^2)t+\sigma W(t)}e^{-\frac{1}{2}\theta^2 t-\theta W(t)} = se^{-\frac{1}{2}(\sigma-\theta)^2 t+(\sigma-\theta)W(t)},$$

which in combination with Lemma 5.32 (Exponential Martingales) shows that $\{\overline{S}^1(t)Z(t)\}_{t\in[0,T]}$ is a $\mathbb{P}$-martingale. Thus

$$\mathbb{E}_s^{\mathbb{Q}}\left[\overline{S}^1(t)\right] = \frac{1}{Z(s)}\mathbb{E}_s\left[\overline{S}^1(t)Z(t)\right] = \overline{S}^1(s),$$

i.e. $\overline{S}^1$ is a $\mathbb{Q}$-martingale and hence $\mathbb{Q}$ is an EMM. □

We subsequently keep the EMM $\mathbb{Q}$ and the corresponding density process $Z = \{Z(t)\}_{t\in[0,T]}$ constructed in Proposition 5.33 (EMM in the Black-Scholes Model) fixed. It turns out that $\mathbb{Q}$ is the **unique EMM** in this model, although we shall not prove this result. We also note that the parameter $\theta = \lambda/\sigma$ in the construction of Z is often referred to as the **market price of risk**.

Under the equivalent martingale measure $\mathbb{Q}$, it is no longer true that W is a Brownian motion. We do have, however, have the following result.

Theorem 5.34 (Girsanov)**.** *The process $B = \{B(t)\}_{t\in[0,T]}$ given by*

$$B(t) \triangleq \theta t + W(t), \qquad t \in [0, T],$$

is a Brownian motion under the equivalent martingale measure $\mathbb{Q}$. ◇

Proof. It is clear that B is continuous, adapted, $B(0) = 0$ almost surely (with respect to $\mathbb{Q}$), and that $B(t) - B(s)$ is independent of $\mathfrak{F}(s)$ whenever $s, t \in [0, T]$ with $s \leq t$. We therefore only have to argue that $B(t) - B(s)$ is normally distributed with mean zero and variance $t - s$ under $\mathbb{Q}$. For this, we compute the moment generating function of $B(t) - B(s)$ under $\mathbb{Q}$: For any $u \in \mathbb{R}$, using independence of $B(t) - B(s)$ and $W(t) - W(s)$ of $\mathfrak{F}(s)$ and $\mathbb{E}[Z(s)] = 1$, we have

$$\begin{aligned}
\mathbb{E}^{\mathbb{Q}}\left[e^{u(B(t)-B(s))}\right] &= \mathbb{E}\left[e^{u(B(t)-B(s))}Z(t)\right] \\
&= \mathbb{E}\left[e^{u(B(t)-B(s))}\frac{Z(t)}{Z(s)}\right]\mathbb{E}[Z(s)] \\
&= \mathbb{E}\left[e^{u\theta(t-s)+u(W(t)-W(s))}e^{-\frac{1}{2}\theta^2(t-s)-\theta(W(t)-W(s))}\right] \\
&= e^{(u\theta-\frac{1}{2}\theta^2)(t-s)}\mathbb{E}\left[e^{(u-\theta)(W(t)-W(s))}\right] \\
&= e^{(u\theta-\frac{1}{2}\theta^2)(t-s)}e^{\frac{1}{2}(u-\theta)^2(t-s)} = e^{\frac{1}{2}u^2(t-s)},
\end{aligned}$$

which we recognize as the moment generating function of a normally distributed random variable with mean zero and variance $t - s$, implying that B is a Brownian motion under $\mathbb{Q}$. □

One of the key result concerning equivalent martingale measures in discrete time was Lemma 3.4 (Characterization of EMMs), which stated, in particular, that any **discounted wealth process** $\overline{X}^{\varphi}$ is a martingale under $\mathbb{Q}$ as well. In continuous time, the same need no longer be true as integrability might fail.

Lemma 5.35 (Discounted Wealth Processes under EMMs)**.** *Let φ be an admissible trading strategy and denote by $B = \{B(t)\}_{t\in[0,T]}$ the $\mathbb{Q}$-Brownian motion constructed in Theorem 5.34 (Girsanov). Then the discounted wealth process $\overline{X}^{\varphi}$ can be written as*

$$\mathrm{d}\overline{X}^{\varphi}(t) = \sigma\varphi_1(t)\overline{S}^1(t)\mathrm{d}B(t), \qquad t \in [0, T],$$

and it is a supermartingale under $\mathbb{Q}$. ◇

Proof. Let φ be an admissible trading strategy. By Itō's product rule, the corresponding discounted wealth process $\overline{X}^{\varphi} = X^{\varphi}(S^0)^{-1}$ has dynamics

$$\begin{aligned} d\overline{X}^{\varphi}(t) &= \frac{1}{S^0(t)}dX^{\varphi}(t) + X^{\varphi}(t)d\frac{1}{S^0(t)} + dX^{\varphi}(t)d\frac{1}{S^0(t)} \\ &= \lambda\varphi_1(t)\overline{S}^1(t)dt + \sigma\varphi_1(t)\overline{S}^1(t)dW(t), \qquad t \in [0,T]. \end{aligned}$$

Now the $\mathbb{Q}$-Brownian motion B may be written as

$$dB(t) = \theta dt + dW(t), \qquad t \in [0,T], \qquad B(0) = 0,$$

which together with $\lambda = \sigma\theta$ implies that

$$\begin{aligned} d\overline{X}^{\varphi}(t) &= \lambda\varphi_1(t)\overline{S}^1(t)dt + \sigma\varphi_1(t)\overline{S}^1(t)dW(t) \\ &= \left[\lambda\varphi_1(t)\overline{S}^1(t) - \sigma\theta\varphi_1(t)\overline{S}^1(t)\right]dt + \sigma\varphi_1(t)\overline{S}^1(t)dB(t) \\ &= \sigma\varphi_1(t)\overline{S}^1(t)dB(t), \qquad t \in [0,T], \end{aligned}$$

which proves the first part of the claim. By construction of the Itō integral, we can find an increasing sequence $\{\tau_k\}_{k\in\mathbb{N}}$ of stopping times such that each

$$\overline{X}^{\varphi}(\cdot \wedge \tau_k) = \overline{X}^{\varphi}(0) + (\sigma\varphi_1\overline{S}^1) \bullet B(\cdot \wedge \tau_k), \qquad k \in \mathbb{N},$$

is a martingale. Moreover, since φ is admissible, $\overline{X}^{\varphi}$ is lower bounded. Thus, for any $s,t \in [0,T]$ with $s \leq t$, an application of Fatou's lemma and continuity of $\overline{X}$ imply that

$$\overline{X}^{\varphi}(s) = \liminf_{k\to\infty} \overline{X}^{\varphi}(s \wedge \tau_k) = \liminf_{k\to\infty} \mathbb{E}_s\left[\overline{X}^{\varphi}(t \wedge \tau_k)\right] \geq \mathbb{E}_s\left[\overline{X}^{\varphi}(t)\right],$$

i.e. $\overline{X}^{\varphi}$ is indeed a supermartingale. □

The previous result sheds some more light on the restriction to admissible trading strategies: The lower boundedness of the discounted wealth process guarantees that it is **at least a supermartingale**, which in the absence of the lower bound need not be true. Moreover, from the representation

$$d\overline{X}^{\varphi}(t) = \sigma\varphi_1(t)\overline{S}^1(t)dB(t), \qquad t \in [0,T],$$

we immediately find that $\overline{X}^{\varphi}$ is even an honest $\mathbb{Q}$-martingale provided that $\sigma\varphi_1\overline{S}^1 \in \mathcal{I}^2$, which is equivalent to $\varphi_1 S^1 \in \mathcal{I}^2$.

Since not all discounted wealth processes are honest $\mathbb{Q}$-martingales in the Black-Scholes model, it makes sense to give these strategies a special name.

Definition 5.36 (Efficient Strategies)**.** An admissible trading strategy φ is called **efficient** provided that

$$\overline{X}^{\varphi} \text{ is a martingale under } \mathbb{Q}. \qquad \diamond$$

The idea behind **efficiency** is the following: Suppose that φ and ψ are admissible strategies with $X^{\varphi}(T) = X^{\psi}(T)$. If φ is efficient, then

$$\overline{X}^{\varphi} \text{ is a martingale, while } \overline{X}^{\pi} \text{ is, in general, only a supermartingale.}$$

In particular, it holds that

$$X^{\varphi}(0) = \overline{X}^{\varphi}(0) = \mathbb{E}^{\mathbb{Q}}\big[\overline{X}^{\varphi}(T)\big] = \mathbb{E}^{\mathbb{Q}}\big[\overline{X}^{\psi}(T)\big] \le \overline{X}^{\psi}(0) = X^{\psi}(0),$$

i.e. the efficient strategy φ is able to attain the same terminal wealth as ψ, but with (possibly) less initial wealth.

With Lemma 5.35 (Discounted Wealth Processes under EMMs) at hand, we are now ready to prove that the Black-Scholes model is **free of arbitrage** opportunities.

Theorem 5.37 (No Arbitrage in the Black-Scholes Model)**.** *There are no arbitrage opportunities in the Black-Scholes model.* $\diamond$

Proof. We argue by contradiction and assume that φ is an arbitrage opportunity, i.e. a self-financing, admissible trading strategy with

$$X^{\varphi}(0) = 0, \qquad X^{\varphi}(T) \ge 0, \qquad \mathbb{P}\big[X^{\varphi}(T) > 0\big] > 0.$$

Since $S^0 > 0$ and $\mathbb{Q} \sim \mathbb{P}$, this is equivalent to

$$\overline{X}^{\varphi}(0) = 0, \qquad \overline{X}^{\varphi}(T) \ge 0, \qquad \mathbb{Q}\big[\overline{X}^{\varphi}(T) > 0\big] > 0.$$

In particular, we must have $\mathbb{E}^{\mathbb{Q}}[\overline{X}^{\varphi}(T)] > 0$, which contradicts the supermartingale property of $\overline{X}^{\varphi}$ under $\mathbb{Q}$ as

$$0 = \overline{X}^{\varphi}(0) \ge \mathbb{E}^{\mathbb{Q}}\big[\overline{X}^{\varphi}(T)\big] > 0. \qquad \square$$

We can now turn to option pricing. Just as in discrete time, we identify options with positive random variables; see Definition 3.8.

Definition 5.38 (Option; Attainable; Replication Strategy)**.** Any $\mathbb{R}_+$-valued and $\mathfrak{F}(T) = \mathfrak{A}$-measurable random variable ξ is referred to as a (European) **option**. An admissible trading strategy φ with $X^\varphi(T) = \xi$ is referred to as a **replication strategy**, in which case we refer to ξ as being **attainable**. ◇

With this, we can turn to **risk neutral pricing** in the Black-Scholes model.

Theorem 5.39 (Risk Neutral Pricing in the Black-Scholes Model)**.** *Let ξ be an option with $\mathbb{E}^{\mathbb{Q}}[\xi] < \infty$ and define $C = \{C(t)\}_{t\in[0,T]}$ by*

$$C(t) \triangleq S^0(t)\mathbb{E}_t^{\mathbb{Q}}\Big[\frac{\xi}{S^0(T)}\Big], \qquad t \in [0,T].$$

Then there are no arbitrage opportunities in the extended financial market (S, C). Moreover, if φ is a replication strategy for ξ, it holds that

$$C(t) \leq X^\varphi(t), \qquad t \in [0,T],$$

with equality if and only if φ is efficient. ◇

Proof. Observe that

$$\overline{C}(s) = \frac{C(s)}{S^0(s)} = \mathbb{E}_s^{\mathbb{Q}}\Big[\frac{C(T)}{S^0(T)}\Big] = \mathbb{E}_s^{\mathbb{Q}}\Big[\mathbb{E}_t^{\mathbb{Q}}\Big[\frac{C(T)}{S^0(T)}\Big]\Big] = \mathbb{E}_s^{\mathbb{Q}}\Big[\frac{C(t)}{S^0(t)}\Big] = \mathbb{E}_s^{\mathbb{Q}}\big[\overline{C}(t)\big]$$

for all $s, t \in [0,T]$ with $s \leq t$, i.e. $\overline{C}$ is a $\mathbb{Q}$-martingale. Absence of arbitrage in the extended market (S, C) then follows as in Theorem 5.37 (No Arbitrage in the Black-Scholes Model). If φ is a replication strategy against ξ, we have $\overline{X}^\varphi(T) = \overline{C}(T)$ and thus, by Lemma 5.35 (Discounted Wealth Processes under EMMs), it follows that

$$X^\varphi(t) = S^0(t)\overline{X}^\varphi(t) \geq S^0(t)\mathbb{E}_t^{\mathbb{Q}}\big[\overline{X}^\varphi(T)\big] = S^0(t)\mathbb{E}_t^{\mathbb{Q}}\big[\overline{C}(T)\big] = C(t)$$

for all $t \in [0,T]$, with equality if and only if φ is efficient. □

With risk neutral pricing at hand, we are able to compute arbitrage free prices of any sufficiently integrable option ξ. Note, however, that given our results it is not clear if these are the only arbitrage free prices. This turns out to be true, i.e. the Black-Scholes model as defined in these notes turns out to be **complete**, but we shall not prove this result here.

