2

Models of Financial Markets on Finite Probability Spaces

2.1 Description of the Model

In this section we shall develop the theory of pricing and hedging of derivative securities in financial markets.

In order to reduce the technical difficulties of the theory of option pricing to a minimum, we assume throughout this chapter that the probability space Ω underlying our model will be finite, say, $\Omega = \{\omega_1, \omega_2, \ldots, \omega_N\}$ equipped with a probability measure $\mathbf{P}$ such that $\mathbf{P}[\omega_n] = p_n > 0$, for $n = 1, \ldots, N$. This assumption implies that all functional-analytic delicacies pertaining to different topologies on $L^\infty(\Omega, \mathcal{F}, \mathbf{P})$, $L^1(\Omega, \mathcal{F}, \mathbf{P})$, $L^0(\Omega, \mathcal{F}, \mathbf{P})$ etc. evaporate, as all these spaces are simply $\mathbb{R}^N$ (we assume w.l.o.g. that the σ-algebra $\mathcal{F}$ is the power set of Ω). Hence all the functional analysis, which we shall need in later chapters for the case of more general processes, reduces in the setting of the present chapter to simple linear algebra. For example, the use of the Hahn-Banach theorem is replaced by the use of the separating hyperplane theorem in finite dimensional spaces.

Nevertheless we shall write $L^\infty(\Omega, \mathcal{F}, \mathbf{P})$, $L^1(\Omega, \mathcal{F}, \mathbf{P})$ etc. (knowing very well that in the present setting these spaces are all isomorphic to $\mathbb{R}^N$) to indicate, which function spaces we shall encounter in the setting of the general theory. It also helps to see if an element of $\mathbb{R}^N$ is a contingent claim or an element of the dual space, i.e. a price vector.

In addition to the probability space $(\Omega, \mathcal{F}, \mathbf{P})$ we fix a natural number $T \geq 1$ and a filtration $(\mathcal{F}_t)_{t=0}^T$ on Ω, i.e., an increasing sequence of σ-algebras. To avoid trivialities, we shall always assume that $\mathcal{F}_T = \mathcal{F}$; on the other hand, we shall *not* assume that $\mathcal{F}_0$ is trivial, i.e. $\mathcal{F}_0 = \{\emptyset, \Omega\}$, although this will be the case in most applications. But for technical reasons it will be more convenient to allow for general σ-algebras $\mathcal{F}_0$.

We now introduce a model of a financial market in not necessarily discounted terms. The rest of Sect. 2.1 will be devoted to reducing this situation to a model in discounted terms which, as we shall see, will make life much easier.

Readers who are not so enthusiastic about this mainly formal and elementary reduction might proceed directly to Definition 2.1.4. On the other hand, we know from sad experience that often there is a lot of myth and confusion arising in this operation of discounting; for this reason we decided to devote this section to the clarification of this issue.

Definition 2.1.1. *A model of a* financial market *is an* $\mathbb{R}^{d+1}$*-valued stochastic process* $\widehat{S} = (\widehat{S}_t)_{t=0}^T = (\widehat{S}_t^0, \widehat{S}_t^1, \ldots, \widehat{S}_t^d)_{t=0}^T$*, based on and adapted to the filtered stochastic base* $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t=0}^T, \mathbf{P})$*. We shall assume that the zero coordinate* $\widehat{S}^0$ *satisfies* $\widehat{S}_t^0 > 0$ *for all* $t = 0, \ldots, T$ *and* $\widehat{S}_0^0 = 1$*.*

The interpretation is the following. The prices of the assets $0, \ldots, d$ are measured in a fixed money unit, say Euros. For $1 \leq j \leq d$ they are not necessarily non-negative (think, e.g., of forward contracts). The asset 0 plays a special role. It is supposed to be strictly positive and will be used as a numéraire. It allows us to compare money (e.g., Euros) at time 0 to money at time $t > 0$. In many elementary models, $\widehat{S}^0$ is simply a bank account which in case of constant interest rate r is then defined as $\widehat{S}_t^0 = e^{rt}$. However, it might also be more complicated, e.g. $\widehat{S}_t^0 = \exp(r_0 h + r_1 h + \cdots + r_{t-1} h)$ where $h > 0$ is the length of the time interval between $t-1$ and t (here kept fixed) and where r_{t-1} is the stochastic interest rate valid between $t-1$ and t. Other models are also possible and to prepare the reader for more general situations, we only require $\widehat{S}_t^0$ to be strictly positive. Notice that we only require that $\widehat{S}_t^0$ to be $\mathcal{F}_t$-measurable and that it is not necessarily $\mathcal{F}_{t-1}$-measurable. In other words, we assume that the process $\widehat{S}^0 = (\widehat{S}_t^0)_{t=0}^T$ is adapted, but not necessarily predictable.

An economic agent is able to buy and sell financial assets. The decision taken at time t can only use information available at time t which is modelled by the σ-algebra $\mathcal{F}_t$.

Definition 2.1.2. *A trading strategy* $(\widehat{H}_t)_{t=1}^T = (\widehat{H}_t^0, \widehat{H}_t^1, \ldots, \widehat{H}_t^d)_{t=1}^T$ *is an* $\mathbb{R}^{d+1}$*-valued process which is predictable, i.e.* $\widehat{H}_t$ *is* $\mathcal{F}_{t-1}$*-measurable.*

The interpretation is that between time $t-1$ and time t, the agent holds a quantity equal to $\widehat{H}_t^j$ of asset j. The decision is taken at time $t-1$ and therefore, $\widehat{H}_t$ is required to be $\mathcal{F}_{t-1}$-measurable.

Definition 2.1.3. *A strategy* $(\widehat{H}_t)_{t=1}^T$ *is called self financing if for every* $t = 1, \ldots, T-1$*, we have*

$$\left(\widehat{H}_t, \widehat{S}_t\right) = \left(\widehat{H}_{t+1}, \widehat{S}_t\right) \tag{2.1}$$

or, written more explicitly,

$$\sum_{j=0}^{d} \widehat{H}_t^j \widehat{S}_t^j = \sum_{j=0}^{d} \widehat{H}_{t+1}^j \widehat{S}_t^j. \tag{2.2}$$

The initial investment required for a strategy is $\widehat{V}_0 = (\widehat{H}_1, \widehat{S}_0) = \sum_{j=0}^d \widehat{H}_1^j \widehat{S}_0^j$*.*

The interpretation goes as follows. By changing the portfolio from $\widehat{H}_{t-1}$ to $\widehat{H}_t$ there is no input/outflow of money. We remark that we assume that changing a portfolio does not trigger transaction costs. Also note that $\widehat{H}_t^j$ may assume negative values, which corresponds to short selling asset j during the time interval $]t_{j-1}, t_j]$.

The $\mathcal{F}_t$-measurable random variable defined in (2.1) is interpreted as the value $\widehat{V}_t$ of the portfolio at time t defined by the trading strategy $\widehat{H}$:

$$\widehat{V}_t = (\widehat{H}_t, \widehat{S}_t) = (\widehat{H}_{t+1}, \widehat{S}_t).$$

The way in which the value $(\widehat{H}_t, \widehat{S}_t)$ evolves can be described much easier when we use discounted prices using the asset $\widehat{S}^0$ as numéraire. Discounting allows us to compare money at time t to money at time 0. For instance we could say that $\widehat{S}_t^0$ units of money at time t are the "same" as 1 unit of money, e.g., Euros, at time 0. So let us see what happens if we replace prices $\widehat{S}$ by discounted prices $\left(\frac{\widehat{S}}{\widehat{S}^0}\right) = \left(\frac{\widehat{S}^0}{\widehat{S}^0}, \frac{\widehat{S}^1}{\widehat{S}^0}, \ldots, \frac{\widehat{S}^d}{\widehat{S}^0}\right)$. We will use the notation

$$S_t^j := \frac{\widehat{S}_t^j}{\widehat{S}_t^0}, \quad \text{for } j = 1, \ldots, d \text{ and } t = 0, \ldots, T. \tag{2.3}$$

There is no need to include the coordinate 0, since obviously $S_t^0 = 1$. Let us now consider $(\widehat{H}_t)_{t=1}^T = (\widehat{H}_t^0, \widehat{H}_t^1, \ldots, \widehat{H}_t^d)_{t=1}^T$ to be a self financing strategy with initial investment $\widehat{V}_0$; we then have

$$\widehat{V}_0 = \sum_{j=0}^d \widehat{H}_1^j \widehat{S}_0^j = \widehat{H}_1^0 + \sum_{j=1}^d \widehat{H}_1^j \widehat{S}_0^j = \widehat{H}_1^0 + \sum_{j=1}^d \widehat{H}_1^j S_0^j,$$

since by definition $\widehat{S}_0^0 = 1$.

We now write $(H_t)_{t=1}^T = (H_t^1, \ldots, H_t^d)_{t=1}^T$ for the $\mathbb{R}^d$-valued process obtained by discarding the 0'th coordinate of the $\mathbb{R}^{d+1}$-valued process $(\widehat{H}_t)_{t=1}^T = (\widehat{H}_t^0, \widehat{H}_t^1, \ldots, \widehat{H}_t^d)_{t=1}^T$, i.e., $H_t^j = \widehat{H}_t^j$ for $j = 1, \ldots, d$. The reason for dropping the 0'th coordinate is, as we shall discover in a moment, that the holdings $\widehat{H}_t^0$ in the numéraire asset S_t^0 will be no longer of importance when we do the book-keeping in terms of the numéraire asset, i.e., in discounted terms.