Instead, let us compute the prices of **European calls and puts** and see if the resulting prices are in line with the prices in the **CRR(N) model** as computed in Theorem 5.3 (Black-Scholes Formula). For this, we recall that the European call option on S^1 with maturity T and strike price K is given by the positive random variable ξ_C with

$$\xi_C \triangleq \big(S^1(T) - K\big)_+.$$

Similarly, the European put on S^1 with the same maturity and strike can be modeled by the random variable ξ_P given by

$$\xi_P \triangleq \big(K - S^1(T)\big)_+.$$

Theorem 5.40 (Black-Scholes Formula Revisited)**.** *Let ξ_C and ξ_P be the payoffs of a European call and put on S^1 with maturity T and strike price K. Define $C = \{C(t)\}_{t\in[0,T]}$ and $P = \{P(t)\}_{t\in[0,T]}$ by*

$$C(t) \triangleq S^1(t)\Phi\big(d_+(t, S^1(t))\big) - Ke^{-r(T-t)}\Phi\big(d_-(t, S^1(t))\big), \qquad t \in [0,T],$$
$$P(t) \triangleq Ke^{-r(T-t)}\Phi\big(-d_-(t, S^1(t))\big) - S^1(t)\Phi\big(-d_+(t, S^1(t))\big), \quad t \in [0,T],$$

for $\Phi : \mathbb{R} \to [0,1]$ denoting the cumulative distribution function of a standard normal random variable and

$$d_+(t,s) \triangleq \frac{\log(s/K) + (r + \sigma^2/2)(T-t)}{\sigma\sqrt{T-t}}, \qquad (t,s) \in [0,T] \times \mathbb{R}_+,$$
$$d_-(t,s) \triangleq \frac{\log(s/K) + (r - \sigma^2/2)(T-t)}{\sigma\sqrt{T-t}}, \qquad (t,s) \in [0,T] \times \mathbb{R}_+.$$

Then C and P define arbitrage free prices of ξ_C and ξ_P, respectively. ◇

Proof. By Theorem 5.39 (Risk Neutral Pricing in the Black-Scholes Model), it suffices to show that

$$C(t) = S^0(t)\mathbb{E}_t^{\mathbb{Q}}\Big[\frac{\xi_C}{S^0(T)}\Big] \quad \text{and} \quad P(t) = S^0(t)\mathbb{E}_t^{\mathbb{Q}}\Big[\frac{\xi_P}{S^0(T)}\Big], \qquad t \in [0,T].$$

We only perform the computations in the case of the put, the case of the call being analogous. For this, we observe that

$$\begin{aligned}
S^0(t)\mathbb{E}_t^{\mathbb{Q}}\Big[\frac{\xi_P}{S^0(T)}\Big] &= e^{-r(T-t)}\mathbb{E}_t^{\mathbb{Q}}\Big[\big(K - S^1(T)\big)_+\Big] \\
&= S^1(t)\mathbb{E}_t^{\mathbb{Q}}\Big[\Big(\frac{Ke^{-r(T-t)}}{S^1(t)} - e^{-r(T-t)}\frac{S^1(T)}{S^1(t)}\Big)_+\Big] \\
&= S^1(t)\mathbb{E}_t^{\mathbb{Q}}\Big[\Big(\frac{Ke^{-r(T-t)}}{S^1(t)} - e^{(\lambda-\frac{1}{2}\sigma^2)(T-t)+\sigma(W(T)-W(t))}\Big)_+\Big] \\
&= S^1(t)\mathbb{E}_t^{\mathbb{Q}}\Big[\Big(\frac{Ke^{-r(T-t)}}{S^1(t)} - e^{-\frac{1}{2}\sigma^2(T-t)+\sigma(B(T)-B(t))}\Big)_+\Big],
\end{aligned}$$

where $B = \{B(t)\}_{t\in[0,T]}$ with $B(t) \triangleq W(t) + \theta t$, $t \in [0,T]$, is a $\mathbb{Q}$-Brownian motion by Theorem 5.34 (Girsanov). In particular, $B(T) - B(t)$ is independent of $\mathfrak{F}(S)$ under $\mathbb{Q}$ and $S^1(t)$ is $\mathfrak{F}(t)$-measurable. Thus

$$S^0(t)\mathbb{E}_t^{\mathbb{Q}}\Big[\frac{\xi_P}{S^0(T)}\Big] = S^1(t) f\Big(\frac{Ke^{-r(T-t)}}{S^1(t)}\Big),$$

where

$$f : \mathbb{R} \to \mathbb{R}_+, \qquad x \mapsto f(x) \triangleq \mathbb{E}^{\mathbb{Q}}\Big[\Big(x - e^{-\frac{1}{2}\sigma^2(T-t)+\sigma(B(T)-B(t))}\Big)_+\Big]$$

Since $B(T)-B(s)$ is normally distributed with zero mean and variance $T-t$, we can rewrite

$$f(x) = \mathbb{E}^{\mathbb{Q}}\Big[\Big(x - e^{-\frac{1}{2}\sigma^2(T-t)+\sigma\sqrt{T-t}Z}\Big)_+\Big], \qquad x \in \mathbb{R},$$

for a standard normal random variable Z. The same calculations as in the proof of Theorem 5.3 (Black-Scholes Formula) allow to compute this function f, showing that

$$\begin{aligned}
S^0(t)\mathbb{E}_t^{\mathbb{Q}}\Big[\frac{\xi_P}{S^0(T)}\Big] &= S^1(t) f\Big(\frac{Ke^{-r(T-t)}}{S^1(t)}\Big) \\
&= Ke^{-r(T-t)}\Phi\big(-d_+(t, S^1(t))\big) - S^1(t)\Phi\big(-d_-(t, S^1(t))\big),
\end{aligned}$$

and thus completing the proof. □

We note that for $t = 0$, we recover the result in Theorem 5.3 (Black-Scholes Formula). This justifies once again the Black-Scholes model as the **continuous time limit** of the CRR model.

We also make the observation that the excess return λ **does not enter** the pricing equations. In other words, in terms of pricing, the choice of λ does not matter at all. This is one of the reasons why the Black-Scholes model is so popular. In practice, estimating the parameters r (interest rate) and σ (volatility) is quite feasible, whereas the estimation of λ is statistically very unstable.

As a matter of fact, in practice, nobody believes in the Black-Scholes model, because it seems too unrealistic that interest rate and volatility remain constant over time. On the other hand, everybody uses the Black-Scholes formula. E.g., in practice it is often the case that vanilla option prices are not quoted in monetary units, but rather in terms of the volatility parameter σ which make the Black-Scholes formula valid. This particular volatility parameter is referred to as the **implied volatility** of the option. More precisely, consider a vanilla option trading at a market price of $p > 0$. The implied volatility $\hat{\sigma} > 0$ is the volatility parameter in the Black-Scholes model for which the market price p coincides with the theoretical price given by Theorem 5.3 (Black-Scholes Formula).

The implied volatility hence provides a means to **check the validity** of the model. One can compute the implied volatility for a range of vanilla options with different strikes and maturities. If the resulting implied volatilities are approximately constant, the Black-Scholes model could be considered reasonable. In practice, however, this is usually not the case: The implied volatility as a function of the strike tends to look like a **smile** or **smirk**, which is why this effect is referred to as the **volatility smile**; see Figure 5.1.

In practice, a common approach is to choose the volatility parameter in the Black-Scholes model in a way to minimize the distance between theoretical and market prices. This process is referred to as **model calibration**. Note, however, that as the Black-Scholes model is oftentimes too simplistic, more general models are employed. We will learn about a few of them in **Mathematical Finance II**.

We conclude this chapter with a presentation of the original derivation of the

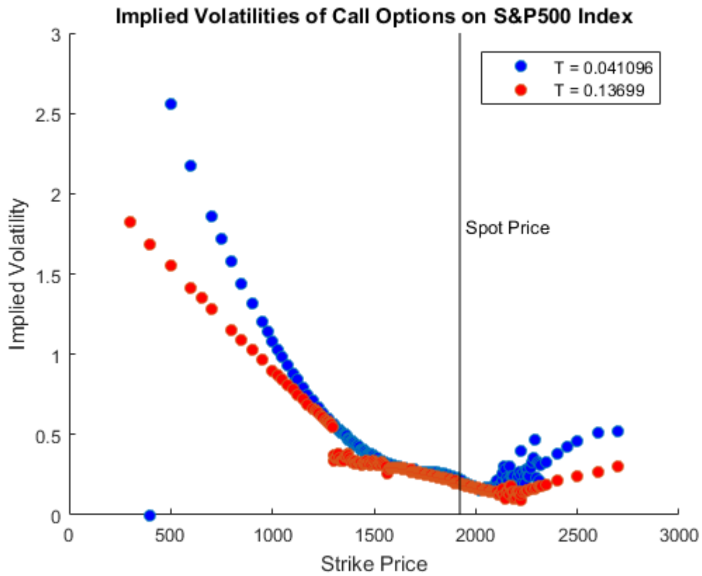

Figure 5.1. Implied volatility as a function of the strike for European calls written on the S&P500 index.

Black-Scholes formula. At the time the Black-Scholes formula was found, notions such as equivalent martingale measures were not around, and a more **analytic approach** towards option pricing was chosen. For this, we first give a heuristic argument to fix ideas.

Warning: The following arguments are purely heuristic and constitute in no way a rigorous mathematical derivation.

More precisely, let us take as given an option of the form

$$\xi \triangleq g\big(S^1(T)\big), \qquad \text{where} \qquad g : \mathbb{R}_+ \to \mathbb{R}_+, \qquad s \mapsto g(s),$$

i.e. the payoff of ξ is a **deterministic function** of $S^1(T)$ just as in the case of a European call or put. Let us now **assume** that an arbitrage free price $C = \{C(t)\}_{t\in[0,T]}$ is given by a function of S^1 as well, in the sense that

$$C(t) = u\big(t, S^1(t)\big), \qquad t \in [0,T].$$

Again, we observe that this is the case for European calls and puts. If the function u is sufficiently smooth, we can apply Itō's formula and obtain

$$\mathrm{d}C(t) = \Big[u_t\big(t, S^1(t)\big) + (r+\lambda)S^1(t)u_x\big(t, S^1(t)\big) + \frac{1}{2}\sigma^2 S^1(t)^2 u_{xx}\big(t, S^1(t)\big)\Big]\mathrm{d}t + \sigma S^1(t)u_x\big(t, S^1(t)\big)\mathrm{d}W(t).$$

Now consider a self-financing trading strategy $\varphi = (\varphi_0, \varphi_1, \varphi_C)$ in the extended market (S, C) and assume that $\varphi_C = -1$. In other words, we consider a portfolio in which the trader has **sold one unit** of the option. Since the trading strategy is self-financing and using the explicit expression for $\mathrm{d}C(t)$, we have

$$\begin{aligned}
\mathrm{d}X^\varphi(t) &= \varphi_0(t)\mathrm{d}S^0(t) + \varphi_1(t)\mathrm{d}S^1(t) + \varphi_C(t)\mathrm{d}C(t)\\
&= r\varphi_0(t)S^0(t)\mathrm{d}t\\
&\qquad + (r+\lambda)\varphi_1(t)S^1(t)\mathrm{d}t + \sigma\varphi_1(t)S^1(t)\mathrm{d}W(t)\\
&\qquad - \Big[\big[u_t\big(t, S^1(t)\big) + (r+\lambda)S^1(t)u_x\big(t, S^1(t)\big)\\
&\qquad\qquad + \frac{1}{2}\sigma^2 S^1(t)^2 u_{xx}\big(t, S^1(t)\big)\big]\mathrm{d}t + \sigma S^1(t)u_x\big(t, S^1(t)\big)\mathrm{d}W(t)\Big]\\
&= \Big[r\varphi_0(t)S^0(t) + (r+\lambda)\varphi_1(t)S^1(t) - u_t\big(t, S^1(t)\big)\\
&\qquad\qquad - (r+\lambda)S^1(t)u_x\big(t, S^1(t)\big) - \frac{1}{2}\sigma^2 S^1(t)^2 u_{xx}\big(t, S^1(t)\big)\Big]\mathrm{d}t\\
&\qquad + \sigma S^1(t)\big[\varphi_1(t) - u_x\big(t, S^1(t)\big)\big]\mathrm{d}W(t).
\end{aligned}$$

If the trader aims to construct a hedging strategy against the option, the aim should be to make X^φ as risk-free as possible. One way to achieve this is to choose

$$\varphi_1(t) = u_x\big(t, S^1(t)\big), \qquad t \in [0, T],$$

because this ensures that the $\mathrm{d}W(t)$ term in the dynamics of X^φ vanishes. More precisely, for this specific choice of φ_1, we obtain

$$\mathrm{d}X^\varphi(t) = \Big[r\varphi_0(t)S^0(t) - u_t\big(t, S^1(t)\big) - \frac{1}{2}\sigma^2 S^1(t)^2 u_{xx}\big(t, S^1(t)\big)\Big]\mathrm{d}t.$$

We also recall that the wealth $X^\varphi(t)$ is given by

$$\begin{aligned}
X^\varphi(t) &= \varphi_0(t)S^0(t) + \varphi_1(t)S^1(t) + \varphi_C(t)C(t)\\
&= \varphi_0(t)S^0(t) + u_x\big(t, S^1(t)\big)S^1(t) - u\big(t, S^1(t)\big),
\end{aligned}$$

and thus our equation for $\mathrm{d}X^\varphi(t)$ can also be written as

$$\mathrm{d}X^\varphi(t) = \Big[r\Big(X^\varphi(t) - u_x\big(t, S^1(t)\big)S^1(t) + u\big(t, S^1(t)\big)\Big)$$

$$
\begin{aligned}
&\quad - u_t\big(t, S^1(t)\big) - \frac{1}{2}\sigma^2 S^1(t)^2 u_{xx}\big(t, S^1(t)\big)\Big]\mathrm{d}t \\
&= rX^\varphi(t)\mathrm{d}t + \Big[-u_t\big(t, S^1(t)\big) + ru\big(t, S^1(t)\big) - rS^1(t)u_x\big(t, S^1(t)\big) \\
&\qquad - \frac{1}{2}\sigma^2 S^1(t)^2 u_{xx}\big(t, S^1(t)\big)\Big]\mathrm{d}t.
\end{aligned}
$$

Since X^φ evolves risk-free, in order to avoid arbitrage, the second $\mathrm{d}t$ integral should vanish, which requires

$$
u_t(t,x) - ru(t,x) + rxu_x(t,x) + \frac{1}{2}\sigma^2 x^2 u_{xx}(t,x) = 0 \tag{5.10}
$$

for all $(t,x) \in [0,T) \times (0,+\infty)$. That is to say that the option price u should solve a **partial differential equation** (PDE). The PDE is complemented with the terminal condition

$$
u(T,x) = g(x), \qquad x \in (0,+\infty),
$$

which guarantees that $C(T) = u(T, S^1(T)) = g(S^1(T)) = \xi$. Moreover, to ensure unique solvability of the PDE, one has to specify suitable growth condition on u as $x \to +\infty$.

While this argument is not rigorous in that it is not guaranteed that solving the PDE yields an arbitrage free price, we can make a few observations:

- The arguments do not rely on any of the machinery we have developed so far. In particular there is **no mention of EMMs**.
- Observe that the **excess return** parameter λ disappears in the derivation, i.e. it does not show up in the PDE. Thus, the option price will be independent of λ.
- The derivation automatically yields a replication strategy for the option, namely
  $$\varphi_1(t) = u_x\big(t, S^1(t)\big), \qquad t \in [0,T].$$
  This strategy is referred to as the **Δ-hedging strategy** in the literature. The fact that it delivers a replication strategy is a significant **advantage** of this approach over the risk neutral pricing methodology.
- The disadvantage of this approach is its applicability: It is only possible to apply it to options given in terms of **deterministic functions** of S^1 and having a sufficiently **smooth price** to apply Itō's formula.

▷ One can show that in case of European calls and puts, the prices given in Theorem 5.40 (Black-Scholes Formula Revisited) **solve the PDE**. The calculations are, however, a bit tedious so we skip them here.

The remainder of this chapter is dedicated to turning this heuristic derivation into rigorous mathematics. To begin with, we have to introduce the notion of **stochastic differential equations**. In what follows, we fix two deterministic functions

$$b : \mathbb{R} \to \mathbb{R} \qquad \text{and} \qquad a : \mathbb{R} \to \mathbb{R}.$$

We are interested in making sense of **stochastic differential equations** (SDEs) of the form

$$\mathrm{d}X(t) = b\big(X(t)\big)\mathrm{d}t + a\big(X(t)\big)\mathrm{d}W(t), \qquad t \in [t_0, T], \qquad X(t_0) = \eta.$$

Here, $t_0 \in [0, T)$ plays the role of the initial time whereas η is some $\mathfrak{F}(t_0)$-measurable random variable playing the role of the initial value of the SDE. Our first aim is to define what it means for a process $X = \{X(t)\}_{t\in[t_0,T]}$ to be a solution of such an equation. Before doing so, however, we once again take note of the fact that calling such an equation a **differential** equation is an abuse of wording, as it actually is a stochastic **integral** equation.

Definition 5.41 ((Square-Integrable) Solution of an SDE)**.** Let $t_0 \in [0, T)$ and η be an $\mathfrak{F}(t_0)$-measurable random variable. We say that a continuous adapted process $X = \{X(t)\}_{t\in[t_0,T]}$ is a **solution** of the SDE

$$\mathrm{d}X(t) = b\big(X(t)\big)\mathrm{d}t + a\big(X(t)\big)\mathrm{d}W(t), \quad t \in [t_0, T], \quad X(t_0) = \eta. \tag{5.11}$$

provided that

$$\int_{t_0}^{T} \big|b\big(X(t)\big)\big| + \big|a\big(X(t)\big)\big|^2 \mathrm{d}t < +\infty$$

and it holds that

$$X(t) = \eta + \int_{t_0}^{t} b\big(X(s)\big)\mathrm{d}s + \big(a(X)\mathbb{1}_{(t_0,T]}\big) \bullet W(t), \qquad t \in [0, T].$$

Moreover, we say that X is a **square-integrable solution** if, in addition,

$$\mathbb{E}\Big[\sup_{t\in[t_0,T]} |X(t)|^2\Big] < +\infty. \qquad \diamond$$

Note that square-integrable solutions can only exist if the initial condition is square-integrable as well since

$$\mathbb{E}\big[|\eta|^2\big] = \mathbb{E}\big[|X(t_0)|^2\big] \leq \mathbb{E}\Big[\sup_{t\in[t_0,T]} |X(t)|^2\Big] < +\infty.$$

The next natural question to ask is under which conditions a stochastic differential equation admits a solution. Very similar to the case of non-stochastic, i.e. **ordinary differential equations**, a Lipschitz condition on b and a suffices. Since the proof is a bit out of scope, we omit it here.

Theorem 5.42 (Existence and Uniqueness for SDEs). *Suppose that b and a are Lipschitz continuous, i.e. suppose there exists a constant $L > 0$ such that*

$$|b(x) - b(y)| + |a(x) - a(y)| \leq L|x - y|, \qquad x, y \in \mathbb{R}.$$

Fix $t_0 \in [0, T)$ and let η be $\mathfrak{F}(t_0)$-measurable with $\mathbb{E}[|\eta|^2] < +\infty$. Then there exists a unique square-integrable solution $X = \{X(t)\}_{t\in[0,T]}$ of the SDE

$$\mathrm{d}X(t) = b\big(X(t)\big)\mathrm{d}t + a\big(X(t)\big)\mathrm{d}W(t), \qquad t \in [t_0, T], \qquad X(t_0) = \eta. \qquad \diamond$$

In the previous theorem, uniqueness is to be understood **pathwise**. More precisely, if $\hat{X} = \{\hat{X}(t)\}_{t\in[0,T]}$ is another solution of the SDE, then $X = \hat{X}$.

Note that we have encountered stochastic differential equations before. For example, we saw that S^0 and S^1 solve the SDEs

$$\mathrm{d}S^0(t) = rS^0(t)\mathrm{d}t, \qquad t \in [0, T], \qquad S^0(0) = 1,$$

as well as

$$\mathrm{d}S^1(t) = (r + \lambda)S^1(t)\mathrm{d}t + \sigma S^1(t), \qquad t \in [0, T], \qquad S^1(0) = s,$$

respectively. Observe that, in particular, these equations fit into the setting of Theorem 5.42 (Existence and Uniqueness for SDEs).

Let us now turn towards **partial differential equations** (PDEs). For this, we fix a subset $\mathcal{O} \subseteq \mathbb{R}$. Typically, we will choose $\mathcal{O} = (0, +\infty)$ or $\mathcal{O} = \mathbb{R}$. We furthermore take as given continuous functions

$$k : \mathcal{O} \to \mathbb{R} \qquad \text{and} \qquad g : \mathcal{O} \to \mathbb{R}$$

and assume that k is bounded. We are subsequently interested in PDEs of the form

$$u_t(t,x) + \mathcal{L}[u](t,x) - k(x)u(t,x) = 0, \quad (t,x) \in [0,T) \times \mathcal{O}, \tag{5.12}$$

with terminal condition

$$u(T,x) = g(x), \qquad x \in \mathcal{O}, \tag{5.13}$$

where $\mathcal{L}$ denotes the so-called **infinitesimal generator** corresponding to the general SDE in Equation (5.11) and takes the form

$$\mathcal{L}[u](t,x) \triangleq b(x)u_x(t,x) + \frac{1}{2}a(x)^2 u_{xx}(t,x), \qquad (t,x) \in [0,T) \times \mathcal{O}.$$

Note that the PDE above generalizes the PDE we derived heuristically in Equation (5.10). We say that a function

$$u : [0,T] \times \mathcal{O} \to \mathbb{R}, \qquad (t,x) \mapsto u(t,x)$$

is a **classical solution** of the PDE if u is continuous on $[0,T] \times \mathcal{O}$, once continuously differentiable in its first and twice continuously differentiable in its second argument on $[0,T) \times \mathcal{O}$ and satisfies both Equation (5.12) and Equation (5.13). With this, we have the following remarkable result linking solutions of PDEs to solutions of SDEs.

Theorem 5.43 (Feynman-Kac Representation). *Suppose that u is a classical solution of the PDE in Equation* (5.12) *satisfying the terminal condition in Equation* (5.13) *such that there exists a constant $C > 0$ with*

$$|u(t,x)| \le C(1 + |x|^2), \qquad (t,x) \in [0,T] \times \mathcal{O}.$$

Fix moreover $t_0 \in [0,T)$ and let η be a $\mathfrak{F}(t_0)$-measurable and $\mathcal{O}$-valued random variable with $\mathbb{E}[|\eta|^2] < \infty$. Assume that there exists an $\mathcal{O}$-valued square-integrable solution of the SDE

$$\mathrm{d}X(t) = b\big(X(t)\big)\mathrm{d}t + a\big(X(t)\big)\mathrm{d}W(t), \qquad t \in [t_0,T], \qquad X(t_0) = \eta.$$

Then $u(t_0,\eta)$ admits the representation

$$u(t_0,\eta) = \mathbb{E}_{t_0}\Big[g\big(X(T)\big)e^{-\int_{t_0}^T k(X(t))\mathrm{d}t}\Big]. \qquad \diamond$$

Proof. Define $\beta = \{\beta(t)\}_{t\in[t_0,T]}$ by

$$\beta(t) \triangleq e^{-\int_{t_0}^{t} k(X(s))\mathrm{d}s}, \qquad t \in [t_0, T].$$

Since k is bounded, $k(X)$ is integrable and it follows from Itō's formula that β is an Itō process with

$$\mathrm{d}\beta(t) = -k\big(X(t)\big)\beta(t)\mathrm{d}t, \qquad t \in [t_0, T], \qquad \beta(t_0) = 1.$$

Also by Itō's formula, for any $T_0 \in [t_0, T)$, we have

$$\mathrm{d}u\big(t, X(t)\big) = \big[u_t\big(t, X(t)\big) + \mathcal{L}[u]\big(t, X(t)\big)\big]\mathrm{d}t \\ + a\big(X(t)\big)u_x\big(t, X(t)\big)\mathrm{d}W(t), \quad t \in [t_0, T_0].$$

Since u is a solution of the PDE, this rewrites as

$$\mathrm{d}u\big(t, X(t)\big) = \big[k\big(X(t)\big)u\big(t, X(t)\big)\big]\mathrm{d}t + a\big(X(t)\big)u_x\big(t, X(t)\big)\mathrm{d}W(t), \quad t \in [t_0, T_0].$$

Itō's product rule then yields

$$\begin{aligned} \mathrm{d}\big(\beta u(\cdot, X)\big)(t) &= \beta(t)\mathrm{d}u\big(t, X(t)\big) + u\big(t, X(t)\big)\mathrm{d}\beta(t) \\ &= \beta(t)a\big(X(t)\big)u_x\big(t, X(t)\big)\mathrm{d}W(t), \qquad t \in [t_0, T_0]. \end{aligned}$$

Now let $\{\tau_k\}_{k\in\mathbb{N}}$ be an increasing sequence of stopping times converging to $+\infty$ such that each of the stopped Itō integrals

$$\beta a(X)u_x(\cdot, X)\mathbb{1}_{(t_0,T]} \bullet W(t \wedge \tau_k), \qquad t \in [t_0, T_0],\ k \in \mathbb{N},$$

is a martingale. Setting $T_k \triangleq T_0 \wedge \tau_k$ and writing $\mathrm{d}\big(\beta u(\cdot, X)\big)(t)$ in integral form, we find that

$$\beta(T_k)u\big(T_k, X(T_k)\big) - u\big(t_0, X(t_0)\big) = \beta a(X)u_x(\cdot, X)\mathbb{1}_{(t_0,T]} \bullet W(T_k).$$

Rearranging and taking conditional expectations shows that

$$u(t_0, \eta) = u\big(t_0, X(t_0)\big) = \mathbb{E}_{t_0}\Big[\beta(T_k)u\big(T_k, X(T_k)\big)\Big],$$

where the Itō intregral vanishes due to its martingale property. Since $k(X)$ is bounded, so is β. Moreover,

$$\sup_{T_0\in[t_0,T)}\ \sup_{k\in\mathbb{N}}\big|u\big(T_k, X(T_k)\big)\big| \le \sup_{t\in[t_0,T]}\big|u\big(t, X(t)\big)\big| \le C\big(1 + \sup_{t\in[t_0,T]}|X(t)|^2\big),$$

and the rightmost side is finite in expectation since X is a square-integrable solution. We may hence apply dominated convergence to arrive at

$$\begin{aligned} u(t_0, \eta) &= \lim_{T_0 \uparrow T} \lim_{k \to \infty} \mathbb{E}_{t_0}\Big[\beta(T_k) u\big(T_k, X(T_k)\big)\Big] \\ &= \lim_{T_0 \uparrow T} \mathbb{E}_{t_0}\Big[\beta(T_0) u\big(T_0, X(T_0)\big)\Big] \\ &= \mathbb{E}_{t_0}\Big[\beta(T) u\big(T, X(T)\big)\Big] = \mathbb{E}_{t_0}\Big[\beta(T) g\big(X(T)\big)\Big], \end{aligned}$$

and the proof is complete by the definition of β. □

The **Feynman-Kac representation** can be used to make the PDE pricing argument rigorous.