One can make the following easy, but crucial observation: for *every* $\mathbb{R}^d$-valued, predictable process $(H_t)_{t=1}^T = (H_t^1, \ldots, H_t^d)_{t=1}^T$ there exists a unique *self financing* $\mathbb{R}^{d+1}$-valued predictable process $(\widehat{H}_t)_{t=1}^T = (\widehat{H}_t^0, \widehat{H}_t^1, \ldots, \widehat{H}_t^d)_{t=1}^T$ such that $(\widehat{H}_t^j)_{t=1}^T = (H_t^j)_{t=1}^T$ for $j = 1, \ldots, d$ and $\widehat{H}_1^0 = 0$. Indeed, one determines the values of $\widehat{H}_{t+1}^0$, for $t = 1, \ldots, T-1$, by inductively applying (2.2). The strict positivity of $(\widehat{S}_t^0)_{t=0}^{T-1}$ implies that there is precisely one function $\widehat{H}_{t+1}^0$ such that equality (2.2) holds true. Clearly such a function $\widehat{H}_{t+1}^0$ is

$\mathcal{F}_t$-measurable. In economic terms the above argument is rather obvious: for any given trading strategy $(H_t)_{t=1}^T = (H_t^1, \ldots, H_t^d)_{t=1}^T$ in the "risky" assets $j = 1, \ldots, d$, we may always add a trading strategy $(\widehat{H}_t^0)_{t=1}^T$ in the numéraire asset 0 such that the total strategy becomes self financing. Moreover, by normalising $\widehat{H}_1^0 = 0$, this trading strategy becomes unique. This can be particularly well visualised when interpreting the asset 0 as a cash account, into which at all times $t = 1, \ldots, T-1$, the gains and losses occurring from the investments in the d risky assets are absorbed and from which the investments in the risky assets are financed. If we normalise this procedure by requiring $\widehat{H}_1^0 = 0$, i.e., by starting with an empty cash account, then clearly the subsequent evolution of the holdings in the cash account is uniquely determined by the holdings in the "risky" assets $1, \ldots, d$. From now on we fix two processes $(\widehat{H}_t)_{t=1}^T = (\widehat{H}_t^0, \widehat{H}_t^1, \ldots, \widehat{H}_t^d)_{t=1}^T$ and $(H_t)_{t=1}^T = (H_t^1, \ldots, H_t^d)_{t=1}^T$ corresponding uniquely one to each other in the above described way.

Now one can make a second straightforward observation: the investment $(\widehat{H}_t^0)_{t=1}^T$ in the numéraire asset does not change the *discounted* value $(V_t)_{t=0}^T$ of the portfolio. Indeed, by definition — and rather trivially — the numéraire asset remains constant in discounted terms (i.e., expressed in units of itself).

Hence the discounted value V_t of the portfolio

$$V_t = \frac{\widehat{V}_t}{\widehat{S}_t^0}, \quad t = 0, \ldots, T,$$

depends only on the $\mathbb{R}^d$-dimensional process $(H_t)_{t=1}^T = (H_t^1, \ldots, H_t^d)_{t=1}^T$.

More precisely, in view of the normalisation $\widehat{S}_0^0 = 1$ and $\widehat{H}_1^0 = 0$ we have

$$\widehat{V}_0 = V_0 = \sum_{j=1}^{d} H_1^j S_0^j.$$

For the increment $\Delta V_{t+1} = V_{t+1} - V_t$ we find, using (2.2),

$$\begin{aligned}
\Delta V_{t+1} = V_{t+1} - V_t &= \frac{\widehat{V}_{t+1}}{\widehat{S}_{t+1}^0} - \frac{\widehat{V}_t}{\widehat{S}_t^0} \\
&= \sum_{j=0}^{d} \widehat{H}_{t+1}^j \frac{\widehat{S}_{t+1}^j}{\widehat{S}_{t+1}^0} - \sum_{j=0}^{d} \widehat{H}_{t+1}^j \frac{\widehat{S}_t^j}{\widehat{S}_t^0} \\
&= \widehat{H}_{t+1}^0 (1-1) + \sum_{j=1}^{d} \widehat{H}_{t+1}^j \left(S_{t+1}^j - S_t^j \right) \\
&= \left(H_{t+1}^j, \Delta S_{t+1}^j \right),
\end{aligned}$$

where $(. , .)$ now denotes the inner product in $\mathbb{R}^d$.

In particular, the final value V_T of the portfolio becomes (in discounted units)

$$V_T = V_0 + \sum_{t=1}^{T} (H_t, \Delta S_t) = V_0 + (H \cdot S)_T,$$

where $(H \cdot S)_T = \sum_{t=1}^{T} (H_t, \Delta S_t)$ is the notation for a stochastic integral familiar from the theory of stochastic integration. In our discrete time framework the "stochastic integral" is simply a finite Riemann sum.

In order to know the value V_T of the portfolio in real money, we still would have to multiply by $\widehat{S}_T^0$, i.e., we have $\widehat{V}_T = V_T \widehat{S}_T^0$. This, however, is rarely needed.

We can therefore replace Definition 2.1.2 by the following definition in discounted terms, which will turn out to be much easier to handle.

Definition 2.1.4. *Let $S = (S^1, \ldots, S^d)$ be a model of a financial market in discounted terms. A* trading strategy *is an $\mathbb{R}^d$-valued process $(H_t)_{t=1}^T = (H_t^1, H_t^2, \ldots, H_t^d)_{t=1}^T$ which is predictable, i.e., each H_t is $\mathcal{F}_{t-1}$-measurable. We denote by $\mathcal{H}$ the set of all such trading strategies.*

We then define the stochastic integral $H \cdot S$ as the $\mathbb{R}$-valued process $((H \cdot S)_t)_{t=0}^T$ given by

$$(H \cdot S)_t = \sum_{u=1}^{t} (H_u, \Delta S_u), \quad t = 0, \ldots, T, \tag{2.4}$$

where $(. , .)$ denotes the inner product in $\mathbb{R}^d$. The random variable

$$(H \cdot S)_t = \sum_{u=1}^{t} (H_u, \Delta S_u)$$

models — when following the trading strategy H — the gain or loss occurred up to time t in discounted terms.

Summing up: by following the good old actuarial tradition of discounting, i.e. by passing from the process $\widehat{S}$, denoted in units of money, to the process S, denoted in terms of the numéraire asset (e.g., the cash account), things become considerably simpler and more transparent. In particular the value process V of an agent starting with initial wealth $V_0 = 0$ and subsequently applying the trading strategy H, is given by the stochastic integral $V_t = (H \cdot S)_t$ defined in (2.4).

We still emphasize that the choice of the numéraire is not unique; only for notational convenience we have fixed it to be the asset indexed by 0. But it may be chosen as any traded asset, provided only that it always remains strictly positive. We shall deal with this topic in more detail in Sect. 2.5 below.

From now on we shall work in terms of the discounted $\mathbb{R}^d$-valued process, denoted by S.

2.2 No-Arbitrage and the Fundamental Theorem of Asset Pricing

Definition 2.2.1. *We call the subspace K of $L^0(\Omega, \mathcal{F}, \mathbf{P})$ defined by*

$$K = \{(H \cdot S)_T \mid H \in \mathcal{H}\},$$

the set of contingent claims attainable at price 0, *where $\mathcal{H}$ denotes the set of predictable, $\mathbb{R}^d$-valued processes $H = (H_t)_{t=1}^T$.*

We leave it to the reader to check that K is indeed a vector space.

The economic interpretation is the following: the random variables $f = (H \cdot S)_T$ are precisely those contingent claims, i.e., the pay-off functions at time T, depending on $\omega \in \Omega$, that an economic agent may replicate with zero initial investment by pursuing some predictable trading strategy H.

For $a \in \mathbb{R}$, we call the *set of contingent claims attainable at price a* the affine space $K_a = a + K$, obtained by shifting K by the constant function a, in other words, the space of all the random variables of the form $a + (H \cdot S)_T$, for some trading strategy H. Again the economic interpretation is that these are precisely the contingent claims that an economic agent may replicate with an initial investment of a by pursuing some predictable trading strategy H.

Definition 2.2.2. *We call the convex cone C in $L^\infty(\Omega, \mathcal{F}, \mathbf{P})$ defined by*

$$C = \{g \in L^\infty(\Omega, \mathcal{F}, \mathbf{P}) \mid \textit{there exists } f \in K \textit{ with } f \geq g\}.$$

the set of contingent claims super-replicable at price 0.

Economically speaking, a contingent claim $g \in L^\infty(\Omega, \mathcal{F}, \mathbf{P})$ is *super-replicable at price* 0, if we can achieve it with zero net investment by pursuing some predictable trading strategy H. Thus we arrive at some contingent claim f and if necessary we "throw away money" to arrive at g. This operation of "throwing away money" or "free disposal" may seem awkward at this stage, but we shall see later that the set C plays an important role in the development of the theory. Observe that C is a convex cone containing the negative orthant $L^\infty_-(\Omega, \mathcal{F}, \mathbf{P})$. Again we may define $C_a = a + C$ as the *contingent claims super-replicable at price a*, if we shift C by the constant function a.

Definition 2.2.3. *A financial market S satisfies the* no-arbitrage condition *(NA) if*

$$K \cap L^0_+(\Omega, \mathcal{F}, \mathbf{P}) = \{0\}$$

or, equivalently,

$$C \cap L^0_+(\Omega, \mathcal{F}, \mathbf{P}) = \{0\}$$

where 0 *denotes the function identically equal to zero.*

Recall that $L^0(\Omega, \mathcal{F}, \mathbf{P})$ denotes the space of all $\mathcal{F}$-measurable real-valued functions and $L^0_+(\Omega, \mathcal{F}, \mathbf{P})$ its positive orthant.

We now have formalised the concept of an arbitrage possibility: it means the existence of a trading strategy H such that — starting from an initial investment zero — the resulting contingent claim $f = (H \cdot S)_T$ is non-negative and not identically equal to zero. Such an opportunity is of course the dream of every arbitrageur. If a financial market does not allow for arbitrage opportunities, we say it satisfies the *no-arbitrage condition (NA).*

Proposition 2.2.4. *Assume S satisfies (NA) then*

$$C \cap (-C) = K.$$

Proof. Let $g \in C \cap (-C)$ then $g = f_1 - h_1$ with $f_1 \in K$, $h_1 \in L^\infty_+$ and $g = f_2 + h_2$ with $f_2 \in K$ and $h_2 \in L^\infty_+$. Then $f_1 - f_2 = h_1 + h_2 \in L^\infty_+$ and hence $f_1 - f_2 \in K \cap L^\infty_+ = \{0\}$. It follows that $f_1 = f_2$ and $h_1 + h_2 = 0$, hence $h_1 = h_2 = 0$. This means that $g = f_1 = f_2 \in K$. □

Definition 2.2.5. *A probability measure* $\mathbf{Q}$ *on* $(\Omega, \mathcal{F})$ *is called an* equivalent martingale measure *for S, if $\mathbf{Q} \sim \mathbf{P}$ and S is a martingale under $\mathbf{Q}$, i.e., $\mathbf{E}_\mathbf{Q}[S_{t+1}|\mathcal{F}_t] = S_t$ for $t = 0, \ldots, T-1$.*

We denote by $\mathcal{M}^e(S)$ the set of equivalent martingale measures and by $\mathcal{M}^a(S)$ the set of all (not necessarily equivalent) martingale probability measures. The letter a stands for "absolutely continuous with respect to $\mathbf{P}$" which in the present setting (finite Ω and $\mathbf{P}$ having full support) automatically holds true, but which will be of relevance for general probability spaces $(\Omega, \mathcal{F}, \mathbf{P})$ later. Note that in the present setting of a finite probability space Ω with $\mathbf{P}[\omega] > 0$ for each $\omega \in \Omega$, we have that $\mathbf{Q} \sim \mathbf{P}$ iff $\mathbf{Q}[\omega] > 0$, for each $\omega \in \Omega$. We shall often identify a measure $\mathbf{Q}$ on $(\Omega, \mathcal{F})$ with its Radon-Nikodým derivative $\frac{d\mathbf{Q}}{d\mathbf{P}} \in L^1(\Omega, \mathcal{F}, \mathbf{P})$. In the present setting of finite Ω, this simply means

$$\frac{d\mathbf{Q}}{d\mathbf{P}}(\omega) = \frac{\mathbf{Q}[\omega]}{\mathbf{P}[\omega]}.$$

In statistics this quantity is also called the likelihood ratio.

Lemma 2.2.6. *For a probability measure* $\mathbf{Q}$ *on* $(\Omega, \mathcal{F})$ *the following are equivalent:*

(i) $\mathbf{Q} \in \mathcal{M}^a(S)$,
(ii) $\mathbf{E}_\mathbf{Q}[f] = 0$, *for all* $f \in K$,
(iii) $\mathbf{E}_\mathbf{Q}[g] \leq 0$, *for all* $g \in C$.

Proof. The equivalences are rather trivial. (ii) is tantamount to the very definition of S being a martingale under $\mathbf{Q}$, i.e., to the validity of

$$\mathbf{E}_\mathbf{Q}[S_t \mid \mathcal{F}_{t-1}] = S_{t-1}, \quad \text{for } t = 1, \ldots, T. \tag{2.5}$$

Indeed, (2.5) holds true iff for each $\mathcal{F}_{t-1}$-measurable set A we have $\mathbf{E}_{\mathbf{Q}}[\chi_A(S_t - S_{t-1})] = 0 \in \mathbb{R}^d$, in other words $\mathbf{E}_{\mathbf{Q}}[(x\chi_A, \Delta S_t)] = 0$, for each x. By linearity this relation extends to K which shows (ii).

The equivalence of (ii) and (iii) is straightforward. □

After having fixed these formalities we may formulate and prove the central result of the theory of pricing and hedging by no-arbitrage, sometimes called the "Fundamental Theorem of Asset Pricing", which in its present form (i.e., finite Ω) is due to M. Harrison and S.R. Pliska [HP 81].

Theorem 2.2.7 (Fundamental Theorem of Asset Pricing). *For a financial market S modelled on a finite stochastic base $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t=0}^T, \mathbf{P})$, the following are equivalent:*

(i) *S satisfies (NA),*
(ii) $\mathcal{M}^e(S) \neq \emptyset$.

Proof. (ii) $\Rightarrow$ (i): This is the obvious implication. If there is some $\mathbf{Q} \in \mathcal{M}^e(S)$ then by Lemma 2.2.6 we have that

$$\mathbf{E}_{\mathbf{Q}}[g] \leq 0, \quad \text{for } g \in C.$$

On the other hand, if there were $g \in C \cap L_+^\infty$, $g \neq 0$, then, using the assumption that $\mathbf{Q}$ is equivalent to $\mathbf{P}$, we would have

$$\mathbf{E}_{\mathbf{Q}}[g] > 0,$$

a contradiction.

(i) $\Rightarrow$ (ii) This implication is the important message of the theorem which will allow us to link the no-arbitrage arguments with martingale theory. We give a functional analytic existence proof, which will be extendable — in spirit — to more general situations.

By assumption the space K intersects L_+^∞ only at 0. We want to separate the disjoint convex sets $L_+^\infty \setminus \{0\}$ and K by a hyperplane induced by a linear functional $\mathbf{Q} \in L^1(\Omega, \mathcal{F}, \mathbf{P})$. In order to get a strict separation of K and $L_+^\infty \setminus \{0\}$ we have to be a little careful since the standard separation theorems do not directly apply.

One way to overcome this difficulty (in finite dimension) is to consider the convex hull of the unit vectors $\left(\mathbf{1}_{\{\omega_n\}}\right)_{n=1}^N$ in $L^\infty(\Omega, \mathcal{F}, \mathbf{P})$ i.e.

$$P := \left\{ \sum_{n=1}^N \mu_n \mathbf{1}_{\{\omega_n\}} \;\middle|\; \mu_n \geq 0, \sum_{n=1}^N \mu_n = 1 \right\}.$$

This is a convex, compact subset of $L_+^\infty(\Omega, \mathcal{F}, \mathbf{P})$ and, by the *(NA)* assumption, disjoint from K. Hence we may strictly separate the convex compact set P from the convex closed set K by a linear functional $\mathbf{Q} \in L^\infty(\Omega, \mathcal{F}, \mathbf{P})^* = L^1(\Omega, \mathcal{F}, \mathbf{P})$, i.e., find $\alpha < \beta$ such that

$$(\mathbf{Q}, f) \leq \alpha, \quad \text{for } f \in K,$$
$$(\mathbf{Q}, h) \geq \beta, \quad \text{for } h \in P.$$

Since K is a linear space, we have $\alpha \geq 0$ and may replace α by 0. Hence $\beta > 0$. Defining by I the constant vector $I = (1, \ldots, 1)$, we have $(\mathbf{Q}, I) > 0$. We may normalise $\mathbf{Q}$ such that $(\mathbf{Q}, I) = 1$. As $\mathbf{Q}$ is strictly positive on each $\mathbf{1}_{\{\omega_n\}}$, we therefore have found a probability measure $\mathbf{Q}$ on $(\Omega, \mathcal{F})$ equivalent to $\mathbf{P}$ such that condition (ii) of Lemma 2.2.6 holds true. In other words, we found an equivalent martingale measure $\mathbf{Q}$ for the process S. □

The name "Fundamental Theorem of Asset Pricing" was, as far as we are aware, first used in [DR 87]. We shall see that it plays a truly fundamental role in the theory of pricing and hedging of derivative securities (or, synonymously, contingent claims, i.e., elements of $L^0(\Omega, \mathcal{F}, \mathbf{P})$) by no-arbitrage arguments.

It seems worthwhile to discuss the intuitive interpretation of this basic result: a martingale S (say, under the original measure $\mathbf{P}$) is a mathematical model for a *perfectly fair* game. Applying any strategy $H \in \mathcal{H}$ we always have $\mathbf{E}[(H \cdot S)_T] = 0$, i.e., an investor can neither win nor lose in expectation.

On the other hand, a process S allowing for arbitrage, is a model for an utterly unfair game: choosing a good strategy $H \in \mathcal{H}$, an investor can make "something out of nothing". Applying H, the investor is sure not to lose, but has strictly positive probability to gain something.

In reality, there are many processes S which do not belong to either of these two extreme classes. Nevertheless, the above theorem tells us that there is a sharp dichotomy by allowing to *change the odds*. Either a process S is utterly unfair, in the sense that it allows for arbitrage. In this case there is no remedy to make the process fair by changing the odds: it never becomes a martingale. In fact, the possibility of making an arbitrage is not affected by changing the odds, i.e., by passing to an equivalent probability $\mathbf{Q}$. On the other hand, discarding this extreme case of processes allowing for arbitrage, we can always pass from $\mathbf{P}$ to an equivalent measure $\mathbf{Q}$ under which S is a martingale, i.e., a perfectly fair game. Note that the passage from $\mathbf{P}$ to $\mathbf{Q}$ may change the probabilities (the "odds") but not the impossible events (i.e. the null sets).

We believe that this dichotomy is a remarkable fact, also from a purely intuitive point of view.

Corollary 2.2.8. *Let S satisfy (NA) and let $f \in L^\infty(\Omega, \mathcal{F}, \mathbf{P})$ be an attainable contingent claim. In other words f is of the form*

$$f = a + (H \cdot S)_T, \tag{2.6}$$

for some $a \in \mathbb{R}$ and some trading strategy H. Then the constant a and the process $(H \cdot S)_t$ are uniquely determined by (2.6) and satisfy, for every $\mathbf{Q} \in \mathcal{M}^e(S)$,

$$a = \mathbf{E}_{\mathbf{Q}}[f], \quad \text{and} \quad a + (H \cdot S)_t = \mathbf{E}_{\mathbf{Q}}[f \mid \mathcal{F}_t], \quad \text{for} \quad 0 \leq t \leq T. \tag{2.7}$$

Proof. As regards the uniqueness of the constant $a \in \mathbb{R}$, suppose that there are two representations $f = a^1 + (H^1 \cdot S)_T$ and $f = a^2 + (H^2 \cdot S)_T$ with $a^1 \neq a^2$. Assuming w.l.o.g. that $a^1 > a^2$ we find an obvious arbitrage possibility by considering the trading strategy $H^2 - H^1$. We have $a^1 - a^2 = ((H^2 - H^1) \cdot S)_T$, i.e. the trading strategy $H^2 - H^1$ produces a strictly positive result at time T, a contradiction to *(NA)*.

As regards the uniqueness of the process $H \cdot S$, we simply apply a conditional version of the previous argument: assume that $f = a + (H^1 \cdot S)_T$ and $f = a + (H^2 \cdot S)_T$ and suppose that the processes $H^1 \cdot S$ and $H^2 \cdot S$ are not identical. Then there is $0 \leq t \leq T$ such that $(H^1 \cdot S)_t \neq (H^2 \cdot S)_t$ and without loss of generality we may suppose that $A := \{(H^1 \cdot S)_t > (H^2 \cdot S)_t\}$ is a non-empty event, which clearly is in $\mathcal{F}_t$. Hence, using the fact hat $(H^1 \cdot S)_T = (H^2 \cdot S)_T$, the trading strategy $H := (H^2 - H^1)\mathbf{1}_A \cdot \mathbf{1}_{]t,T]}$ is a predictable process producing an arbitrage, as $(H \cdot S)_T = 0$ outside A, while $(H \cdot S)_T = (H^1 \cdot S)_t - (H^2 \cdot S)_t > 0$ on A, which again contradicts *(NA)*.