Theorem 5.44 (PDE Pricing). *Let $\mathcal{O} = (0, +\infty)$, suppose that $g : \mathcal{O} \to \mathcal{O}$ is continuous, and let u be a classical solution of*

$$u_t(t,x) + rxu_x(t,x) + \frac{1}{2}\sigma^2 x^2 u_{xx}(t,x) - ru(t,x) = 0, \quad (t,x) \in [0,T) \times \mathcal{O},$$

with terminal condition

$$u(T,x) = g(x), \qquad x \in \mathcal{O}.$$

Suppose moreover that there exists $C > 0$ such that

$$|u(t,x)| \leq C(1 + |x|^2), \qquad (t,x) \in [0,T] \times \mathcal{O}.$$

Then an arbitrage free price of the European option $\xi \triangleq g(S^1(T))$ is given by $C = \{C(t)\}_{t \in [0,T]}$ with

$$C(t) \triangleq u\big(t, S^1(t)\big), \qquad t \in [0,T]. \quad \diamond$$

Proof. Note that, for each $t_0 \in [0,T)$, the process $S^1 = \{S^1(t)\}_{t \in [t_0,T]}$ is a (square-integrable) solution of the SDE

$$\mathrm{d}S^1(t) = rS^1(t)\mathrm{d}t + \sigma S^1(t)\mathrm{d}B(t), \qquad t \in [t_0, T],$$

with (square-integrable) initial value $S^1(t_0)$ at time initial t_0 and where $B = \{B(t)\}_{t\in[0,T]}$ is a Brownian motion under $\mathbb{Q}$. The infinitesimal generator associated with this SDE is

$$\mathcal{L}[u](t,x) = rxu_x(t,x) + \frac{1}{2}\sigma^2 x^2 u_{xx}(t,x), \qquad (t,x) \in [0,T) \times \mathcal{O}.$$

Hence we can apply Theorem 5.43 (Feynman-Kac Representation) to obtain

$$C(t_0) = u\big(t_0, S^1(t_0)\big) = \mathbb{E}^{\mathbb{Q}}_{t_0}\Big[g\big(S^1(T)\big)e^{-\int_{t_0}^T r\mathrm{d}t}\Big] = S^0(t_0)\mathbb{E}^{\mathbb{Q}}_{t_0}\Big[\frac{\xi}{S^0(T)}\Big]$$

for all $t_0 \in [0,T]$ and the proof is completed by applying Theorem 5.39 (Risk Neutral Pricing in the Black-Scholes Model). □

This is it. Let us wrap up these notes by appreciating the last results. The Feynman-Kac representation builds an interesting bridge between the mathematical fields of analysis (partial differential equations) and probability theory (stochastic processes) as it opens up the door to study PDEs by **probabilistic techniques**. Moreover, we have seen the close connection between probability theory and mathematical finance throughout the entirety of these notes.

As the following example shows, keeping such connections in mind can be quite useful. A classical problem in PDE theory is the following: Suppose that u_1 and u_2 are classical solutions of the PDE in Theorem 5.44 (PDE Pricing) with terminal conditions g_1 and g_2, respectively. Suppose furthermore that $g_1 \le g_2$, is it then true that $u_1 \le u_2$ as well? Results of this form are useful to obtain uniqueness of solutions, since $g_1 = g_2$ must then imply $u_1 = u_2$, i.e. there is only one solution given the terminal condition g_1. Now from a **probabilistic point of view**, the proof of this statement is straightforward as Theorem 5.44 (PDE Pricing) shows that

$$u_1(t,x) = e^{-r(T-t)}\mathbb{E}^{\mathbb{Q}}_t\big[g_1\big(S^1(T)\big)\big] \le e^{-r(T-t)}\mathbb{E}^{\mathbb{Q}}_t\big[g_2\big(S^1(T)\big)\big] = u_2(t,x).$$

Finally, taking on the **mathematical finance** perspective, the problem is only worth a weary smile: Since g_1 and g_2 are payoffs of options and u_1 and u_2 are their prices, absence of arbitrage dictates that u_1 must be smaller than u_2. Intriguing, right?

Appendix A — Stopping Times and Hitting Times

Throughout this section, we work on a general (not necessarily finite) probability space $(\Omega, \mathfrak{A}, \mathbb{P})$. Moreover, we fix a general time index set $\mathcal{T}$ as well as a filtration $\mathfrak{F} = \{\mathfrak{F}(t)\}_{t \in \mathcal{T}}$.

A.1 Stopping Times and the Sigma Field of the τ-Past

The aim of this section is to give a brief introduction to stopping times, which are random times which, in a sense, **do not take into account future information**. Let us try to clarify what we mean by this by considering an example from mathematical finance: Clearly, a well-defined time to exercise an American option would be obtained by a rule of the form *'Exercise when S^1 has risen two trading dates in a row.'*, i.e.

$$\tau(\omega) \triangleq \min\big\{t \in \{2, \dots, T\} : S^1(t, \omega) \geq S^1(t-1, \omega) \geq S^1(t-2, \omega)\big\}, \quad \omega \in \Omega.$$

This is a well-defined exercise date since at any time $t = 0, 1, \dots, T$ we can decide if we should exercise by looking at the prices of S^1 we have observed up to time t and there is no need to know future prices.

On the other hand, we would like to rule out exercise dates of the form *'Exercise when S^1 attains its largest value for the first time.'*, i.e.

$$\rho(\omega) \triangleq \min\Big\{\arg\max\big\{S^1(t, \omega) : t = 0, 1, \dots, T\big\}\Big\}, \qquad \omega \in \Omega.$$

The point here is that at time $t < T$ we know $S^1(0), S^1(1), \dots, S^1(t)$, but we do not know the future prices at times $t+1, \dots, T$. In other words, we cannot tell at time t if S^1 has already attained its overall maximum, so we do not know if we should exercise now or not.

The difference between τ and ρ is that at any time $t = 0, 1, \ldots, T$, we are able to decide if τ has already occurred or not, whereas this is not the case for ρ. Recalling that $\mathfrak{F}(t)$ models the information available at time t, this means that $\{\tau \leq t\} \in \mathfrak{F}(t)$, whereas the same is in general not true for ρ.

Definition A.1 (Stopping Time)**.** A mapping

$$\tau : \Omega \to \mathcal{T} \cup \{+\infty\}, \qquad \omega \mapsto \tau(\omega),$$

is referred to as a **stopping time** with respect to $\mathfrak{F}$ if

$$\{\tau \leq t\} \in \mathfrak{F}(t), \qquad t \in \mathcal{T}. \qquad \diamond$$

Observe that we allow τ to take the value $+\infty$ in case we run into a stopping rule which has the possibility of never happening.[1] We furthermore observe that we do not require τ to be a random variable as this is implied by the condition $\{\tau \leq t\} \in \mathfrak{F}(t) \subseteq \mathfrak{A}$ for all $t \in \mathcal{T}$. Finally, we note that every deterministic time $t_0 \in \mathcal{T}$ is also a stopping time since

$$\{t_0 \leq t\} = \begin{cases} \emptyset & \text{if } t_0 > t, \\ \Omega & \text{if } t_0 \leq t, \end{cases} \qquad t \in \mathcal{T}.$$

In the case when τ takes at most countably many values (which holds, in particular, if $\mathcal{T}$ is at most countable), there is a weaker condition to check if a random time is a stopping time.

Lemma A.2 (Characterization of Discrete Stopping Times)**.** *Let $\tau : \Omega \to \mathcal{T} \cup \{\infty\}$ be such that $\tau(\Omega) \triangleq \{\tau(\omega) : \omega \in \Omega\}$ is at most countable and let $\mathfrak{F} = \{\mathfrak{F}(t)\}_{t \in \mathcal{T}}$ be a filtration. Then τ is a stopping time with respect to $\mathfrak{F}$ if and only if*

$$\{\tau = t\} \in \mathfrak{F}(t) \qquad \textit{for all } t \in \tau(\Omega) \textit{ with } t < \infty. \qquad \diamond$$

[1] Think e.g. of the stopping time *'The first time all students fail the course on Mathematical Finance.'* — hopefully this will never happen, i.e. hopefully this stopping time is equal to infinity.

Proof. Step 1: Suppose that $\{\tau = t\} \in \mathfrak{F}(t)$ for all $t \in \tau(\Omega)$ with $t < \infty$ and let $s \in \mathcal{T}$. We have to check that $\{\tau \leq s\} \in \mathfrak{F}(s)$, in which case we conclude that τ is a stopping time. If $s = \infty$, then $\{\tau \leq s\} = \Omega \in \mathfrak{F}(s)$, so suppose that $s < \infty$. In that case, we have

$$\{\tau \leq s\} = \bigcup_{t \in \tau(\Omega),\ t \leq s} \{\tau = t\}.$$

Now for every $t \in \tau(\Omega)$ with $t \leq s$, we have $\{\tau = t\} \in \mathfrak{F}(t)$ by assumption and hence, since $t \leq s$ implies $\mathfrak{F}(t) \subseteq \mathfrak{F}(s)$, we find that $\{\tau = t\} \in \mathfrak{F}(s)$. But now, since $\tau(\Omega)$ is countable, this implies that $\{\tau \leq s\} \in \mathfrak{F}(s)$.

Step 2: Suppose that τ is a stopping time and let $s \in \tau(\Omega)$ with $s < \infty$. We have to show that $\{\tau = s\} \in \mathfrak{F}(s)$. For this, we first observe that, for every $t \in \mathcal{T}$ with $t \leq s$, we have $\{\tau \leq t\} \in \mathfrak{F}(t) \subseteq \mathfrak{F}(s)$. But then, again since $\tau(\Omega)$ is countable, we conclude since

$$\{\tau = s\} = \{\tau \leq s\} \setminus \bigcup_{t \in \tau(\Omega),\ t < s} \{\tau \leq t\} \in \mathfrak{F}(s).$$

□

Lemma A.3 (Maxima and Minima of Stopping Times). *Let τ and σ be two stopping times. Then*

$$\tau \wedge \sigma : \Omega \to \mathcal{T} \cup \{+\infty\}, \qquad \tau \wedge \sigma(\omega) \triangleq \min\{\tau(\omega), \sigma(\omega)\},$$
$$\tau \vee \sigma : \Omega \to \mathcal{T} \cup \{+\infty\}, \qquad \tau \vee \sigma(\omega) \triangleq \max\{\tau(\omega), \sigma(\omega)\},$$

are stopping times with respect to $\mathfrak{F}$ as well. ◇

Proof. For all $t \in \mathcal{T}$ it holds that

$$\{\tau \wedge \sigma \leq t\} = \{\tau \leq t\} \cup \{\sigma \leq t\} \in \mathfrak{F}(t)$$

as well as

$$\{\tau \vee \sigma \leq t\} = \{\tau \leq t\} \cap \{\sigma \leq t\} \in \mathfrak{F}(t).$$

Thus both $\tau \wedge \sigma$ and $\tau \vee \sigma$ are stopping times. □

A natural question is to ask **which information is available** up to the stopping time τ. More precisely, if $F \in \mathfrak{A}$ is some event, we would like to know if F is known at time τ. In other words, at any given time $t \in \mathcal{T}$, we would like to know if F is known at time t, provided that τ has already occurred. Mathematically, this means that F is known at time τ if and only if $F \cap \{\tau \leq t\} \in \mathfrak{F}(t)$ for all $t \in \mathcal{T}$.

Definition A.4 (σ-Field of the τ-Past)**.** Let τ be a stopping time. Then

$$\mathfrak{F}(\tau) \triangleq \big\{F \in \mathfrak{A} : F \cap \{\tau \leq t\} \in \mathfrak{F}(t) \text{ for all } t \in \mathcal{T}\big\}$$

is called the ***σ*-field of the *τ*-past**. ◇

We have to be a bit careful with the previous definition as it is quite suggestive. First, the name '*σ-field of the τ-past*' suggests that $\mathfrak{F}(\tau)$ is a σ-field. This is indeed true and can easily be verified. Moreover, the notation $\mathfrak{F}(\tau)$ suggests some sort of **compatibility with the filtration** $\mathfrak{F}$. Of course, we should **never ever** think of $\mathfrak{F}(\tau)$ as a mapping $\omega \mapsto \mathfrak{F}(\tau(\omega))$. Nevertheless, $\mathfrak{F}(\tau)$ is compatible with the filtration in that $\mathfrak{F}(\tau) = \mathfrak{F}(t)$ if τ is constant and equal to t. Another property of the σ-field of the τ-past which justifies the notation $\mathfrak{F}(\tau)$ is that the mapping $\tau \mapsto \mathfrak{F}(\tau)$ respects the order on $\mathfrak{F}$ induced by the order on $\mathcal{T}$, i.e. if σ and τ are two stopping times with $\sigma \leq \tau$, then it holds that $\mathfrak{F}(\sigma) \subseteq \mathfrak{F}(\tau)$.