Finally, the equations in (2.7) result from the fact that, for every predictable process H and every $\mathbf{Q} \in \mathcal{M}^a(S)$, the process $H \cdot S$ is a $\mathbf{Q}$-martingale. □

We denote by $\operatorname{cone}(\mathcal{M}^e(S))$ and $\operatorname{cone}(\mathcal{M}^a(S))$ the cones generated by the convex sets $\mathcal{M}^e(S)$ and $\mathcal{M}^a(S)$ respectively. The subsequent Proposition 2.2.9 clarifies the polar relation between these cones and the cone C.

Let $\langle E, E' \rangle$ be two vector spaces in separating duality. This means that there is a bilinear form $\langle\, .\, ,\, .\, \rangle : E \times E' \to \mathbb{R}$, so that if $\langle x, x' \rangle = 0$ for all $x \in E$, we must have $x' = 0$. Similarly if $\langle x, x' \rangle = 0$ for all $x' \in E'$, we must have $x = 0$. Recall (see, e.g., [Sch 99]) that, for a pair (E, E') of vector spaces in separating duality via the scalar product $\langle\, .\, ,\, .\, \rangle$, the polar C^0 of a set C in E is defined by

$$C^0 = \{g \in E' \mid \langle f, g \rangle \leq 1 \text{ for all } f \in C\}.$$

In the case when C is closed under multiplication by positive scalars (e.g., if C is a convex cone) the polar C^0 may equivalently be defined as

$$C^0 = \{g \in E' \mid \langle f, g \rangle \leq 0 \text{ for all } f \in C\}.$$

The *bipolar theorem* (see, e.g., [Sch 99]) states that the bipolar $C^{00} := (C^0)^0$ of a set C in E is the $\sigma(E, E')$-closed convex hull of C.

In the present, finite dimensional case, $E = L^\infty(\Omega, \mathcal{F}_T, \mathbf{P}) = \mathbb{R}^N$ and $E' = L^1(\Omega, \mathcal{F}_T, \mathbf{P}) = \mathbb{R}^N$ the bipolar theorem is easier. In this case there is only one topology on $\mathbb{R}^N$ compatible with its vector space structure, so that we don't have to speak about different topologies such as $\sigma(E, E')$. However, the proof of the bipolar theorem is in the finite dimensional case and in the infinite dimensional case almost the same and follows from the separating hyperplane resp. the Hahn-Banach theorem.

After these general observations we pass to the concrete setting of the cone $C \subseteq L^\infty(\Omega, \mathcal{F}, \mathbf{P})$ of contingent claims super-replicable at price 0. Note that in our finite dimensional setting this convex cone is closed as it is the algebraic sum of the closed linear space K (a linear space in $\mathbb{R}^N$ is always closed) and the closed polyhedral cone $L^\infty_-(\Omega, \mathcal{F}, \mathbf{P})$ (the verification, that the algebraic sum of a space and a polyhedral cone in $\mathbb{R}^N$ is closed, is an easy, but not completely trivial exercise). We deduce from the bipolar theorem, that C equals its bipolar C^{00}.

Proposition 2.2.9. *Suppose that S satisfies (NA). Then the polar of C is equal to* $\mathrm{cone}(\mathcal{M}^a(S))$*, the cone generated by $\mathcal{M}^a(S)$, and $\mathcal{M}^e(S)$ is dense in $\mathcal{M}^a(S)$. Hence the following assertions are equivalent for an element $g \in L^\infty(\Omega, \mathcal{F}, \mathbf{P})$:*

(i) $g \in C$,
(ii) $\mathbf{E}_\mathbf{Q}[g] \leq 0$*, for all* $\mathbf{Q} \in \mathcal{M}^a(S)$,
(iii) $\mathbf{E}_\mathbf{Q}[g] \leq 0$*, for all* $\mathbf{Q} \in \mathcal{M}^e(S)$.

Proof. The fact that the polar C^0 and the set $\mathrm{cone}(\mathcal{M}^a(S))$ coincide, follows from Lemma 2.2.6 and the observation that $C \supseteq L^\infty_-(\Omega, \mathcal{F}, \mathbf{P})$ and $C^0 \subseteq L^1_+(\Omega, \mathcal{F}, \mathbf{P})$. Hence the equivalence of (i) and (ii) follows from the bipolar theorem.

As regards the density of $\mathcal{M}^e(S)$ in $\mathcal{M}^a(S)$ we first deduce from Theorem 2.2.7 that there is at least one $\mathbf{Q}^* \in \mathcal{M}^e(S)$. For any $\mathbf{Q} \in \mathcal{M}^a(S)$ and $0 < \mu \leq 1$ we have that $\mu\mathbf{Q}^* + (1-\mu)\mathbf{Q} \in \mathcal{M}^e(S)$, which clearly implies the density of $\mathcal{M}^e(S)$ in $\mathcal{M}^a(S)$. The equivalence of (ii) and (iii) is now obvious. □

Similarly we can show the following:

Proposition 2.2.10. *Suppose S satisfies (NA). Then for $f \in L^\infty$, the following assertions are equivalent*

(i) *$f \in K$, i.e. $f = (H \cdot S)_T$ for some strategy $H \in \mathcal{H}$.*
(ii) *For all $\mathbf{Q} \in \mathcal{M}^e(S)$ we have $\mathbf{E}_\mathbf{Q}[f] = 0$.*
(iii) *For all $\mathbf{Q} \in \mathcal{M}^a(S)$ we have $\mathbf{E}_\mathbf{Q}[f] = 0$.*

Proof. By Proposition 2.2.4 we have that $f \in K$ iff $f \in C \cap (-C)$. Hence the result follows from the preceding Proposition 2.2.9. □

Corollary 2.2.11. *Assume that S satisfies (NA) and that $f \in L^\infty$ satisfies $\mathbf{E}_\mathbf{Q}[f] = a$ for all $\mathbf{Q} \in \mathcal{M}^e(S)$, then $f = a + (H \cdot S)_T$ for some strategy H.* □

Corollary 2.2.12 (complete financial markets). *For a financial market S satisfying the no-arbitrage condition (NA), the following are equivalent:*

(i) *$\mathcal{M}^e(S)$ consists of a single element $\mathbf{Q}$.*
(ii) *Each $f \in L^\infty(\Omega, \mathcal{F}, \mathbf{P})$ may be represented as*

$$f = a + (H \cdot S)_T \quad \textit{for some } a \in \mathbb{R} \textit{ and } H \in \mathcal{H}.$$

In this case $a = \mathbf{E}_{\mathbf{Q}}[f]$, the stochastic integral $H \cdot S$ is unique and we have that

$$\mathbf{E}_{\mathbf{Q}}[f \mid \mathcal{F}_t] = \mathbf{E}_{\mathbf{Q}}[f] + (H \cdot S)_t, \quad t = 0, \ldots, T. \qquad \square$$

The Fundamental Theorem of Asset Pricing 2.2.7 allows us to prove the following proposition, which we shall need soon.

Proposition 2.2.13. *Assume that S satisfies (NA) and let $H \cdot S$ be the process obtained from S by means of a fixed strategy $H \in \mathcal{H}$. Fix $a \in \mathbb{R}$ and define the $\mathbb{R}$-valued process $S^{d+1} = (S_t^{d+1})_{t=0}^T$ by $S^{d+1} = a + H \cdot S$. Then the process $\overline{S} = (S^1, S^2, \ldots, S^d, S^{d+1})$ also satisfies the (NA) property and the sets $\mathcal{M}^e(S)$ and $\mathcal{M}^e(\overline{S})$ (as well as $\mathcal{M}^a(S)$ and $\mathcal{M}^a(\overline{S})$) coincide.*

Proof. If $\mathbf{Q} \in \mathcal{M}^e(S)$ then $H \cdot S$ is a $\mathbf{Q}$-martingale. Consequently $\overline{S}$ satisfies *(NA).* $\square$

2.3 Equivalence of Single-period with Multiperiod Arbitrage

The aim of this section is to describe the relation between one-period no-arbitrage and multiperiod no-arbitrage. At the same time we will be able to give somewhat more detailed information on the set of risk neutral measures (this term is often used in the finance literature in a synonymous way for martingale measures). We start off with the following observation. Recall that we did not assume that $\mathcal{F}_0$ is trivial.

Proposition 2.3.1. *If S satisfies the no-arbitrage condition, $\mathbf{Q} \in \mathcal{M}^e(S)$ is an equivalent martingale measure, and $Z_t = \mathbf{E}_{\mathbf{P}}\left[\frac{d\mathbf{Q}}{d\mathbf{P}} \,\middle|\, \mathcal{F}_t\right]$ denotes the density process associated with $\mathbf{Q}$, then the process $L_t = \frac{Z_t}{Z_0}$ defines the density process of an equivalent measure $\mathbf{Q}'$ such that $\frac{d\mathbf{Q}'}{d\mathbf{P}} = L_T$, $\mathbf{Q}' \in \mathcal{M}^e(S)$ and $\mathbf{Q}'|_{\mathcal{F}_0} = \mathbf{P}|_{\mathcal{F}_0}$.*

Proof. This is rather straightforward. Since $\mathbf{Q} \in \mathcal{M}^e(S)$ we have that SZ is a $\mathbf{P}$-martingale. Since $Z_0 > 0$ and since it is $\mathcal{F}_0$-measurable the process $S\frac{Z}{Z_0}$ is still a $\mathbf{P}$-martingale. Since SL is now a $\mathbf{P}$-martingale and since the density $L_T > 0$, we necessarily have $\mathbf{Q}' \in \mathcal{M}^e(S)$. As $L_0 = 1$ we obtain $\mathbf{Q}'|_{\mathcal{F}_0} = \mathbf{P}|_{\mathcal{F}_0}$. $\square$

Theorem 2.3.2. *Let $S = (S_t)_{t=0}^T$ be a price process. Then the following are equivalent:*

(i) *S satisfies the no-arbitrage property.*
(ii) *For each $0 \le t < T$, we have that the one-period market (S_t, S_{t+1}) with respect to the filtration $(\mathcal{F}_t, \mathcal{F}_{t+1})$ satisfies the no-arbitrage property.*

Proof. Obviously (i) implies (ii), since there are less strategies in each single period market than in the multiperiod market. So let us show that (ii) implies (i). By the fundamental theorem applied to (S_t, S_{t+1}), we have that for each t there is a probability measure $\mathbf{Q}_t$ on $\mathcal{F}_{t+1}$ equivalent to $\mathbf{P}$, so that under $\mathbf{Q}_t$ the process (S_t, S_{t+1}) is a $\mathbf{Q}_t$-martingale. This means that $\mathbf{E}_{\mathbf{Q}_t}[S_{t+1} \mid \mathcal{F}_t] = S_t$. By the previous proposition we may take $\mathbf{Q}_t|_{\mathcal{F}_t} = \mathbf{P}|_{\mathcal{F}_t}$. Let $f_{t+1} = \frac{d\mathbf{Q}_t}{d\mathbf{P}}$ and define $L_t = f_1 \dots f_{t-1} f_t$ and $L_0 = 1$. Clearly $(L_t)_{t=0}^T$ is the density process of an equivalent measure $\mathbf{Q}$ defined by $\frac{d\mathbf{Q}}{d\mathbf{P}} = L_T$. One can easily check that, for all $t = 0, \dots, T-1$ we have $\mathbf{E}_{\mathbf{Q}}[S_{t+1} \mid \mathcal{F}_t] = S_t$, i.e., $\mathbf{Q} \in \mathcal{M}^e(S)$. $\square$

Remark 2.3.3. The equivalence between one-period no-arbitrage and multi-period no-arbitrage can also be checked directly by the definition of no-arbitrage. We invite the reader to give a direct proof of the following: if H is a strategy so that $(H \cdot S)_T \geq 0$ and $\mathbf{P}[(H \cdot S)_T > 0] > 0$ then there is a $1 \leq t \leq T$ as well as $A \in \mathcal{F}_{t-1}$, $\mathbf{P}[A] > 0$ so that $\mathbf{1}_A(H_t, \Delta S_t) \geq 0$ and $\mathbf{P}[\mathbf{1}_A(H_t, \Delta S_t) > 0] > 0$ (compare Lemma 5.1.5 below).

Remark 2.3.4. We give one more indication, why there is little difference between the one-period and the T period situation; this discussion also reveals a nice economic interpretation. Given $S = (S_t)_{t=0}^T$ as above, we may associate a one-period process $\widetilde{S} = (\widetilde{S}_t)_{t=0}^1$, adapted to the filtration $(\widetilde{\mathcal{F}}_0, \widetilde{\mathcal{F}}_1) := (\mathcal{F}_0, \mathcal{F}_T)$ in the following way: choose any collection $(f_1, \dots, f_m)$ in the finite dimensional linear space K defined in 2.2.1, which linearly spans K. Define the $\mathbb{R}^m$-valued process $\widetilde{S}$ by $\widetilde{S}_0 = 0$, $\widetilde{S}_1 = (f_1, \dots, f_m)$.

Obviously the process $\widetilde{S}$ yields the same space K of stochastic integrals as S. Hence the set of equivalent martingale measures for the processes S and $\widetilde{S}$ coincide and therefore all assertions, depending only on the set of equivalent martingale measures coincide for S and $\widetilde{S}$. In particular S and $\widetilde{S}$ yield the same arbitrage-free prices for derivatives, as we shall see in the next section.

The economic interpretation of the transition from S to $\widetilde{S}$ reads as follows: if we fix the trading strategies H^j yielding $f_j = (H^j \cdot S)_T$, we may think of f_j as a contingent claim at time $t = T$ which may be bought at price 0 at time $t = 0$, by then applying the trading rules given by H^j. By taking sufficiently many of these H^j's, in the sense that the corresponding f_j's linearly span K, we may represent the result $f = (H \cdot S)_T$ of *any* trading strategy H as a linear combination of the f_j's.

The bottom line of this discussion is that in the present framework (i.e. Ω is finite) — from a mathematical as well as from an economic point of view — the T period situation can easily be reduced to the one-period situation.

2.4 Pricing by No-Arbitrage

The subsequent theorem will tell us what the principle of no-arbitrage implies about the possible prices for a contingent claim f. It goes back to the work of D. Kreps [K 81].

For given $f \in L^\infty(\Omega, \mathcal{F}, \mathbf{P})$, we call $a \in \mathbb{R}$ an *arbitrage-free price*, if in addition to the financial market S, the introduction of the contingent claim f at price a does not create an arbitrage possibility. How can we mathematically formalise this economically intuitive idea? We enlarge the financial market S by introducing a new financial instrument which can be bought (or sold) at price a at time $t = 0$ and yields the random cash flow $f(\omega)$ at time $t = T$. We don't postulate anything about the price of this financial instrument at the intermediate times $t = 1, \dots, T-1$. The reader might think of an "over the counter" option where the two parties agree on certain payments at times $t = 0$ and $t = T$. So if we look at the linear space generated by K and the vector $(f-a)$ we obtain an enlarged space $K^{f,a}$ of attainable claims. The price a should be such that arbitrage opportunities are inexistent. Mathematically speaking this means that we still should have $K^{f,a} \cap L^\infty_+ = \{0\}$. In this case we say that a is an *arbitrage free price* for the contingent claim f.

Theorem 2.4.1 (Pricing by no-arbitrage). *Assume that S satisfies (NA) and let $f \in L^\infty(\Omega, \mathcal{F}, \mathbf{P})$. Define*

$$\begin{aligned} \underline{\pi}(f) &= \inf \{\mathbf{E}_{\mathbf{Q}}[f] \mid \mathbf{Q} \in \mathcal{M}^e(S)\}, \\ \overline{\pi}(f) &= \sup \{\mathbf{E}_{\mathbf{Q}}[f] \mid \mathbf{Q} \in \mathcal{M}^e(S)\}, \end{aligned} \tag{2.8}$$

Either $\underline{\pi}(f) = \overline{\pi}(f)$, in which case f is attainable at price $\pi(f) := \underline{\pi}(f) = \overline{\pi}(f)$, i.e. $f = \pi(f) + (H \cdot S)_T$ for some $H \in \mathcal{H}$ and therefore $\pi(f)$ is the unique arbitrage-free price for f.

Or $\underline{\pi}(f) < \overline{\pi}(f)$, in which case

$$]\underline{\pi}(f), \overline{\pi}(f)[= \{\mathbf{E}_{\mathbf{Q}}[f] \mid \mathbf{Q} \in \mathcal{M}^e(S)\}$$

and a is an arbitrage-free price for f iff a lies in the open *interval $]\underline{\pi}(f), \overline{\pi}(f)[$.*

Proof. The case $\underline{\pi}(f) = \overline{\pi}(f)$ follows from corollary 2.2.11 and so we only have to concentrate on the case $\underline{\pi}(f) < \overline{\pi}(f)$. First observe that the set $\{\mathbf{E}_{\mathbf{Q}}[f] \mid \mathbf{Q} \in \mathcal{M}^e(S)\}$ forms a bounded non-empty interval in $\mathbb{R}$, which we denote by I.

We claim that a number a is in I iff a is an arbitrage-free price for f. Indeed, supposing that $a \in I$ we may find $\mathbf{Q} \in \mathcal{M}^e(S)$ s.t. $\mathbf{E}_{\mathbf{Q}}[f - a] = 0$ and therefore $K^{f,a} \cap L^\infty_+(\Omega, \mathcal{F}, \mathbf{P}) = \{0\}$.

Conversely suppose that $K^{f,a} \cap L^\infty_+ = \{0\}$. Then exactly as in the proof of the Fundamental Theorem 2.2.7, we find a probability measure $\mathbf{Q}$ so that $\mathbf{E}_{\mathbf{Q}}[g] = 0$ for all $g \in K^{f,a}$ and so that $\mathbf{Q}$ is equivalent to $\mathbf{P}$. This, of course, implies that $\mathbf{Q} \in \mathcal{M}^e(S)$ and that $a = \mathbf{E}_{\mathbf{Q}}[f]$.

Now we deal with the boundary case: suppose that a equals the right boundary of I, i.e., $a = \overline{\pi}(f) \in I$, and consider the contingent claim $f - \overline{\pi}(f)$. By definition we have $\mathbf{E}_{\mathbf{Q}}[f - \overline{\pi}(f)] \leq 0$, for all $\mathbf{Q} \in \mathcal{M}^e(S)$, and therefore by Proposition 2.2.9, that $f - \overline{\pi}(f) \in C$. We may find $g \in K$ such that $g \geq f - \overline{\pi}(f)$. If the sup in (2.8) is attained, i.e., if there is $\mathbf{Q}^* \in \mathcal{M}^e(S)$ such that $\mathbf{E}_{\mathbf{Q}^*}[f] = \overline{\pi}(f)$, then we have $0 = \mathbf{E}_{\mathbf{Q}^*}[g] \geq \mathbf{E}_{\mathbf{Q}^*}[f - \overline{\pi}(f)] = 0$ which in view of $\mathbf{Q}^* \sim \mathbf{P}$ implies that $f - \overline{\pi}(f) \equiv g$; in other words f is attainable at price $\overline{\pi}(f)$. This in turn implies that $\mathbf{E_Q}[f] = \overline{\pi}(f)$ for all $\mathbf{Q} \in \mathcal{M}^e(S)$, and therefore I is reduced to the singleton $\{\overline{\pi}(f)\}$.

Hence, if $\underline{\pi}(f) < \overline{\pi}(f)$, $\overline{\pi}(f)$ cannot belong to the interval I, which is therefore open on the right hand side. Passing from f to $-f$, we obtain the analogous result for the left hand side of I, which is therefore equal to $I =]\underline{\pi}(f), \overline{\pi}(f)[$. □

The argument in the proof of the preceding theorem can be recast to yield the following duality theorem. The reader familiar with the duality theory of linear programming will recognise the primal-dual relation.