Lemma A.5 (Properties of the σ-Field of the τ-Past)**.** *Let σ and τ be stopping times. Then*

$$\mathfrak{F}(\tau \wedge \sigma) = \mathfrak{F}(\tau) \cap \mathfrak{F}(\sigma).$$

In particular, if $\sigma \leq \tau$, then $\mathfrak{F}(\sigma) \subseteq \mathfrak{F}(\tau)$. Finally, we have

$$\{\sigma \leq \tau\}, \{\sigma < \tau\}, \{\sigma = \tau\} \in \mathfrak{F}(\tau \wedge \sigma).$$ ◇

Proof. Step 1: We show that $\mathfrak{F}(\tau \wedge \sigma) \subseteq \mathfrak{F}(\tau) \cap \mathfrak{F}(\sigma)$. Let $F \in \mathfrak{F}(\tau \wedge \sigma)$ and $t \in \mathcal{T}$. Since $\{\sigma \leq t\} = \{\sigma \wedge \tau \leq t\} \cap \{\sigma \leq t\}$, it follows that

$$F \cap \{\sigma \leq t\} = \big(F \cap \{\sigma \wedge \tau \leq t\}\big) \cap \{\sigma \leq t\}.$$

As $F \in \mathfrak{F}(\tau \wedge \sigma)$, we have $F \cap \{\sigma \wedge \tau \leq t\} \in \mathfrak{F}(t)$, and, since σ is a stopping time, we see moreover that $\{\sigma \leq t\} \in \mathfrak{F}(t)$. Thus $F \cap \{\sigma \leq t\} \in \mathfrak{F}(t)$, implying that $F \in \mathfrak{F}(\sigma)$ since $t \in \mathcal{T}$ was chosen arbitrarily. Interchanging the roles of τ and σ furthermore yields $F \in \mathfrak{F}(\tau)$, i.e. $F \in \mathfrak{F}(\tau) \cap \mathfrak{F}(\sigma)$. This in turn implies $\mathfrak{F}(\tau \wedge \sigma) \subseteq \mathfrak{F}(\tau) \cap \mathfrak{F}(\sigma)$.

Step 2: We show that $\mathfrak{F}(\tau) \cap \mathfrak{F}(\sigma) \subseteq \mathfrak{F}(\tau \wedge \sigma)$. Let $F \in \mathfrak{F}(\tau) \cap \mathfrak{F}(\sigma)$ and $t \in \mathcal{T}$. Using $\{\tau \wedge \sigma \leq t\} = \{\tau \leq t\} \cup \{\sigma \leq t\}$ and $F \in \mathfrak{F}(\tau)$ and $F \in \mathfrak{F}(\sigma)$, we find

$$F \cap \{\tau \wedge \sigma \leq t\} = \big(F \cap \{\tau \leq t\}\big) \cup \big(F \cap \{\sigma \leq t\}\big) \in \mathfrak{F}(t).$$

However, this implies that $F \in \mathfrak{F}(\tau \wedge \sigma)$, i.e. $\mathfrak{F}(\tau) \cap \mathfrak{F}(\sigma) \subseteq \mathfrak{F}(\tau \wedge \sigma)$.

Step 3: We are left with showing that $\{\sigma < \tau\} \in \mathfrak{F}(\tau \wedge \sigma)$ since from this it follows that

$$\begin{aligned} \{\sigma \leq \tau\} &= \{\tau < \sigma\}^c \in \mathfrak{F}(\tau \wedge \sigma) \\ \{\sigma = \tau\} &= \{\sigma \leq \tau\} \setminus \{\sigma < \tau\} \in \mathfrak{F}(\tau \wedge \sigma). \end{aligned}$$

To see that $\{\sigma < \tau\} \in \mathfrak{F}(\tau \wedge \sigma)$, let $t \in \mathcal{T}$ be arbitrary and denote by $\mathcal{Q}$ a countable and dense subset of $\mathcal{T}$ with $t \in \mathcal{Q}$. Then it holds that

$$\begin{aligned} \{\sigma < \tau\} \cap \{\tau \leq t\} &= \bigcup\nolimits_{q \in \mathcal{Q}, q < t} \{\sigma \leq q < \tau \leq t\} \\ &= \bigcup\nolimits_{q \in \mathcal{Q}, q < t} \{\sigma \leq q\} \cap \big(\{\tau \leq t\} \setminus \{\tau \leq q\}\big). \end{aligned}$$

Since $\mathcal{Q}$ is countable, we must have $\{\sigma < \tau\} \cap \{\tau \leq t\} \in \mathfrak{F}(t)$, which implies that $\{\sigma < \tau\} \in \mathfrak{F}(\tau)$. Conversely, we have

$$\{\sigma < \tau\} \cap \{\sigma \leq t\} = \bigcup\nolimits_{q \in \mathcal{Q}, q \leq t} \{\sigma \leq q < \tau\} = \bigcup\nolimits_{q \in \mathcal{Q}, q \leq t} \{\sigma \leq q\} \cap \{\tau \leq q\}^c.$$

By countability of $\mathcal{Q}$, this implies that $\{\sigma < \tau\} \cap \{\sigma \leq t\} \in \mathfrak{F}(t)$ and hence $\{\sigma < \tau\} \in \mathfrak{F}(\sigma)$. But then $\{\sigma < \tau\} \in \mathfrak{F}(\tau) \cap \mathfrak{F}(\sigma) \subseteq \mathfrak{F}(\tau \wedge \sigma)$ by step 2. □

Recalling that each deterministic time $t \in \mathcal{T}$ is also a stopping time and using that $\mathfrak{F}(\tau \wedge t) \subseteq \mathfrak{F}(\tau)$ and $\mathfrak{F}(\tau \wedge t) \subseteq \mathfrak{F}(t)$, Lemma A.5 (Properties of the σ-Field of the τ-Past) implies that the events $\{\tau \leq t\}$, $\{\tau < t\}$, $\{\tau = t\}$ are contained in both $\mathfrak{F}(\tau)$ and $\mathfrak{F}(t)$.

A.2 Hitting Times and Stopped Processes

Thus far, we have not seen any nontrivial example of a stopping time and we have not drawn any connection to stochastic processes. This is about to change: The prime example of a stopping time, which we shall encounter over and over again, is the first time that a stochastic process enters into a measurable subset of its state space.

Definition A.6 (Hitting Time)**.** Let $X = \{X(t)\}_{t\in\mathcal{T}}$ be a stochastic process taking values in a measurable space $(S, \mathfrak{S})$ and let $B \in \mathfrak{S}$. Then the mapping

$$\tau_B : \Omega \to \mathcal{T} \cup \{+\infty\}, \qquad \tau_B(\omega) \triangleq \inf\big\{t \in \mathcal{T} : X(t, \omega) \in B\big\},$$

is called the (first) **hitting time** of B by X. ◇

For the remainder of this section, we fix a stochastic process $X = \{X(t)\}_{t\in\mathcal{T}}$ taking values in some measurable space $(S, \mathfrak{S})$. A priori, it is **not clear** if every hitting time of X is a stopping time with respect $\mathfrak{F}$ even though X is adapted. This turns out to be a very deep issue which goes far beyond the scope of this course (and is in general not true). Under some additional assumptions on X, $\mathfrak{F}$, and/or B, we can however give a positive answer to this question, and this will be sufficient for our purposes. The first, most basic, case is when X is a discrete time process.

Lemma A.7 (Hitting Times with Finite Time Index Sets)**.** *Fix $B \in \mathfrak{S}$ and assume that the time index set $\mathcal{T}$ is such that*

for each $t \in \mathcal{T}$, the set $\{s \in \mathcal{T} : s \leq t\}$ is finite.

Then the first hitting time τ_B of B by X is a stopping time. ◇

Proof. The condition that $\{s \in \mathcal{T} : s \leq t\}$ is finite for all $t \in \mathcal{T}$ implies that, for all $\omega \in \Omega$, we either have $\tau_B(\omega) = \infty$ or

$$\tau_B(\omega) = \inf\big\{t \in \mathcal{T} : X(t, \omega) \in B\big\} = \min\big\{t \in \mathcal{T} : X(t, \omega) \in B\big\} \in \mathcal{T}.$$

Let now $t \in \mathcal{T}$ with $t < \infty$. Then

$$\{\tau_B = t\} = \{X(t) \in B\} \cap \Big(\bigcap_{s\in\mathcal{T},\ s<t} \{X(s) \notin B\}\Big).$$

Since $\{s \in \mathcal{T} : s \leq t\}$ is finite and X is adapted, this shows that $\{\tau_B = t\} \in \mathfrak{F}(t)$ for all $t \in \mathcal{T}$ with $t < \infty$, and we conclude by Lemma A.2 (Characterization of Discrete Stopping Times). □

In Lemma A.7, the finiteness of the sets $\{s \in \mathcal{T} : s \leq t\}$ is needed to guarantee that the hitting time τ_B takes values in $\mathcal{T} \cup \{\infty\}$ (in particular $\mathcal{T}$ has a minimal element). Indeed, if for example $\mathcal{T} = \{1/n : n \in \mathbb{N}\}$ and we choose $B = S$, then $\tau_B(\omega) = 0$ for all $\omega \in \Omega$, but $0 \notin \mathcal{T}$.

For continuous time processes, the issue of whether hitting times are stopping times is more delicate. Imagine, e.g., that X is a continuous process and B is an open set. Then at time τ_B, the process X will be on the boundary of B, but not in B itself. So, whenever the process X is on the boundary of B, we need to be able to look an infinitesimal amount of time into the future for τ_B to be a stopping time (compare with Figure A.1), since otherwise we will not be able to decide whether we have to stop or not.

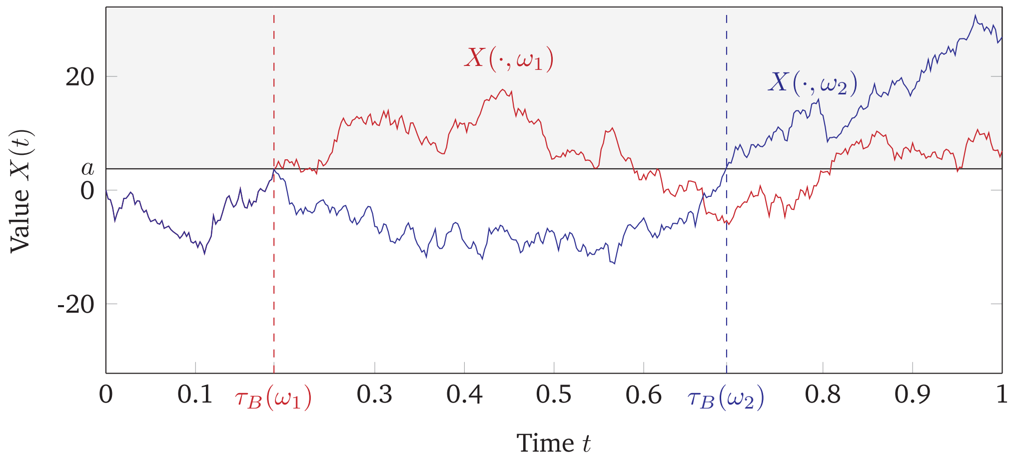

Figure A.1. Hitting time of the open set $B = (a, \infty)$ by a continuous process X. Observe that while both paths take the value a at time $\tau_B(\omega_1)$, only the path $X(\cdot, \omega_1)$ enters B at this time, whereas $X(\cdot, \omega_2)$ moves away from B. Hence, in order to decide if we should stop at any given time, we need to look an infinitesimal amount of time into the future.

Definition A.8 (Right Continuous Filtration)**.** We say that a filtration $\mathfrak{F} = \{\mathfrak{F}(t)\}_{t \in \mathcal{T}}$ is **right continuous** if

$$\mathfrak{F}(t+) \triangleq \bigcap_{s \in \mathcal{T},\ s > t} \mathfrak{F}(s) = \mathfrak{F}(t), \qquad t \in \mathcal{T}. \qquad \diamond$$

Right continuity of a filtration $\mathfrak{F} = \{\mathfrak{F}(t)\}_{t \in \mathcal{T}}$ is only interesting if the time index set is sufficiently rich. Indeed, if each $t \in \mathcal{T}$ has a distinct next element $t+1$ (which holds, e.g., for $\mathcal{T} = \mathbb{N}$ but not for $\mathcal{T} = \mathbb{Q}$), then $\mathfrak{F}(t+) = \mathfrak{F}(t+1)$ for all $t \in \mathcal{T}$ and hence the notion of right continuity of $\mathfrak{F}$ boils down to the requirement $\mathfrak{F}(t) = \mathfrak{F}(t+1)$ for all $t \in \mathcal{T}$.

Proposition A.9 (Hitting Times of Continuous Time Processes)**.** *Let $X = \{X(t)\}_{t\in[0,\infty)}$ be a stochastic process adapted to a filtration $\mathfrak{F} = \{\mathfrak{F}(t)\}_{t\in[0,\infty)}$. Suppose that S is a metric space with metric d and assume that $\mathfrak{S} = \mathfrak{B}(S)$ is the Borel σ-field on S. Then the hitting time τ_B of a set $B \in \mathfrak{S}$ by X is a stopping time with respect to $\mathfrak{F}$ in each of the following cases:*

(i) *X is continuous and B is closed.*

(ii) *X is right continuous, B is open, and $\mathfrak{F}$ is right continuous.* ◇

Proof. Step 1: Assume that X is continuous and B is closed. For every $t \in [0,\infty)$, we claim that

$$\{\tau_B \leq t\} = \{X(t) \in B\} \cup \Big(\bigcup_{s\in[0,t)} \{X(s) \in B\} \Big) \tag{A.1}$$

$$= \{X(t) \in B\} \cup \Big(\bigcap_{k=1}^{\infty} \bigcup_{q\in[0,t)\cap\mathbb{Q}} \big\{ d\big(X(q), B\big) \leq 1/k \big\} \Big). \tag{A.2}$$

From this and the adaptedness of X, it follows that $\{\tau_B \leq t\} \in \mathfrak{F}(t)$, and hence τ_B is a stopping time. To see that this representation of $\{\tau_B \leq t\}$ indeed holds, let us first fix $\omega \in \Omega$ with $\tau_B(\omega) < \infty$. Then the continuity of X and the closedness of B imply that $X(\tau_B(\omega), \omega) \in B$. But then $\omega \in \{\tau_B \leq t\}$ implies that there exists $s \in [0,t]$ such that $X(s,\omega) \in B$. On the other hand, if $\omega \in \{X(s) \in B\}$ for some $s \in [0,t]$, then $\tau_B(\omega) \leq s \leq t$. Thus Equation (A.1) is argued for. To see that Equation (A.2) holds, let us first assume that $\omega \in \{X(s) \in B\}$ for some $s \in [0,t)$. Then the continuity of $X(\cdot,\omega)$ implies that there exists a sequence $\{q_k\}_{k\in\mathbb{N}}$ in $[0,t) \cap \mathbb{Q}$ such that $d(X(q_k,\omega), X(s,\omega)) \leq 1/k$. But then, since $X(s,\omega) \in B$, we must necessarily have $d(X(q_k,\omega), B) \leq 1/k$. This shows that the set on the right hand side of Equation (A.1) is contained in the set in Equation (A.2). For the reverse containment, assume that $\omega \in \Omega$ is such that, for each $k \in \mathbb{N}$, there exists $q_k \in [0,t) \cap \mathbb{Q}$ such that $d(X(q_k,\omega), B) \leq 1/k$. Since the sequence

$\{q_k\}_{k\in\mathbb{N}} \subset [0,t]$ and $[0,t]$ is compact, we may without loss of generality assume that $s \triangleq \lim_{k\to\infty} q_k$ exists in $[0,t]$ (drop to a subsequence if $\{q_k\}_{k\in\mathbb{N}}$ does not converge). But then $X(s,\omega) \in B$ since by continuity of X we have

$$0 \le d\big(X(s,\omega), B\big) = \lim_{k\to\infty} d\big(X(q_k,\omega), B\big) \le \lim_{k\to\infty} 1/k = 0.$$

Step 2: Assume that X and $\mathfrak{F}$ are right continuous and B is open. Observe that, in this situation, we do not necessarily have $X(\tau_B) \in B$ on $\{\tau_B < \infty\}$, so we must be careful with events of the form $\{\tau_B = t\}$. Nevertheless, if $t \in [0,\infty)$, we have

$$\{\tau_B < t\} = \bigcup_{q\in[0,t)\cap\mathbb{Q}} \{X(q) \in B\} \in \mathfrak{F}(t). \tag{A.3}$$

Indeed, if $\omega \in \{X(q) \in B\}$ for some $q \in [0,t)\cap\mathbb{Q}$, then clearly $\tau_B(\omega) \le q < t$ and hence $\omega \in \{\tau_B < t\}$. On the other hand, if $\omega \in \{\tau_B < t\}$, then by definition of $\tau_B(\omega)$ there exists $s \in [0,t)$ such that $X(s,\omega) \in B$. Now let $\{q_k\}_{k\in\mathbb{N}}$ be a sequence in $[0,t)\cap\mathbb{Q}$ with $q_k > s$ and $q_k \downarrow s$ as $k \to \infty$. Then $X(q_k,\omega) \to X(s,\omega)$ by right continuity of X. But then, since B is open and $X(s,\omega) \in B$, we must have $X(q_k,\omega) \in B$ for all $k \in \mathbb{N}$ sufficiently large, i.e. $\omega \in \bigcup_{q\in[0,t)\cap\mathbb{Q}}\{X(q) \in B\}$, which establishes Equation (A.3). But then

$$\{\tau_B \le t\} = \bigcap_{q\in(t,\infty)\cap\mathbb{Q}} \{\tau_B < q\} = \bigcap_{q\in(t,s]\cap\mathbb{Q}} \{\tau_B < q\} \in \mathfrak{F}(s) \quad \text{for all } s > t$$

since the events $\{\tau_B < q\}$ are increasing in q. But then $\{\tau_B \le t\} \in \mathfrak{F}(t+) = \mathfrak{F}(t)$ by right continuity of $\mathfrak{F}$ and hence τ_B is a stopping time. □

The previous proposition extends immediately to processes $X = \{X(t)\}_{t\in\mathcal{T}}$ with time index set of the form $\mathcal{T} = [0,T]$ for some $T > 0$. Indeed, given X, we can simply define a new process $Y = \{Y(t)\}_{t\in[0,\infty)}$ by setting $Y(t) \triangleq X(t\wedge T)$ for all $t \in [0,\infty)$ and extend the filtration $\mathfrak{F} = \{\mathfrak{F}(t)\}_{t\in[0,T]}$ to a filtration on $[0,\infty)$ by setting $\mathfrak{F}(t) = \mathfrak{F}(T)$ for all $t \ge T$. Since X and Y coincide on $[0,T]$ and Y is constant outside of $[0,T]$, any hitting time of the process X coincides with any hitting time of the process Y. Thus Proposition A.9 remain valid also for the process X.

Now what are stopping times good for? Well, as the name suggests, they can be used to stop stochastic processes.

Definition A.10 (Stopped Process)**.** For any $\mathcal{T}$-valued (i.e. finite) stopping time τ, we set

$$X(\tau) : \Omega \to S, \qquad \omega \mapsto X(\tau, \omega) \triangleq X(\tau(\omega), \omega). \tag{A.4}$$

With this, if τ is an arbitrary $\mathcal{T} \cup \{+\infty\}$-valued stopping time, we refer to $X^\tau = X(\cdot \wedge \tau) = \{X(t \wedge \tau)\}_{t \in \mathcal{T}}$ as the **process X stopped at time τ**, or **stopped process** for short. ◇

We observe that Equation (A.4) is ill-posed if $\tau(\omega) = +\infty$ as $+\infty \notin \mathcal{T}$, which is why we restrict this part of the definition to finite stopping times. Next, for any $t \in \mathcal{T}$ and any $\mathcal{T} \cup \{+\infty\}$-valued stopping time τ, Lemma A.3 (Maxima and Minima of Stopping Times) implies that $t \wedge \tau$ is a finite stopping time and hence $X(t \wedge \tau)$ is well defined. Finally, the stopped process X^τ simply freezes the original process X at time τ; see Figure A.2.

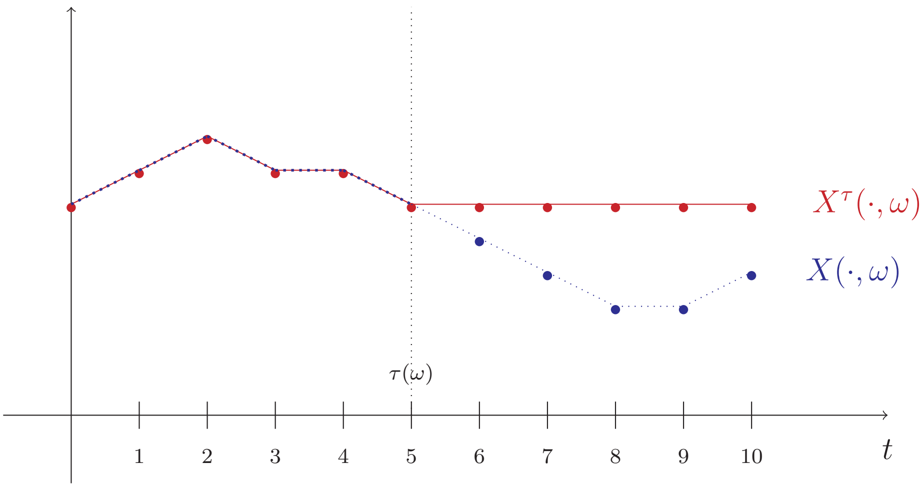

Figure A.2. Trajectories of X and the stopped process X^τ.

Now take another close look at Equation (A.4). Since $\omega \mapsto X(\tau(\omega), \omega)$ has a double dependence on ω, **it is not clear** without any additional conditions if $X(\tau) : \Omega \to S$ is a random variable and hence it is not clear if X^τ is a stochastic process. If X is adapted to $\mathfrak{F}$, then one would naturally expect $X^\tau(t) = X(t \wedge \tau)$ to be $\mathfrak{F}(t \wedge \tau)$-measurable for all $t \in \mathcal{T}$, i.e. one would expect the stopped process $X^\tau = \{X^\tau(t)\}_{t \in \mathcal{T}}$ to be adapted to the **stopped filtration** $\mathfrak{F}^\tau = \mathfrak{F}(\cdot \wedge \tau) = \{\mathfrak{F}(t \wedge \tau)\}_{t \in \mathcal{T}}$ (and hence also adapted to $\mathfrak{F}$ as $\mathfrak{F}(t \wedge \tau) \subset \mathfrak{F}(t)$, $t \in \mathcal{T}$). A **warning** though: this is not true in general! So under which conditions is the stopped process adapted? The answer is once again easy in discrete time, or, more generally, if the stopping time takes only countably many values.

Lemma A.11 (Measurability of Stopped Processes (Discrete Case)). *Let τ be a stopping time. If $\tau(\Omega)$ is countable, the stopped process X^τ is adapted to the stopped filtration $\mathfrak{F}^\tau = \{\mathfrak{F}^\tau(t)\}_{t\in\mathcal{T}}$ given by $\mathfrak{F}^\tau(t) \triangleq \mathfrak{F}(t\wedge\tau)$ for all $t\in\mathcal{T}$. In particular, X^τ is adapted to $\mathfrak{F}$ and, if τ is $\mathcal{T}$-valued, $X(\tau)$ is $\mathfrak{F}(\tau)$-measurable.* ◇

Proof. Assume that τ is $\mathcal{T}$-valued. We argue that in this case $X(\tau)$ is $\mathfrak{F}(\tau)$-measurable. From this, it follows immediately that for arbitrary stopping times τ, the random variable $X^\tau(t) = X(t\wedge\tau)$ is $\mathfrak{F}(t\wedge\tau) = \mathfrak{F}^\tau(t)$-measurable, and hence X^τ is $\mathfrak{F}^\tau$-adapted. Since $\mathfrak{F}^\tau(t)\subseteq\mathfrak{F}(t)$ for all $t\in\mathcal{T}$, this furthermore implies that X^τ is $\mathfrak{F}$-adapted.

Let $B\in\mathfrak{S}$ and observe that, for any $s,t\in\mathcal{T}$ with $s\le t$, we have

$$\{X(\tau)\in B\}\cap\{\tau=s\} = \{X(s)\in B\}\cap\{\tau=s\}\in\mathfrak{F}(s)\subseteq\mathfrak{F}(t)$$

by adaptedness of X and Lemma A.5 (Properties of the σ-Field of the τ-Past). But then, using that $\tau(\Omega)$ is countable, we find

$$\{X(\tau)\in B\}\cap\{\tau\le t\} = \bigcup_{s\in\tau(\Omega),\ s\le t}\Big(\{X(\tau)\in B\}\cap\{\tau=s\}\Big)\in\mathfrak{F}(t)$$

for all $t\in\mathcal{T}$, i.e. $\{X(\tau)\in B\}\in\mathfrak{F}(\tau)$ and thus $X(\tau)$ is $\mathfrak{F}(\tau)$-measurable. □

In continuous time, as you may have guessed, the question of adaptedness of the stopped process is more delicate and the answer is **in general not affirmative**. We shall prove the adaptedness of the stopped process in the case of right continuous processes (which is, once again, not the most general result but sufficient for our purposes). The proof is based on the very useful fact that any stopping time can be approximated from the right by a monotone sequence of stopping times with a finite range.