Theorem 2.4.2 (Superreplication). *Assume that S satisfies (NA). Then, for $f \in L^\infty$, we have*

$$\begin{aligned} \overline{\pi}(f) &= \sup\{\mathbf{E_Q}[f] \mid \mathbf{Q} \in \mathcal{M}^e(S)\} \\ &= \max\{\mathbf{E_Q}[f] \mid \mathbf{Q} \in \mathcal{M}^a(S)\} \\ &= \min\{a \mid \textit{there exists } k \in K, a + k \geq f\}. \end{aligned}$$

Proof. As shown in the previous proof we have $f - \overline{\pi}(f) \in C$ and hence

$$\begin{aligned} f &= \overline{\pi}(f) + g, && \text{for some } g \in C \\ &= \overline{\pi}(f) + k - h, && \text{for some } k \in K \text{ and } h \in L^\infty_+ \\ &\leq \overline{\pi}(f) + k, && \text{for some } k \in K. \end{aligned}$$

This shows that $\overline{\pi}(f) \geq \inf\{a \mid \text{there exists } k \in K, a + k \geq f\}$.

Let now $a < \overline{\pi}(f)$. We will show that there is no element $k \in K$ with $a + k \geq f$. This shows that $\overline{\pi}(f) = \inf\{a \mid \text{there exists } k \in K, a + k \geq f\}$ and moreover establishes that the infimum is a minimum. Since $a < \overline{\pi}(f)$ there is $\mathbf{Q} \in \mathcal{M}^e(S)$ with $\mathbf{E_Q}[f] > a$. But this implies that for all $k \in K$ we have that $\mathbf{E_Q}[a + k] = a < \mathbf{E_Q}[f]$, in contradiction to the relation $a + k \geq f$. □

Remark 2.4.3. Theorem 2.4.2 may be rephrased in economic terms: in order to superreplicate f, i.e., to find $a \in \mathbb{R}$ and $H \in \mathcal{H}$ s.t. $a + (H \cdot S)_T \geq f$, we need at least an initial investment a equal to $\overline{\pi}(f)$.

We now give a conditional version of the duality theorem that allows us to use initial investments that are not constant and to possibly use the information $\mathcal{F}_0$ available at time $t = 0$. This is relevant when the initial σ-algebra $\mathcal{F}_0$ is not trivial.

Theorem 2.4.4. *Let us assume that S satisfies (NA). Denote by $\mathcal{M}^e(S, \mathcal{F}_0)$ the set of equivalent martingale measures $\mathbf{Q} \in \mathcal{M}^e(S)$ so that $\mathbf{Q}|_{\mathcal{F}_0} = \mathbf{P}$. Then, for $f \in L^\infty$, we have*

$$\begin{aligned} &\sup\{\mathbf{E_Q}[f \mid \mathcal{F}_0] \mid \mathbf{Q} \in \mathcal{M}^e(S, \mathcal{F}_0)\} \\ &\quad = \min\{h \mid h \text{ is } \mathcal{F}_0\text{-measurable and there exists } g \in K \text{ such that } h + g \geq f\}. \end{aligned}$$

Remark 2.4.5. Before we prove the theorem let us remark that the "sup" and the "min" are taken in the space $L^0(\Omega, \mathcal{F}_0, \mathbf{P})$ of $\mathcal{F}_0$-measurable functions. Both sets are lattice ordered. Indeed, if $\mathbf{E}_{\mathbf{Q}_1}[f \mid \mathcal{F}_0]$ and $\mathbf{E}_{\mathbf{Q}_2}[f \mid \mathcal{F}_0]$ are given, where $\mathbf{Q}_1, \mathbf{Q}_2 \in \mathcal{M}^e(S, \mathcal{F}_0)$, then there is an element $\mathbf{Q}_3 \in \mathcal{M}^e(S, \mathcal{F}_0)$ so that $\mathbf{E}_{\mathbf{Q}_3}[f \mid \mathcal{F}_0] = \max\{\mathbf{E}_{\mathbf{Q}_1}[f \mid \mathcal{F}_0], \mathbf{E}_{\mathbf{Q}_2}[f \mid \mathcal{F}_0]\}$. The construction is rather straightforward. Let $A = \{\mathbf{E}_{\mathbf{Q}_1}[f \mid \mathcal{F}_0] > \mathbf{E}_{\mathbf{Q}_2}[f \mid \mathcal{F}_0]\} \in \mathcal{F}_0$ and let $\mathbf{Q}_3[B] = \mathbf{Q}_1[A \cap B] + \mathbf{Q}_2[A^c \cap B]$. Because $\mathbf{Q}_1|_{\mathcal{F}_0} = \mathbf{Q}_2|_{\mathcal{F}_0} = \mathbf{P}$ we get that $\mathbf{Q}_3$ is a probability and that $\mathbf{Q}_3 \in \mathcal{M}^e(S, \mathcal{F}_0)$. Also $\mathbf{E}_{\mathbf{Q}_3}[f \mid \mathcal{F}_0] = \mathbf{E}_{\mathbf{Q}_1}[f \mid \mathcal{F}_0] \vee \mathbf{E}_{\mathbf{Q}_2}[f \mid \mathcal{F}_0]$.

Similarly, the set on the right is stable for the "min" operation. Indeed, let $h_1 + g_1 \geq f$ and $h_2 + g_2 \geq f$. For $A = \{h_1 < h_2\}$, an $\mathcal{F}_0$-measurable set, we define $h = h_1 \mathbf{1}_A + h_2 \mathbf{1}_{A^c}$ and $g_1 \mathbf{1}_A + g_2 \mathbf{1}_{A^c} = g$. The function h is $\mathcal{F}_0$-measurable and $g \in K$ (because $A \in \mathcal{F}_0$). Clearly $h + g \geq f$.

Proof of Theorem 2.4.4. If $f \leq h + g$, where h is $\mathcal{F}_0$-measurable and $g \in K$, then for $\mathbf{Q} \in \mathcal{M}^e(S, \mathcal{F}_0)$ we have $\mathbf{E}_{\mathbf{Q}}[f \mid \mathcal{F}_0] \leq h + 0 = h$. This shows that

$$\begin{aligned} a_1 &:= \sup \{\mathbf{E}_{\mathbf{Q}}[f \mid \mathcal{F}_0] \mid \mathbf{Q} \in \mathcal{M}^e(S, \mathcal{F}_0)\} \\ &\leq \inf \{h \mid h\ \mathcal{F}_0\text{-measurable}, h + g \geq f, \text{ for some } g \in K\} \\ &=: a_2. \end{aligned}$$

To prove the converse inequality, we show that there is $g \in K$ with $a_1 + g \geq f$. If this were not be true then $(a_1 + K) \cap (f + L^\infty_+) = \emptyset$ and we could find, using the separating hyperplane theorem, a linear functional φ and $\varepsilon > 0$, so that $\forall g \in K$, $\forall\, l \geq 0$ we have $\varepsilon + \varphi(a_1 + g) < \varphi(f + l)$. This implies that $\varphi \geq 0$ and $\varphi(g) = 0$ for all $g \in K$. Of course we can normalise φ so that it comes from a probability measure $\mathbf{Q}$. So we get $\mathbf{E}_{\mathbf{Q}}[a_1] + \varepsilon' < \mathbf{E}_{\mathbf{Q}}[f]$ and $\mathbf{Q} \in \mathcal{M}^a(S)$, where $\varepsilon' > 0$.

By the density of $\mathcal{M}^e(S)$ in $\mathcal{M}^a(S)$ we may perturb $\mathbf{Q}$ a little bit to make it an element of $\mathcal{M}^e(S)$. We still get $\mathbf{E}_{\mathbf{Q}}[a_1] + \varepsilon < \mathbf{E}_{\mathbf{Q}}[f]$, but this time for a measure $\mathbf{Q} \in \mathcal{M}^e(S)$. Let now $Z_t = \frac{d\mathbf{Q}}{d\mathbf{P}}\Big|_{\mathcal{F}_t}$ and set $L_t = \frac{Z_t}{Z_0}$. The process $(L_t)_{t=0}^\infty$ defines a measure $\mathbf{Q}^0 \in \mathcal{M}^e(S, \mathcal{F}_0)$ via $\frac{d\mathbf{Q}^0}{d\mathbf{P}} = L_T$. Furthermore

$$\begin{aligned} \mathbf{E}_{\mathbf{Q}^0}[f \mid \mathcal{F}_0] &= \mathbf{E}_{\mathbf{P}}[f L_T \mid \mathcal{F}_0] \\ &= \frac{\mathbf{E}_{\mathbf{P}}[f Z_T \mid \mathcal{F}_0]}{Z_0} = \mathbf{E}_{\mathbf{Q}}[f \mid \mathcal{F}_0] \end{aligned}$$

Therefore $\mathbf{E}_{\mathbf{Q}}[f \mid \mathcal{F}_0] \leq a_1$ and hence $\mathbf{E}_{\mathbf{Q}}[f] \leq \mathbf{E}_{\mathbf{Q}}[a_1]$, contradicting the choice of $\mathbf{Q}$. □

Corollary 2.4.6. *Under the assumptions of Theorem 2.4.4 we have*

$$\{\mathbf{E}_{\mathbf{Q}}[f \mid \mathcal{F}_0] \mid \mathbf{Q} \in \mathcal{M}^e(S)\} = \{\mathbf{E}_{\mathbf{Q}}[f \mid \mathcal{F}_0] \mid \mathbf{Q} \in \mathcal{M}^e(S, \mathcal{F}_0)\}.$$

Hence, for $f \in L^\infty_+(\Omega, \mathcal{F}, \mathbf{P})$, *we have* $\sup_{\mathbf{Q} \in \mathcal{M}^e(S)} \mathbf{E}_{\mathbf{Q}}[f] = \|a_1\|_\infty$ *where*

$$a_1 = \sup \{\mathbf{E}_{\mathbf{Q}}[f \mid \mathcal{F}_0] \mid \mathbf{Q} \in \mathcal{M}^e(S, \mathcal{F}_0)\}.$$

Proof. As observed in the proof of Theorem 2.3.2 and Proposition 2.3.1, every $\mathbf{Q} \in \mathcal{M}^e(S)$ can be written as $\frac{d\mathbf{Q}}{d\mathbf{P}} = f_0 \frac{d\mathbf{Q}^0}{d\mathbf{P}}$ where $\mathbf{Q}^0|_{\mathcal{F}_0} = \mathbf{P}|_{\mathcal{F}_0}$, $\mathbf{Q}^0 \in \mathcal{M}^e(S, \mathcal{F}_0)$ and where f_0 is $\mathcal{F}_0$-measurable, strictly positive and $\mathbf{E}_\mathbf{P}[f_0] = 1$. But otherwise f_0 is arbitrary. Now for $\mathbf{Q} \in \mathcal{M}^e(S)$ we have $\frac{d\mathbf{Q}}{d\mathbf{P}} = f_0 \frac{d\mathbf{Q}^0}{d\mathbf{P}}$ and hence

$$\begin{aligned} \mathbf{E}_\mathbf{Q}[f] &= \mathbf{E}_\mathbf{Q}\left[\mathbf{E}_{\mathbf{Q}^0}[f \mid \mathcal{F}_0]\right] \\ &\leq \mathbf{E}_\mathbf{Q}[a_1] = \mathbf{E}_\mathbf{P}[a_1 f_0]. \end{aligned}$$

Thus $\sup_{\mathbf{Q} \in \mathcal{M}^e(S)} \mathbf{E}_\mathbf{Q}[f] \leq \|a_1\|_\infty$.

To prove the converse inequality we need some more approximations. First for given $\varepsilon > 0$, we choose f_0, $\mathcal{F}_0$-measurable, $f_0 > 0$, $\mathbf{E}_\mathbf{P}[f_0] = 1$ and so that $\mathbf{E}_\mathbf{P}[f_0 a_1] \geq \|a_1\|_\infty - \varepsilon$. Given f_0 we may take $\mathbf{Q}^1 \in \mathcal{M}^e(S, \mathcal{F}_0)$ so that $\mathbf{E}[f_0(a_1 - \mathbf{E}_{\mathbf{Q}^1}[f \mid \mathcal{F}_0])] \leq \varepsilon$. This is possible since the family $\{\mathbf{E}_\mathbf{Q}[f \mid \mathcal{F}_0] \mid \mathbf{Q} \in \mathcal{M}^e(S, \mathcal{F}_0)\}$ is a lattice and since all these functions are in the L^∞-ball with radius $\|f\|_\infty$. Now take $\mathbf{Q}^0$ defined by $\frac{d\mathbf{Q}^0}{d\mathbf{P}} = f_0 \frac{d\mathbf{Q}^1}{d\mathbf{P}}$. Clearly $\mathbf{Q}^0 \in \mathcal{M}^e(S)$ and we have

$$\begin{aligned} \mathbf{E}_{\mathbf{Q}^0}[f] &= \mathbf{E}_\mathbf{P}\left[f_0 \frac{d\mathbf{Q}^1}{d\mathbf{P}} f\right] \\ &= \mathbf{E}_\mathbf{P}\left[f_0 \mathbf{E}_{\mathbf{Q}^1}[f \mid \mathcal{F}_0]\right] \text{ since } \mathbf{Q}^1|_{\mathcal{F}_0} = \mathbf{P} \\ &\geq \mathbf{E}_\mathbf{P}\left[f_0 a_1\right] - \varepsilon \text{ by the choice of } \mathbf{Q}^1 \\ &\geq \|a_1\|_\infty - 2\varepsilon \text{ by the choice of } f_0. \end{aligned}$$

□

2.5 Change of Numéraire

In the previous sections we have developed the basic tools for the pricing and hedging of derivative securities. Recall that we did our analysis in a *discounted model* where we did choose one of the traded assets as numéraire.

How do these things change, when we pass to a new numéraire, i.e., a new unit in which we denote the values of the stocks? Of course, the arbitrage free prices should remain unchanged (after denominating things in the new numéraire), as the notion of arbitrage should not depend on whether we do the book-keeping in € or in \$. On the other hand, we shall see that the risk-neutral measures $\mathbf{Q}$ do depend on the choice of numéraire. We will also show how, conversely, a change of risk neutral measures corresponds to a change of numéraire.

Let us analyse the situation in the appropriate degree of generality: the model of a financial market $\widehat{S} = (\widehat{S}_t^0, \widehat{S}_t^1, \ldots, \widehat{S}_t^d)_{t=0}^T$ is defined as in 2.1 above. Recall that we assumed that the traded asset $\widehat{S}^0$ serves as numéraire, i.e., we have passed from the value $\widehat{S}_t^j$ of the j'th asset at time t to its value $S_t^j = \frac{\widehat{S}_t^j}{\widehat{S}_t^0}$, expressed in units of $\widehat{S}_t^0$. This led us in (2.3) to the introduction of the process

$$S = (S^1, S^2, \dots, S^d) = \left(\frac{\widehat{S}^1}{\widehat{S}^0}, \dots, \frac{\widehat{S}^d}{\widehat{S}^0} \right).$$

Before we prove the theorem, let us first see what assets can be used as numéraire. The crucial requirement on a numéraire is that it is a *traded asset*. We could of course use one of the assets $1, \dots, d$ but we want to be more general and also want to accept, e.g., baskets as new numéraires. So we might use the value $(V_t)_{t=0}^T$ of a portfolio as a numéraire. Of course, we need to assume $V_t > 0$ for all t. Indeed, if the numéraire becomes zero or even negative, then we obviously have a problem in calculating the value of an asset in terms of V. Further, for normalisation reasons, it is convenient to assume that $V_0 = 1$, exactly as we did for $\widehat{S}^0$. So we start with a value process $V = 1 + (H^0 \cdot S)$ satisfying $V_t > 0$ a.s. for all t, where H^0 is a fixed element of $\mathcal{H}$. Observe that the processes V and S are denoted in terms of our originally chosen numéraire asset $\widehat{S}^0$.

As we have seen above (Proposition 2.2.13), the extended market

$$S^{\text{ext}} = (S^1, S^2, \dots, S^d, 1, V) \tag{2.9}$$

is still arbitrage free and $\mathcal{M}^e(S) = \mathcal{M}^e(S^{\text{ext}})$. In real money terms this process is described by the process

$$\begin{aligned} \widehat{S}^{\text{ext}} &= \left(S^1 \widehat{S}^0, \, \dots, \, S^d \widehat{S}^0, \, \widehat{S}^0, \, V\widehat{S}^0 \right) \\ &= \left(\widehat{S}^1, \, \dots, \, \widehat{S}^d, \, \widehat{S}^0, \, V\widehat{S}^0 \right). \end{aligned}$$

If we now use the last coordinate as numéraire, we obtain the process

$$X = \left(\frac{S^1}{V}, \dots, \frac{S^d}{V}, \frac{1}{V}, 1 \right). \tag{2.10}$$

In order to keep the notation more symmetric we will drop the dummy entry 1 and use $(d+1)$-dimensional predictable processes as strategies. Similarly we shall also drop in (2.9) the dummy entry 1 for S^{ext}. This allows us to pass more easily from S^{ext} to X.

The next lemma shows the economically rather obvious fact that when passing from S to S^{ext}, the space K of claims attainable at price 0 does not change.

Lemma 2.5.1. *Using the above notation we have*

$$\begin{aligned} K(S^{\text{ext}}) &= \{(H \cdot S^{\text{ext}})_T \mid H \ \ (d+1)\text{-}\textit{dimensional predictable}\} \\ &= K(S) = \{(H' \cdot S)_T \mid H' \ \ d\text{-}\textit{dimensional predictable}\}. \end{aligned}$$

Proof. The process V is given by the stochastic integral $(H^0 \cdot S)$ with respect to S, so we expect that nothing new can be created by using the additional

V. It suffices to show that, for a one-dimensional predictable process L, the quantities $L_t \Delta V_t$ are in $K(S)$. This is easy, since

$$L_t \Delta V_t = L_t \left(H_t^0, \Delta S_t \right) = \left(L_t H_t^0, \Delta S_t \right) \in K(S)$$

by definition of $K(S)$. This shows that $K(S^{\mathrm{ext}}) = K(S)$. □

Lemma 2.5.2. *Fix $0 \le t \le T$, and let $f \in K(S) = K(S^{\mathrm{ext}})$ be $\mathcal{F}_t$-measurable. Then the random variable $\frac{f}{V_t}$ is of the form $\frac{f'}{V_T}$ where $f' \in K(S)$.*

Proof. Clearly

$$\frac{f}{V_t} - \frac{f}{V_T} = \frac{1}{V_T} \left(f \frac{V_T - V_t}{V_t} \right) = \frac{1}{V_T} \sum_{s=t+1}^{T} \frac{f}{V_t} \left(V_s - V_{s-1} \right).$$

We see that $f'' = \sum_{s=t+1}^{T} \frac{f}{V_t} (V_s - V_{s-1})$ belongs to $K(S^{\mathrm{ext}})$ because $\frac{f}{V_t}$ is $\mathcal{F}_t$-measurable and the summation is on $s > t$. Hence $f' = f'' + f$ does the job. □

Proposition 2.5.3. *Assume that X is defined as in (2.10). Then*

$$K(X) = \left\{ \frac{f}{V_T} \,\middle|\, f \in K(S) \right\}.$$

Proof. We have that $g \in K(X)$ if and only if there is a $(d+1)$-dimensional predictable process H, with $g = \sum_{t=1}^{T} (H_t, \Delta X_t) = \sum_{t=1}^{T} \sum_{j=1}^{d+1} H_t^j \Delta X_t^j$. Clearly, for $j = 1, \dots, d$ and $t = 1, \dots, T$,

$$\begin{aligned} \Delta X_t^j &= \left(\frac{S_t^j}{V_t} - \frac{S_{t-1}^j}{V_{t-1}} \right) \\ &= \frac{\Delta S_t^j}{V_t} + S_{t-1}^j \left(\frac{1}{V_t} - \frac{1}{V_{t-1}} \right) \\ &= \frac{\Delta S_t^j}{V_t} - \frac{S_{t-1}^j}{V_{t-1}} \frac{\Delta V_t}{V_t} \\ &= \frac{1}{V_t} \left(\Delta S_t^j - X_{t-1}^j \Delta V_t \right). \end{aligned}$$

So we get that $H_t^j \Delta X_t^j = \frac{1}{V_t} \left(H_t^j \Delta S_t^j - \left(H_t^j X_{t-1}^j \right) \Delta V_t \right)$, which is of the form $\frac{f}{V_t}$ for some $f \in K(S^{\mathrm{ext}}) = K(S)$. For $j = d+1$ and $t = 1, \dots, T$ the same argument applies by replacing S_t^j and S_{t-1}^j by 1.