Proposition A.12 (Approximation of Stopping Times). *Assume that $\mathcal{T} = [0,\infty)$ and let τ be a stopping time with respect to a filtration $\mathfrak{F} = \{\mathfrak{F}(t)\}_{t\in[0,\infty)}$. Then there exists a sequence $\{\tau_n\}_{n\in\mathbb{N}}$ of stopping times such that* *(i)* $\tau_n(\Omega) \subset [0,\infty]$ *is finite for all* $n \in \mathbb{N}$*,*

(ii) $\tau \leq \tau_{n+1} \leq \tau_n$ *for all* $n \in \mathbb{N}$*,*

(iii) $\tau < \tau_n$ *for all* $n \in \mathbb{N}$ *on* $\{\tau < \infty\}$*,*

(iv) $\inf_{n\in\mathbb{N}} \tau_n = \lim_{n\to\infty} \tau_n = \tau$. ◇

Proof. Given $n \in \mathbb{N}$, we define $\tau_n : \Omega \to [0,\infty]$ by

$$\tau_n(\omega) \triangleq \sum_{k=1}^{n2^n} k2^{-n} \mathbb{1}_{\{\tau \in [(k-1)2^{-n}, k2^{-n})\}} + \infty \mathbb{1}_{\{\tau \geq n\}}, \qquad \omega \in \Omega.$$

Then the sequence $\{\tau_n\}_{n\in\mathbb{N}}$ clearly satisfies the properties (i) to (iv) above. Moreover, each τ_n is a stopping time by Lemma A.2 (Characterization of Discrete Stopping Times) since

$$\{\tau_n = k2^{-n}\} = \{\tau \geq (k-1)2^{-n}\} \cap \{\tau < k2^{-n}\} \in \mathfrak{F}(k2^{-n})$$

for all $k \in \{1, 2, \ldots, n2^n\}$. □

It is crucial that the sequence $\{\tau_n\}_{n\in\mathbb{N}}$ approximates the stopping time τ **from the right**. In general, it is not possible to find an approximating sequence from the left, since otherwise we would be able to anticipate the stopping time τ. With this approximating procedure at hand, we can now show by a limiting argument that if we stop adapted right continuous processes, the stopped process is still adapted.

Proposition A.13 (Measurability of Stopped Processes (Continuous Case))**.** *Let* $X = \{X(t)\}_{t\in[0,\infty)}$ *be a right continuous stochastic process which is adapted to a filtration* $\mathfrak{F} = \{\mathfrak{F}(t)\}_{t\in[0,\infty)}$*. Let moreover* τ *be a stopping time with respect to* $\mathfrak{F}$*. Then the stopped process* X^τ *is adapted to both the stopped filtration* $\mathfrak{F}^\tau$ *and the original filtration* $\mathfrak{F}$*. Moreover, if* τ *takes values in* $[0,\infty)$ *only, then* $X(\tau)$ *is* $\mathfrak{F}(\tau)$*-measurable.* ◇

Proof. Let $\{\tau_n\}_{n\in\mathbb{N}}$ be the approximating sequence of stopping times for τ given by Proposition A.12 (Approximation of Stopping Times). Then it follows from Lemma A.11 (Measurability of Stopped Processes (Discrete Case)) that X^{τ_n} is $\mathfrak{F}$-adapted for each $n \in \mathbb{N}$. From this and the right continuity of X, it follows that

$$X^\tau(t) = X(t \wedge \tau) = \lim_{n\to\infty} X(t \wedge \tau_n) = \lim_{n\to\infty} X^{\tau_n}(t)$$

is $\mathfrak{F}(t)$-measurable for all $t \in [0, \infty)$, i.e. X^τ is $\mathfrak{F}$-adapted. Now assume that τ takes values in $[0, \infty)$ and let $t \in [0, \infty)$. Then, since $\{\tau \leq t\} \in \mathfrak{F}(t)$ and $X^\tau(t)$ is $\mathfrak{F}(t)$-measurable, it follows that, for all $B \in \mathfrak{S}$, we have

$$\{X(\tau) \in B\} \cap \{\tau \leq t\} = \{X^\tau(t) = X(t \wedge \tau) \in B\} \cap \{\tau \leq t\} \in \mathfrak{F}(t),$$

and hence $X(\tau)$ is $\mathfrak{F}(\tau)$-measurable. As in the proof of Lemma A.11 (Measurability of Stopped Processes (Discrete Case)), this implies that $X^\tau(t) = X(t \wedge \tau)$ is $\mathfrak{F}^\tau(t) = \mathfrak{F}(t \wedge \tau)$-measurable for all $t \in [0, \infty)$ and hence X^τ is also adapted to the stopped filtration $\mathfrak{F}^\tau$. □

Appendix

B Martingale Theory

Throughout this section, we work on a general (not necessarily finite) probability space $(\Omega, \mathfrak{A}, \mathbb{P})$. Moreover, we fix a general time index set $\mathcal{T} \subseteq \mathbb{R}$ as well as a filtration $\mathfrak{F} = \{\mathfrak{F}(t)\}_{t\in\mathcal{T}}$.

B.1 Doob's Inequalities

The aim of this section is to establish a series of very **useful inequalities** involving sub-/supermartingales which are due to Joseph Doob. We begin with the so-called maximal inequalities.

Theorem B.1 (Doob's Maximal Inequalities (Discrete)). *Let $X = \{X(t)\}_{t\in\mathcal{T}}$ be a stochastic process, assume that $\mathcal{T}$ is finite, and fix $\alpha > 0$. If X is a submartingale, then*

$$\alpha\mathbb{P}\Big[\max_{t\in\mathcal{T}} X(t) > \alpha\Big] \le \mathbb{E}\Big[\mathbb{1}_{\{\max_{t\in\mathcal{T}} X(t)>\alpha\}} X\big(\max \mathcal{T}\big)\Big] \le \mathbb{E}\big[X\big(\max \mathcal{T}\big)^+\big].$$

If on the other hand X is a supermartingale, then

$$\alpha\mathbb{P}\Big[\max_{t\in\mathcal{T}} X(t) > \alpha\Big] \le \mathbb{E}\big[X\big(\min \mathcal{T}\big)\big] + \mathbb{E}\big[X\big(\max \mathcal{T}\big)^-\big]. \qquad \diamond$$

Proof. Define a $\mathcal{T}$-valued stopping time τ by

$$\tau \triangleq \min\big\{t \in \mathcal{T} : X(t) > \alpha\big\} \wedge \max \mathcal{T}.$$

By Lemma A.11 (Stopped Process with Finite Time Index Set), it follows that

$$F \triangleq \left\{\max_{t\in\mathcal{T}} X(t) > \alpha\right\} = \{X(\tau) > \alpha\} \in \mathfrak{F}(\tau).$$

From this we see in particular that $X(\tau) > \alpha$ on F and thus

$$\mathbb{E}\big[\mathbb{1}_F X(\tau)\big] \geq \alpha\mathbb{E}[\mathbb{1}_F] = \alpha\mathbb{P}[F].$$

But then, if X is a submartingale, it follows from Theorem 4.6 (Optional Stopping) that

$$\begin{aligned}\alpha\mathbb{P}[F] \leq \mathbb{E}\big[\mathbb{1}_F X(\tau)\big] \leq \mathbb{E}\Big[\mathbb{1}_F \mathbb{E}_\tau\big[X\big(\max\mathcal{T}\big)\big]\Big] \\ = \mathbb{E}\big[\mathbb{1}_F X(\max\mathcal{T})\big] \leq \mathbb{E}\big[X(\max\mathcal{T})^+\big]\end{aligned}$$

as asserted. If X is a supermartingale, we apply optional stopping twice to arrive at

$$\begin{aligned}\mathbb{E}\big[X\big(\min\mathcal{T}\big)\big] \geq \mathbb{E}\big[X(\tau)\big] &= \mathbb{E}\big[\mathbb{1}_F X(\tau) + \mathbb{1}_{F^c} X(\tau)\big] \\ &\geq \alpha\mathbb{P}[F] + \mathbb{E}\big[\mathbb{1}_{\{X(\tau)\leq\alpha\}} X(\tau)\big] \\ &\geq \alpha\mathbb{P}[F] + \mathbb{E}\big[\mathbb{1}_{\{X(\tau)\leq\alpha\}} \mathbb{E}_\tau\big[X(\max\mathcal{T})\big]\big] \\ &= \alpha\mathbb{P}[F] + \mathbb{E}\big[\mathbb{1}_{\{X(\tau)\leq\alpha\}} X(\max\mathcal{T})\big].\end{aligned}$$

Upon rearranging and further estimation, this implies that

$$\begin{aligned}\alpha\mathbb{P}[F] &\leq \mathbb{E}\big[X\big(\min\mathcal{T}\big)\big] - \mathbb{E}\big[\mathbb{1}_{\{X(\tau)\leq\alpha\}} X(\max\mathcal{T})\big] \\ &\leq \mathbb{E}\big[X\big(\min\mathcal{T}\big)\big] + \mathbb{E}\big[X(\max\mathcal{T})^-\big]\end{aligned}$$

and the proof is complete. □

Doob's maximal inequalities are remarkable because the right hand sides of the inequalities only depend on the process X at time $\min\mathcal{T}$ and $\max\mathcal{T}$, whereas the left hand side involves **the whole path** $X(\cdot)$. In particular, the cardinality and mesh size of $\mathcal{T}$ do not play any role in these bounds.

Theorem B.2 (Doob's L^p Inequality (Discrete))**.** *Let X be a martingale or a nonnegative submartingale and suppose that $\mathcal{T}$ is finite. If for some $p \in (1,\infty)$ we have $\mathbb{E}[|X(t)|^p] < \infty$ for all $t \in \mathcal{T}$, then*

$$\mathbb{E}\Big[\max_{t\in\mathcal{T}} |X(t)|^p\Big] < \infty$$

and

$$\mathbb{E}\Big[\max_{t\in\mathcal{T}} |X(t)|^p\Big]^{1/p} \leq \tfrac{p}{p-1}\mathbb{E}\Big[\big|X\big(\max\mathcal{T}\big)\big|^p\Big]^{1/p}. \qquad \diamond$$

Proof. If X is a martingale, $|X|$ is a submartingale by Jensen's inequality, and hence it suffices to prove the result for submartingales which are almost surely nonnegative. We define $Y \triangleq \max_{t\in\mathcal{T}} X(t)$ and observe that $Y \leq \sum_{t\in\mathcal{T}} X(t)$. The discrete version of Jensen's inequality thus shows that

$$\mathbb{E}\big[Y^p\big] \leq \mathbb{E}\Big[N^{p-1}\sum_{t\in\mathcal{T}} X(t)^p\Big] = N^{p-1}\sum_{t\in\mathcal{T}}\mathbb{E}\big[X(t)^p\big] < \infty,$$

where N denotes the number of elements in $\mathcal{T}$. Now use $x^p = \int_0^x p\alpha^{p-1}\mathrm{d}\alpha$ and apply Tonelli's theorem to rewrite

$$\begin{aligned}\mathbb{E}\big[Y^p\big] &= p\mathbb{E}\Big[\int_0^Y \alpha^{p-1}\mathrm{d}\alpha\Big]\\ &= p\mathbb{E}\Big[\int_0^\infty \alpha^{p-1}\mathbb{1}_{\{Y>\alpha\}}\mathrm{d}\alpha\Big] = p\int_0^\infty \alpha^{p-1}\mathbb{P}[Y>\alpha]\mathrm{d}\alpha.\end{aligned}$$

Now Doob's maximal inequality for submartingales allows us to estimate this by

$$\mathbb{E}\big[Y^p\big] = p\int_0^\infty \alpha^{p-1}\mathbb{P}[Y>\alpha]\mathrm{d}\alpha \leq p\int_0^\infty \alpha^{p-2}\mathbb{E}\big[\mathbb{1}_{\{Y>\alpha\}}X\big(\max\mathcal{T}\big)\big]\mathrm{d}\alpha.$$

Using Tonelli's theorem once more therefore yields

$$\begin{aligned}\mathbb{E}\big[Y^p\big] &\leq p\int_0^\infty \alpha^{p-2}\mathbb{E}\big[\mathbb{1}_{\{Y>\alpha\}}X\big(\max\mathcal{T}\big)\big]\mathrm{d}\alpha\\ &= p\mathbb{E}\Big[\int_0^\infty \alpha^{p-2}\mathbb{1}_{\{Y>\alpha\}}X\big(\max\mathcal{T}\big)\mathrm{d}\alpha\Big]\\ &= p\mathbb{E}\Big[X\big(\max\mathcal{T}\big)\int_0^Y \alpha^{p-2}\mathrm{d}\alpha\Big] = \frac{p}{p-1}\mathbb{E}\big[Y^{p-1}X\big(\max\mathcal{T}\big)\big].\end{aligned}$$

Setting $q \triangleq p/(p-1)$ so that $1/p + 1/q = 1$ and using Hölder's inequality, we can estimate the last term further and arrive at

$$\begin{aligned}\mathbb{E}\big[Y^p\big] \leq \frac{p}{p-1}\mathbb{E}\big[Y^{p-1}X\big(\max\mathcal{T}\big)\big] &\leq \frac{p}{p-1}\mathbb{E}\big[Y^{q(p-1)}\big]^{1/q}\mathbb{E}\big[X\big(\max\mathcal{T}\big)^p\big]^{1/p}\\ &= \frac{p}{p-1}\mathbb{E}\big[Y^p\big]^{1-1/p}\mathbb{E}\big[X\big(\max\mathcal{T}\big)^p\big]^{1/p}.\end{aligned}$$

Dividing both sides by $\mathbb{E}\big[Y^p\big]^{1-1/p}$ finishes the proof. $\square$

While the discrete version of Doob's L^p inequality is useful, in the main part of these notes we need a version for **continuous time index sets**, which can be obtained by an approximation argument.

Theorem B.3 (Doob's L^p Inequality)**.** *Let $X = \{X(t)\}_{t\in[0,T]}$ be a right continuous martingale or a nonnegative and right continuous submartingale. If for some $p \in (1,\infty)$ we have $\mathbb{E}[|X(t)|^p] < \infty$ for all $t \in [0,T]$, then*

$$\mathbb{E}\Big[\sup_{t\in[0,T]} |X(t)|^p\Big] < \infty$$

and

$$\mathbb{E}\Big[\sup_{t\in[0,T]} |X(t)|^p\Big]^{1/p} \leq \tfrac{p}{p-1}\mathbb{E}\Big[\big|X(T)\big|^p\Big]^{1/p}. \qquad \diamond$$

Proof. Since X is right continuous, we have

$$\sup_{t\in[0,T]} |X(t)| = \sup_{q\in([0,T)\cap\mathbb{Q})\cup\{T\}} |X(q)| = \sup_{q\in\mathcal{F}\cup\{T\}\subset([0,T)\cap\mathbb{Q})\cup\{T\},\ \mathcal{F}\text{ finite}} |X(q)|.$$

By monotone convergence we therefore have

$$\mathbb{E}\Big[\sup_{t\in[0,T]} |X(t)|^p\Big]^{1/p} = \sup_{\mathcal{F}\subset[0,T)\cap\mathbb{Q}\text{ finite}} \mathbb{E}\Big[\max_{t\in\mathcal{F}\cup\{T\}} |X(t)|^p\Big]^{1/p}.$$

Now Theorem B.2 (Doob's L^p-Inequality in Discrete Time) shows that

$$\mathbb{E}\Big[\max_{t\in\mathcal{F}\cup\{T\}} |X(t)|^p\Big]^{1/p} \leq \frac{p}{p-1}\mathbb{E}\Big[\big|X(T)\big|^p\Big]^{1/p}$$

for any finite subset $\mathcal{F}$ of $[0,T)\cap\mathbb{Q}$. Together with the previous identity this concludes the proof. □

B.2 Continuous Martingales of Finite Variation

In this section, we establish the important fact that the only continuous martingales of finite variation are the continuous ones. Effectively, this result is a consequence of the fact that martingale increments are **orthogonal**.

Proposition B.4 (Orthogonality of Martingale Increments). *Let us assume that $M = \{M(t)\}_{t\in[0,T]}$ is a martingale with $\mathbb{E}[|M(t)|^2] < \infty$ for all $t \in [0,T]$.*

(i) For all $s, t \in [0,T]$ with $s < t$, it holds that

$$\mathbb{E}_s\big[|M(t) - M(s)|^2\big] = \mathbb{E}_s\big[|M(t)|^2 - |M(s)|^2\big].$$

(ii) For all $0 \le t_0 < t_1 < \ldots < t_n \le T$, $n \in \mathbb{N}$, we have

$$\mathbb{E}\big[|M(t_n)|^2\big] = \mathbb{E}\big[|M(t_0)|^2\big] + \sum_{k=1}^{n} \mathbb{E}\big[|M(t_k) - M(t_{k-1})|^2\big]. \qquad \diamond$$

Proof. Step 1: We prove (i). For this, let $s, t \in [0,T]$ with $s < t$. By expanding the square and using the martingale property of M, we obtain

$$\begin{aligned}
\mathbb{E}_s\big[|M(t) - M(s)|^2\big] &= \mathbb{E}_s\big[|M(t)|^2\big] - \mathbb{E}_s\big[2M(s)M(t)\big] + \mathbb{E}_s\big[|M(s)|^2\big] \\
&= \mathbb{E}_s\big[|M(t)|^2\big] - 2M(s)\mathbb{E}_s\big[M(t)\big] + |M(s)|^2 \\
&= \mathbb{E}_s\big[|M(t)|^2\big] - 2|M(s)|^2 + |M(s)|^2 \\
&= \mathbb{E}_s\big[|M(t)|^2 - |M(s)|^2\big].
\end{aligned}$$

Step 2: We prove (ii). Let $n \in \mathbb{N}$ and $0 \le t_0 < t_1 < \ldots < t_n \le T$. Then it follows from (i) that

$$\begin{aligned}
\mathbb{E}\big[|M(t_n)|^2\big] &= \mathbb{E}\big[|M(t_0)|^2\big] + \sum_{k=1}^{n}\Big(\mathbb{E}\big[|M(t_k)|^2\big] - \mathbb{E}\big[|M(t_{k-1})|^2\big]\Big) \\
&= \mathbb{E}\big[|M(t_0)|^2\big] + \sum_{k=1}^{n} \mathbb{E}\big[|M(t_k) - M(t_{k-1})|^2\big]. \qquad \square
\end{aligned}$$

With the orthogonality of martingale increments at hand, we can now proceed to show that continuous martingales of **finite variation** are almost surely constant. For this, we denote by $V_X = \{V_X(t)\}_{t\in[0,T]}$ the total variation process of a stochastic process $X = \{X(t)\}_{t\in[0,T]}$ defined by

$$V_X(t,\omega) \triangleq V_{X(\cdot,\omega)}(t), \qquad (t,\omega) \in [0,T] \times \Omega,$$

where we define the total variation $V_f : [0,T] \to [0,\infty]$ of a deterministic function $f : [0,T] \to \mathbb{R}$ by

$$V_f(t) \triangleq \sup\Big\{\sum_{k=1}^{n} |f(t_{k-1}) - f(t_k)| : 0 \le t_0 < \ldots < t_n \le t, n \in \mathbb{N}\Big\}, \quad t \in [0,T].$$

We then say that X is of **finite variation** if $V_X(T) < \infty$.

Theorem B.5 (Continuous Martingales of Finite Variation). *Assume that $M = \{M(t)\}_{t\in[0,T]}$ is a continuous martingale of finite variation. Then*

$$M(t) = M(0), \qquad t \in [0,T]. \qquad \diamond$$

Proof. Step 1: Localization. We can without loss of generality assume that $M(0) = 0$ (consider the process $M - M(0)$ otherwise). We show that we may without loss of generality assume that both M and V_M are uniformly bounded. Indeed, for each $k \in \mathbb{N}$ we can define a stopping time

$$\tau_k \triangleq \inf\big\{t \in [0,T] : |M(t)| \ge k \text{ or } V_M(t) \ge k\big\}.$$

By continuity, the stopped processes $M(\cdot \wedge \tau_k)$ and $V_M(\cdot \wedge \tau_k) = V_{M(\cdot\wedge\tau_k)}$ are bounded by k. If the result now holds for $M(\cdot \wedge \tau_k)$ for all $k \in \mathbb{N}$, then it must also hold for M since $\tau_k \to \infty$.

Step 2: Orthogonality of martingale increments. Let us assume that both $|M|$ and V_M are bounded by some $\alpha > 0$ and fix $t \in (0,T]$. For an arbitrary partition $0 = t_0 < t_1 < \ldots < t_n = t$ of $[0,t]$, Proposition B.4 (Othogonality of Martingale Increments) shows that

$$\begin{aligned}
\mathbb{E}\big[|M(t)|^2\big] &= \mathbb{E}\Big[\sum_{k=1}^{n} |M(t_k) - M(t_{k-1})|^2\Big] \\
&\le \mathbb{E}\Big[\max_{k=1,\ldots,n} |M(t_k) - M(t_{k-1})| \sum_{k=1}^{n} |M(t_k) - M(t_{k-1})|\Big] \\
&\le \mathbb{E}\Big[\max_{k=1,\ldots,n} |M(t_k) - M(t_{k-1})| V_M(T)\Big] \\
&\le \alpha \mathbb{E}\Big[\max_{k=1,\ldots,n} |M(t_k) - M(t_{k-1})|\Big].
\end{aligned}$$

Since M is continuous, it is uniformly continuous on $[0,t]$ and we therefore find that $|M(t_k) - M(t_{k-1})| \to 0$ as $\max_{k=1,\dots,n} |t_k - t_{k-1}| \to 0$. Since M is uniformly bounded it follows moreover that $|M(t_k) - M(t_{k-1})| \leq 2\alpha$ for all $k = 1, \dots, n$. The dominated convergence theorem hence yields

$$\alpha \mathbb{E}\Big[\max_{k=1,\dots,n} |M(t_k) - M(t_{k-1})|\Big] \to 0 \quad \text{as} \max_{k=1,\dots,n} |t_k - t_{k-1}| \to 0.$$

But this implies that $\mathbb{E}[|M(t)|^2] = 0$, which is only possible if $M(t) = 0 = M(0)$ for all $t \in [0,T]$. □

Index

Δ-Hedge, 135
σ-Field of the τ-Past, 140

Adapted, 21
Arbitrage Opportunity, 8
  in Finite Markets, 29
Attainable, 39, 129

Bellman Principle
  Optimal Stopping, 65
Binomial Model, 18
Black-Scholes Model, 118
Bond, 3
  Coupon, 3
  Zero-Coupon, 3
Brownian Motion, 95

CRR Model, 17

Derivative, 4
  Underlying, 4
Doob Decomposition, 69
Doob's L^p Inequality
  Continuous Time, 154
  Discrete Time, 152
Doob's Maximal Inequalities
  Discrete Time, 151
Dynamic Programming Principle,
  see Bellman Principle, 65

Efficient Strategy, 128
EMM, 33
Equivalent Martingale Measure,
  see EMM, 33

Filtered Probability Space, 20
Filtration, 20
  Augmented, 96
  Natural, 21
  Right Continuous, 143
Financial Market, 1
  Complete, 43
  Finite, 21
  Perfect, 2
Financial Securities, 2
FTAP
  First, 34
  Second, 44
Fundamental Theorem of Asset Pricing,
  see FTAP, 34

Hitting Time, 142

Implied Volatility, 132
Integrand
- W-Integrand, 108
- X-Integrand, 111
- L^2-Integrand, 102
- Simple Integrand, 101

Itō Formula, 113
- Product Rule, 117

Itō Integral, 109
- L^2-Integrand, 106
- for Itō Processes, 111
- Simple Integrand, 101

Itō Isometry, 104
Itō Process, 110
- Characteristic Pair, 110

Itō's Product Rule, 117

Law of One Price, 38
Lower Price, 47

Market Friction, 2
Market Model
- Binomial, 17
- CRR, 17
- Finite, 21
- Trinomial, 19

Martingale, 97
- Finite, 32
- Orthogonality of Increments, 155
- Square Integrable, 97

Numéraire, 19

Optimal Stopping, 64
- Bellman Principle, 65
- Optimal Stopping Time, 67
- Value Process, 64

Optimal Stopping Theorem, 67, 75
Option, 4
- American, 63
- American Call, 7
- American Put, 7
- Attainable, 39, 129
- Bermudan, 7
- Black-Scholes Market, 129
- Canary, 7
- Digital Call, 12
- Digital Put, 12
- European Call, 4
- European Put, 4
- Exotic, 8
- Finite Market, 39
- Underlying, 4
- Vanilla, 8

Optional Decomposition, 51
Optional Sampling, 72
- for Martingales, 73

Optional Stopping, 71

Predictable, 100
- Finite Time Index Set, 23

Predictable σ-Field, 100
Put-Call Parity, 11

Quadratic Covariation, 112
Quadratic Variation, 112

Replication Strategy, 39, 129
Risk Neutral Pricing
- Finite Markets, 41

Self-Financing, 23
Short Sale, 2
Snell Envelope, 66
Stochastic Integral
- Discrete, 24

Stochastic Process, 16
- Adapted, 21

Predictable, 23, 100
Stopped Process, 146
Stopping Time, 138
Stopped Process, 146
Submartingale
Finite, 32
Superhedging Duality, 58
Superhedging Price
American, 79
European, 40
Superhedging Strategy
American, 79
European, 40
Supermartingale
Finite, 32

Trading Strategy
Δ-Hedge, 135
Efficient Strategy, 128
Replication Strategy, 39, 129
Self-Financing, 23
Superhedging (American), 79
Superhedging (European), 40
with Initial Wealth, 28
Trinomial Model, 19

Underlying, 4
Upper Price, 47

Volatility Smile, 132

Wealth Process, 23
Wiener Process,
see Brownian Motion, 95

Vocabulary

English	German
σ-field	σ-Algebra
bond	Anleihe
call option	Kaufoption
commodity	Rohstoff
demand	Nachfrage
derivative	Derivat
to discount	abzinsen
face value	Nominalwert
filtration	Filtrierung
fixed-income products	Zinsprodukte
to hedge	absichern
interest rate	Zinsrate
loan	Darlehen
maturity	Fälligkeit
predictable	vorhersagbar
put option	Verkaufoption
security	Wertpapier

English	German
short sale	Leerverkauf
Snell envelope	Snellsche Einhüllende
stock	Aktie
stopping time	Stoppzeit
strike price	Ausübungspreis
supply	Angebot
underlying	Basiswert
valuation	Bewertung
wealth	Vermögen
zero-coupon bond	Nullkuponanleihe