By the previous lemma we have $\frac{f}{V_t} = \frac{f'}{V_T}$ for some $f' \in K(S)$. This shows that $K(X) \subset \frac{1}{V_T} K(S)$.

The converse inclusion follows by symmetry. In the financial market modelled by X we can choose $W_t = \frac{1}{V_t}$ as numéraire. The passage from X to S^{ext} is then done by using W as a new numéraire and the inclusion we just proved then yields

$$K(S) \subset \frac{1}{W_T} K(X) = V_T K(X).$$

This shows that $K(S) = V_T K(X)$ as required. □

Theorem 2.5.4 (change of numéraire). *Let S satisfy the no-arbitrage condition, let $V = 1 + H^0 \cdot S$ be such that $V_t > 0$ for all t, and let $X = \left(\frac{S^1}{V}, \ldots, \frac{S^d}{V}, \frac{1}{V}\right)$. Then X satisfies the no-arbitrage condition too and $\mathbf{Q}$ belongs to $\mathcal{M}^e(S)$ if and only if the measure $\mathbf{Q}'$ defined by $d\mathbf{Q}' = V_T d\mathbf{Q}$ belongs to $\mathcal{M}^e(X)$.*

Proof. Since $K(X) = \frac{1}{V_T} K(S)$ we have that X satisfies the no-arbitrage property by directly verifying Definition 2.2.3. By Proposition 2.2.10 an equivalent probability measure $\mathbf{Q}$ is in $\mathcal{M}^e(S)$ if and only if, for all $f \in K(S)$, we have $\mathbf{E}_\mathbf{Q}[f] = 0$. But this is the same as

$$\mathbf{E}_\mathbf{Q}\left[V_T \frac{f}{V_T}\right] = 0, \quad \text{for all } f \in K(S),$$

which is equivalent to $\mathbf{E}_\mathbf{Q}[V_T g] = 0$ for all $g \in K(X)$. This happens if and only if the probability measure $\mathbf{Q}'$, defined as $d\mathbf{Q}' = V_T d\mathbf{Q}$, is in $\mathcal{M}^e(X)$. (Note that by the martingale property we have $\mathbf{E}_\mathbf{Q}[V_T] = V_0 = 1$.) □

Remark 2.5.5. The process $(V_t)_{t=0}^T$ is a $\mathbf{Q}$-martingale for every $\mathbf{Q} \in \mathcal{M}^e(S)$. Now if $d\mathbf{Q}' = V_T d\mathbf{Q}$, then we have the following so-called Bayes' rule for $f \in L^\infty(\Omega, \mathcal{F}, \mathbf{P})$:

$$\begin{aligned} \mathbf{E}_{\mathbf{Q}'}\left[f \mid \mathcal{F}_t\right] &= \frac{\mathbf{E}_\mathbf{Q}\left[f V_T \mid \mathcal{F}_t\right]}{\mathbf{E}_\mathbf{Q}\left[V_T \mid \mathcal{F}_t\right]} = \frac{\mathbf{E}_\mathbf{Q}\left[f V_T \mid \mathcal{F}_t\right]}{V_t} \\ &= \mathbf{E}_\mathbf{Q}\left[f \frac{V_T}{V_t} \,\middle|\, \mathcal{F}_t\right]. \end{aligned}$$

The previous equality can also be written as

$$V_t \mathbf{E}_{Q'}[f \mid \mathcal{F}_t] = \mathbf{E}_\mathbf{Q}[f V_T \mid \mathcal{F}_t].$$

From this it follows that $(Z_t)_{t=0}^T$ is a $\mathbf{Q}'$-martingale if and only if $(Z_t V_t)_{t=0}^T$ is a $\mathbf{Q}$-martingale. This statement can also be seen as the martingale formulation of Theorem 2.5.4 above.

Remark 2.5.6. The above theorem tells us that the no-arbitrage arguments do not depend on whether we do the accounting in Euros or Dollars. To phrase it more precisely: whether we denote prices in properly discounted Euros (by the risk-free Euro rate) or Dollars (discounted by the risk-free Dollar rate).

Let us illustrate this explicitly by considering $f \in L^\infty(\Omega, \mathcal{F}, \mathbf{P})$ which we view as a contingent claim maturing at time T, denoted in units of the original numéraire $S^0 \equiv 1$. By Theorem 2.4.2 the superreplication price (denoted in units of S_0^0) equals

$$\overline{\pi}(f) = \sup\{\mathbf{E}_{\mathbf{Q}}[f] \mid \mathbf{Q} \in \mathcal{M}^e(S)\}. \tag{2.11}$$

Passing to the numéraire V the above contingent claim pays $\frac{f}{V_T}$ units of the numéraire V_T at time T. Hence applying Theorem 2.4.2 to the process X we get

$$\overline{\pi}(f) = \left\{ \mathbf{E}_{\mathbf{Q}'}\left[\frac{f}{V_T}\right] \,\middle|\, \mathbf{Q}' \in \mathcal{M}^e(X) \right\},$$

which by the above Theorem 2.5.4 gives

$$\overline{\pi}(f) = \left\{ \mathbf{E}_{\mathbf{Q}'}\left[\frac{f}{V_T}\right] \,\middle|\, \frac{d\mathbf{Q}'}{d\mathbf{Q}} = V_T,\ \mathbf{Q} \in \mathcal{M}^e(S) \right\}.$$

This price clearly equals the superreplication price obtained in (2.11).

Hence the interval of arbitrage free prices for a contingent claim does not depend on the choice of numéraire.

2.6 Kramkov's Optional Decomposition Theorem

We now present a dynamic version of Theorem 2.4.2 (superreplication), due to D. Kramkov, who actually proved this theorem in a much more general version (see [K 96a], [FK 98], and Chap. 15 below). An earlier version of this theorem is due to N. El Karoui and M.-C. Quenez [EQ 95]. We refer to Chap. 15 for more detailed references.

Theorem 2.6.1 (Optional Decomposition). *Assume that S satisfies (NA) and let $V = (V_t)_{t=0}^T$ be an adapted process.*

The following assertions are equivalent:

(i) *V is a super-martingale for each $\mathbf{Q} \in \mathcal{M}^e(S)$.*
(i') *V is a super-martingale for each $\mathbf{Q} \in \mathcal{M}^a(S)$*
(ii) *V may be decomposed into $V = V_0 + H \cdot S - C$, where $H \in \mathcal{H}$ and $C = (C_t)_{t=0}^T$ is an increasing adapted process starting at $C_0 = 0$.*

Remark 2.6.2. To clarify the terminology *"optional decomposition"* let us compare this theorem with Doob's celebrated decomposition theorem for non-negative super-martingales $(V_t)_{t=0}^T$ (see, e.g., [P 90]): this theorem asserts that, for a non-negative (adapted, càdlàg) process V defined on a general filtered probability space we have the equivalence of the following two statements:

(i) V is a super-martingale (with respect to the fixed measure $\mathbf{P}$),
(ii) V may be decomposed in a unique way into $V = V_0 + M - C$, where M is a local martingale (with respect to $\mathbf{P}$) and C is an increasing predictable process s.t. $M_0 = C_0 = 0$.

We immediately recognise the similarity in spirit. However, there are significant differences. As to condition (i) the difference is that, in the setting of the optional decomposition theorem, the super-martingale property pertains to *all* martingale measures $\mathbf{Q}$ for the process S. As to condition (ii), the role of the local martingale M in Doob's theorem is taken by the stochastic integral $H \cdot S$.

A decisive difference between the two theorems is that in Theorem 2.6.1, the decomposition is no longer unique and one cannot choose, in general, C to be predictable. The process C can only be chosen to be optional, which in the present setting is the same as adapted.

The economic interpretation of the optional decomposition theorem reads as follows: a process of the form $V = V_0 + H \cdot S - C$ describes the wealth process of an economic agent. Starting at an initial wealth V_0, subsequently investing in the financial market according to the trading strategy H, and consuming as described by the process C where the random variable C_t models the accumulated consumption during the time period $\{1, \ldots, t\}$, the agent clearly obtains the wealth V_t at time t. The message of the optional decomposition theorem is that these wealth processes are characterised by condition (i) (or, equivalently, (i')).

Proof of Theorem 2.6.1. First assume that $T = 1$, i.e., we have a one-period model $S = (S_0, S_1)$. In this case the present theorem is just a reformulation of Theorem 2.4.2: if V is a super-martingale under each $\mathbf{Q} \in \mathcal{M}^e(S)$, then

$$\mathbf{E}_{\mathbf{Q}}[V_1 - V_0] \leq 0, \quad \text{for all } \mathbf{Q} \in \mathcal{M}^e(S).$$

Hence there is a predictable trading strategy H (i.e., an $\mathcal{F}_0$-measurable $\mathbb{R}^d$-valued function - in the present case $T = 1$) such that $(H \cdot S)_1 \geq V_1 - V_0$. Letting $C_0 = 0$ and writing $\Delta C_1 = C_1 = -V_1 + (V_0 + (H \cdot S)_1)$ we get the desired decomposition. This completes the construction for the case $T = 1$.

For general $T > 1$ we may apply, for each fixed $t \in \{1, \ldots, T\}$, the same argument as above to the one-period financial market (S_{t-1}, S_t) based on $(\Omega, \mathcal{F}, \mathbf{P})$ and adapted to the filtration $(\mathcal{F}_{t-1}, \mathcal{F}_t)$. We thus obtain an $\mathcal{F}_{t-1}$-measurable, $\mathbb{R}^d$-valued function H_t and a non-negative $\mathcal{F}_t$-measurable function ΔC_t such that

$$\Delta V_t = (H_t, \Delta S_t) - \Delta C_t,$$

where again $(. , .)$ denotes the inner product in $\mathbb{R}^d$. This will finish the construction of the optional decomposition: define the predictable process H as $(H_t)_{t=1}^T$ and the adapted increasing process C by $C_t = \sum_{u=1}^t \Delta C_u$. This proves the implication (i) $\Rightarrow$ (ii).

The implications (ii) $\Rightarrow$ (i') $\Rightarrow$ (i) are trivial. $\square$